Quantum Algorithms for Variants of Average-Case Lattice Problems via Filtering

Yilei Chen[1], Qipeng Liu[2], and Mark Zhandry[3]

[1] Tsinghua University, Beijing, China 100084
[2] Simons Institute for the Theory of Computing
[3] NTT Research and Princeton University
{chenyilei.ra,qipengliu,mzhandry}@gmail.com

Abstract. We show polynomial-time quantum algorithms for the following problems:

1. Short integer solution (SIS) problem under the *infinity* norm, where the public matrix is very wide, the modulus is a polynomially large prime, and the bound of infinity norm is set to be half of the modulus minus a constant.
2. Learning with errors (LWE) problem given LWE-like quantum states with polynomially large moduli and certain error distributions, including bounded uniform distributions and Laplace distributions.
3. Extrapolated dihedral coset problem (EDCP) with certain parameters.

The SIS, LWE, and EDCP problems in their standard forms are as hard as solving lattice problems in the worst case. However, the variants that we can solve are not in the parameter regimes known to be as hard as solving worst-case lattice problems. Still, no classical or quantum polynomial-time algorithms were known for the variants of SIS and LWE we consider. For EDCP, our quantum algorithm slightly extends the result of Ivanyos et al. (2018).

Our algorithms for variants of SIS and EDCP use the existing quantum reductions from those problems to LWE, or more precisely, to the problem of solving LWE given LWE-like quantum states. Our main contribution is solving LWE given LWE-like quantum states with interesting parameters using a filtering technique.

1 Introduction

Solving the shortest vector problem (SVP) over lattices has been a target for designing efficient quantum algorithms for decades. In the literature, solving approximate SVP for *all* lattices has been (classically or quantumly) reduced to the following problems:

1. The short integer solution (SIS) problem, classically, initially shown by Ajtai [Ajt96].
2. The dihedral coset problem (DCP), quantumly, initially shown by Regev [Reg02].

3. The learning with errors problem (LWE), quantumly, initially shown by Regev [Reg05].

Therefore, to show an efficient quantum algorithm for approximate SVP in the worst-case, it suffices to construct an efficient quantum algorithm for any one of those average-case problems. However, no polynomial (or even subexponential) time quantum algorithms are known for SIS or LWE. For DCP, a subexponential quantum algorithm is given by Kuperberg [Kup05]. But the quantum reduction shown by Regev [Reg02] requires the DCP algorithm to be noise-tolerant, while the algorithm of Kuperberg is not. Let us also mention that over the past few years, efficient quantum algorithms for SVP for ideal lattices in certain parameter regimes have been shown in [CGS14, EHKS14, BS16, CDPR16, CDW17]. Still, showing a polynomial (or even subexponential) time quantum algorithm for SVP with polynomial approximation factors for *all* lattices is widely open.

The SIS and LWE problems are powerful tools for building cryptosystems, thus understanding the quantum hardness of those two problems is interesting in its own right. The SIS problem is typically used in constructing elementary cryptographic primitives such as one-way functions [Ajt96], collision-resistant hash functions [GGH96] digital signatures [GPV08]. The LWE problem is extremely versatile, yielding public-key cryptosystems [Reg05], and advanced cryptographic capabilities such as fully homomorphic encryption (FHE) [BV11], attribute-based encryption [GVW13], and quantum FHE [Mah18]. The conjectured quantum hardness of SIS and LWE has also made lattice-based cryptosystems popular candidates for post-quantum cryptography standardization [DKRV18, BDK+18, DKL+18].

1.1 Background of SIS, LWE, DCP, and our main results

We show polynomial-time quantum algorithms for certain variants of SIS, LWE, and DCP. Our quantum algorithms for the variants of SIS and DCP go through the existing quantum reductions from those problems to LWE, or more precisely, to the problems of *Constructing quantum LWE states* ($\mathsf{C}|\mathsf{LWE}\rangle$) and *Solving LWE given LWE-like quantum states* ($\mathsf{S}|\mathsf{LWE}\rangle$). In fact, the heart of our results is showing a quantum filtering technique for solving those quantum versions of LWE.

Let us now provide more background of SIS, LWE, and DCP, then state our main results.

SIS Let us first recall the standard definition of the SIS problem.

Definition 1 (Short integer solution (SIS) problem [Ajt96]). *Let n, m, q be integers such that $m = \Omega(n \log q) \subseteq \mathsf{poly}(n)$. Let β be a positive real number such that $\beta < q$. Let A be a uniformly random matrix over $\mathbb{Z}_q^{n \times m}$. The SIS problem $\mathsf{SIS}_{n,m,q,\beta}$ asks to find a nonzero vector $x \in \mathbb{Z}^m$ such that $\|x\|_2 \leq \beta$ and $Ax \equiv 0 \pmod q$.*

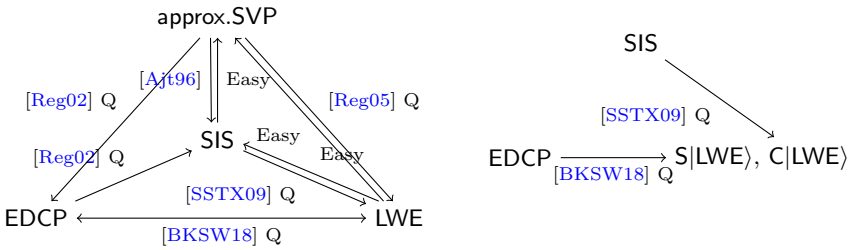

Fig. 1. Left: An overview of the reductions between SVP, SIS, EDCP, and LWE. "$A \to B$" means Problem A reduces to Problem B. "Q" means quantum. Right: The reductions used in our quantum algorithms.

The SIS problem is shown to be as hard as solving approximate SVP for all lattices [Ajt96]. The reductions are improved via a series of works [CN97, Mic02, MR07, GPV08, MP13]. Several variants of the SIS problem are studied in the literature. The most common variant is the one that changes the restriction of the solution. The solution is bounded in ℓ_p norm for some $p \geq 0$, or even the ℓ_∞ norm, instead of bounded in ℓ_2 norm. In this paper, we look at the variant where the solution is bounded by its ℓ_∞ norm. More precisely, we use $\mathsf{SIS}^\infty_{n,m,q,\beta}$ to denote the variant of SIS where the solution x is required to satisfy $\|x\|_\infty \leq \beta$. When $\beta = 1$, it corresponds to the subset-sum problem where the solution is bounded in $\{-1, 0, 1\}$.

Bounding the SIS solution in its ℓ_∞ norm is used quite commonly in cryptography due to its simplicity (it is used, e.g., in [BV15]). When the parameters are set so that $\beta\sqrt{m} > q$, i.e., when m is relatively large compared to q/β, we are not aware of any worst-case problem that is reducible to $\mathsf{SIS}^\infty_{n,m,q,\beta}$. Still, such parameter settings are used in cryptosystems. In a recent practical signature scheme proposed by Ducas et al. [DKL+18], the security of the scheme relies on (the "Module" version of) $\mathsf{SIS}^\infty_{n,m,q,\beta}$ with $\beta\sqrt{m} > q$. In their security analysis, the authors mention that the problem of SIS^∞ by itself has not been studied in-depth. Most of the algorithms they can think of for SIS^∞ are the ones designed for solving SIS or SVP in the ℓ_2 norm, such as BKZ [SE94].

To date, the only algorithm we are aware of that takes advantage of the ℓ_∞-norm bound has the following features. It solves $\mathsf{SIS}^\infty_{n,m,q,\beta}$ with a highly composite q and a very large m. For example, it is a polynomial-time algorithm for $\mathsf{SIS}^\infty_{n,O(n^c),2^c,1}$ when c is a constant. The algorithm is classical, folklore, and we include a formal description of the algorithm in the full version. It was not clear how to solve $\mathsf{SIS}^\infty_{n,m,q,\beta}$ when q is a polynomial prime and β is just slightly smaller than $q/2$, even if m is allowed to be an arbitrary polynomial.

We show a polynomial-time quantum algorithm for $\mathsf{SIS}^\infty_{n,m,q,\beta}$ where q is a polynomial prime modulus, $\beta = \frac{q-c}{2}$ for some constant c, and m is a large polynomial.

Theorem 1. *Let $c > 0$ be a constant integer, $q > c$ be a polynomially large prime modulus. Let $m \in \Omega\left((q-c)^3 \cdot n^{c+1} \cdot q \cdot \log q\right) \subseteq \mathsf{poly}(n)$, there is a polynomial-time quantum algorithm that solves $\mathsf{SIS}^{\infty}_{n,m,q,\frac{q-c}{2}}$.*

Remark 1. Note that if $\beta = q/2$, then a solution can be found classically by simply solving $Ax \equiv 0 \pmod q$ over $\mathbb{Z}_q$ using Gaussian elimination. Then for each entry in x, pick the representative over $\mathbb{Z}$ that lies in the range $[-q/2, q/2)$. This classical algorithm also extends to $\beta = \frac{q-c}{2}$ when $q = \Omega(n)$. In particular, as long as all the entries of x are at least $c/2$ far from $q/2$, x will be a valid solution. In the regime $q = \Omega(n)$, a random solution to $Ax \equiv 0 \pmod q$ will satisfy this with probability at least $O((1-c/n)^n) = O(e^{-c})$, a constant. Theorem 1 thus gives a non-trivial algorithm for $\mathsf{SIS}^{\infty}_{n,m,q,\frac{q-c}{2}}$ when $q \in o(n)$, for which (to the best of our knowledge), no prior classical or quantum algorithm was known.

Remark 2. Our algorithm can also solve a variant of SIS where the each entry of the solution is required to be in an arbitrary subset S of $\mathbb{Z}_q$ such that $q - |S| = c$, where c is a constant (instead of the subset $[-\beta, \beta] \cap \mathbb{Z}$ of $\mathbb{Z}_q$). The width of the A matrix is required to satisfy $m \in \Omega\left((q-c)^3 \cdot n^{c+1} \cdot q \cdot \log q\right) \subseteq \mathsf{poly}(n)$. For example, suppose $q = 3$ and $m \in \Omega(n^2)$, our algorithm is able to provide a $\{0,1\}$-solution for SIS.

Let us remark that our algorithm does not improve upon the existing algorithms for breaking the signature scheme in [DKL+18] since we require m to be very large, while the m used in [DKL+18] is fairly small.

LWE Let us first recall the classical definition of the LWE problem.

Definition 2 (Learning with errors (LWE) [Reg05]). *Let n, m, q be positive integers. Let $u \in \mathbb{Z}_q^n$ be a secret vector. The learning with errors problem $\mathsf{LWE}_{n,m,q,\mathcal{D}_{\mathsf{noise}}}$ asks to find the secret vector u given access to an oracle that outputs a_i, $a_i \cdot u + e_i \pmod q$ on its i^{th} query, for $i = 1, ..., m$. Here each a_i is a uniformly random vector in $\mathbb{Z}_q^n$, and each error term e_i is sampled from a distribution $\mathcal{D}_{\mathsf{noise}}$ over $\mathbb{Z}_q$.*

Regev [Reg05] shows if there is a polynomial-time algorithm that solves $\mathsf{LWE}_{n,m,q,\mathcal{D}_{\mathsf{noise}}}$ where $\mathcal{D}_{\mathsf{noise}}$ is Gaussian and m can be an arbitrary polynomial, then there is a quantum algorithm that solves worst-case approximate SVP. Note that in Regev's definition, the LWE samples are completely classical. In the variants of LWE we consider, the error distribution appears in the amplitude of some quantum states. Those quantum variants of LWE were implicitly used in the quantum reductions in [SSTX09, BKSW18], but they have not been made formal. Looking ahead, our new quantum algorithms make explicit use of the quantum nature of the noise distribution.

Our quantum algorithm for SIS^{∞} adapts the quantum reduction from SIS to the problem of *constructing LWE states* implicitly used in [SSTX09].

Definition 3. *Let n, m, q be positive integers. Let f be a function from $\mathbb{Z}_q$ to $\mathbb{R}$. The problem of constructing LWE states $\mathsf{C|LWE}\rangle_{n,m,q,f}$ asks to construct a quantum state of the form $\sum_{u\in\mathbb{Z}_q^n}\bigotimes_{i=1}^m\left(\sum_{e_i\in\mathbb{Z}_q} f(e_i)|a_i\cdot u+e_i \bmod q\rangle\right)$, given the input $\{a_i\}_{i=1,\ldots,m}$ where each a_i is a uniformly random vector in $\mathbb{Z}_q^n$.*

Our quantum algorithm for EDCP adapts the quantum reduction from EDCP to the problem of *solving LWE given LWE-like quantum states* implicitly used in [BKSW18].

Definition 4. *Let n, m, q be positive integers. Let f be a function from $\mathbb{Z}_q$ to $\mathbb{R}$. Let $u\in\mathbb{Z}_q^n$ be a secret vector. The problem of solving LWE given LWE-like states $\mathsf{S|LWE}\rangle_{n,m,q,f}$ asks to find u given access to an oracle that outputs independent samples a_i, $\sum_{e_i\in\mathbb{Z}_q} f(e_i)|a_i\cdot u+e_i \bmod q\rangle$ on its i^{th} query, for $i=1,...,m$. Here each a_i is a uniformly random vector in $\mathbb{Z}_q^n$.*

We would like to remark that in the problem $\mathsf{C|LWE}\rangle$, there is no secret vector u; the goal is to construct an equal superposition of all LWE states for all possible u. Whereas for the problem $\mathsf{S|LWE}\rangle$, the goal is to find the secret vector u given samples of LWE states for this particular secret vector.

Let us briefly discuss the relations among LWE, $\mathsf{S|LWE}\rangle$, and $\mathsf{C|LWE}\rangle$. If we set f as $\sqrt{\mathcal{D}_{\mathsf{noise}}}$, then an efficient algorithm for solving $\mathsf{LWE}_{n,m,q,\mathcal{D}_{\mathsf{noise}}}$ implies efficient algorithms for solving $\mathsf{C|LWE}\rangle_{n,m,q,f}$ and $\mathsf{S|LWE}\rangle_{n,m,q,f}$. However, solving $\mathsf{C|LWE}\rangle_{n,m,q,f}$ or $\mathsf{S|LWE}\rangle_{n,m,q,f}$ does not necessarily imply solving $\mathsf{LWE}_{n,m,q,\mathcal{D}_{\mathsf{noise}}}$ in general. An algorithm for solving $\mathsf{C|LWE}\rangle_{n,m,q,f}$ only implies an efficient algorithm for solving $\mathsf{LWE}_{n,m,q,\mathcal{D}_{\mathsf{noise}}}$ when m is small compared to the ratio of the "widths" of f and $\mathcal{D}_{\mathsf{noise}}$; we will explain in details in §1.4.

Let us also remark that the $\mathsf{C|LWE}\rangle$ and $\mathsf{S|LWE}\rangle$ problems we define are different from the problem of "LWE with quantum samples" defined in [GKZ19]. In their definition, the quantum LWE samples are of the form $\sum_{a\in\mathbb{Z}_q^n}|a\rangle|a\cdot u+e\rangle$, where the error e is classical and a is in the quantum state. This variant of quantum LWE is easy to solve [GKZ19], but the idea in the algorithm does not carry to the quantum LWE variants we are interested in.

In [Reg05] (followed by [SSTX09, BKSW18] and most of the papers that use LWE), the noise distribution $\mathcal{D}_{\mathsf{noise}}$ or f is chosen to be Gaussian. One of the nice features of a Gaussian function f is that both f and its discrete Fourier transform (DFT) (over $\mathbb{Z}_q$), defined as

$$\hat{f}:\mathbb{Z}_q\to\mathbb{C},\quad \hat{f}:y\mapsto\sum_{x\in\mathbb{Z}_q}\frac{1}{\sqrt{q}}\cdot e^{\frac{2\pi ixy}{q}}\cdot f(x),$$

are negligible beyond their centers. Such a feature of $\hat{f}$ is crucial in establishing the quantum reductions among lattice problems in [Reg05, SSTX09, BKSW18].

Other choices of noise distribution are also used for LWE in the literature. One popular option is to let f be the bounded uniform distribution over $[-B,B]$ for some $0<B<\frac{q}{2}$. For certain choices of n,m,q,B, (classical) LWE with B-bounded uniform error is proven to be as hard as LWE with Gaussian

noise [DM13]. On the other hand, Arora and Ge [AG11] present a classical algorithm for breaking LWE with a prime modulus q when the support S of the LWE error distribution is very small. It requires $m \in \Omega(n^{|S|})$ and runs in time $\mathsf{poly}(n^{|S|})$. When $B \in \omega(1)$ and q is a prime, no polynomial-time quantum algorithm has been published for LWE, $\mathsf{C|LWE}\rangle$, or $\mathsf{S|LWE}\rangle$.

We show when the noise distribution f is chosen such that $\hat{f}$ is *non-negligible* over $\mathbb{Z}_q$, then we can solve both $\mathsf{C|LWE}\rangle$ and $\mathsf{S|LWE}\rangle$ in quantum polynomial-time.

Theorem 2. *Let $n \in \mathbb{N}$ and $q \in \mathsf{poly}(n)$. Let $f : \mathbb{Z}_q \to \mathbb{R}$ be the amplitude for the error state such that the state $\sum_{e\in\mathbb{Z}_q} f(e)|e\rangle$ is efficiently constructible and $\eta := \min_{y\in\mathbb{Z}_q} |\hat{f}(y)|$ is non-negligible. Let $m \in \Omega\left(n \cdot q/\eta^2\right) \subseteq \mathsf{poly}(n)$, there exist polynomial-time quantum algorithms that solve $\mathsf{C|LWE}\rangle_{n,m,q,f}$ and $\mathsf{S|LWE}\rangle_{n,m,q,f}$.*

Although the theorem does not cover the case where f is Gaussian, it does cover some interesting error distributions f, such as when f is super-Gaussian (i.e., when $f(x) = e^{-|x/B|^p}$, for $0 < p < 2$, $0 < B < q$). It also covers the case where f is the bounded uniform distribution. The following is a corollary of Theorem 2 given that the DFT of bounded uniform distribution is non-negligible over $\mathbb{Z}_q$.

Corollary 1. *Let $n \in \mathbb{N}$ and $q \in \mathsf{poly}(n)$. Let $B \in \mathbb{Z}$ such that $0 < 2B + 1 < q$ and $\gcd(2B+1, q) = 1$. Let $f : \mathbb{Z}_q \to \mathbb{R}$ be $f(x) := 1/\sqrt{2B+1}$ when $x \in [-B, B] \cap \mathbb{Z}$ and 0 elsewhere. Let $m \in \Omega\left(n \cdot q^4 \cdot (2B+1)\right) \subseteq \mathsf{poly}(n)$, there exist polynomial-time quantum algorithms that solve $\mathsf{C|LWE}\rangle_{n,m,q,f}$ and $\mathsf{S|LWE}\rangle_{n,m,q,f}$.*

Our quantum algorithms for SIS^∞ and EDCP (i.e., Theorem 1 and Theorem 5) are obtained from the following variant of Theorem 2, where the noise amplitude for the quantum LWE problems is set to be the DFT of the bounded uniform distribution.

Theorem 3. *Let q be a polynomially large prime modulus. Let $B \in \mathbb{Z}$ be such that $q - (2B+1) = c$ is a constant. Let $f : Z_q \to \mathbb{R}$ be the bounded uniform distribution over $[-B, B] \cap \mathbb{Z}$. Let $m \in \Omega\left((q-c)^3 \cdot n^{c+1} \cdot q \cdot \log q\right) \subseteq \mathsf{poly}(n)$. There exist polynomial-time quantum algorithms that solve $\mathsf{C|LWE}\rangle_{n,m,q,\hat{f}}$ and $\mathsf{S|LWE}\rangle_{n,m,q,\hat{f}}$.*

DCP Let us introduce the variant of DCP defined by Brakerski et al. [BKSW18].

Definition 5 (Extrapolated Dihedral Coset Problem (EDCP)). *Let $n \in \mathbb{N}$ be the dimension, $q \geq 2$ be the modulus, and a function $D : \mathbb{Z}_q \to \mathbb{R}$, consists of m input states of the form*

$$\sum_{j\in\mathbb{Z}_q} D(j)|j\rangle|x + j \cdot s\rangle,$$

where $x \in \mathbb{Z}_q^n$ *is arbitrary and* $s \in \mathbb{Z}_q^n$ *is fixed for all* m *states. We say that an algorithm solves* $\mathsf{EDCP}_{n,m,q,D}$ *if it outputs* s *with probability* $\mathsf{poly}(1/(n \log q))$ *in time* $\mathsf{poly}(n \log q)$.

In this paper we are interested in the parameter setting where n is the security parameter and $q \in \mathsf{poly}(n)$. Although not strictly needed in this paper, let us briefly recall how the variants of the dihedral coset problem evolve. The original dihedral coset problem is a special case of EDCP where $n = 1$, q is exponentially large, and D is the uniform distribution over $\{0, 1\}$. Solving DCP implies solving the dihedral hidden subgroup problem. The two-point problem defined by Regev [Reg02] is another special case of EDCP where D is the uniform distribution over $\{0, 1\}$, and n is the security parameter. It was used as an intermediate step for establishing the reduction from approximate SVP to DCP. When the distribution D is non-zero beyond $\{0, 1\}$, the EDCP problem does not necessarily correspond to any versions of the hidden subgroup problem. The reason that Brakerski et al. [BKSW18] considers a distribution D supported beyond $\{0, 1\}$ is to establish a reduction from EDCP to LWE. Therefore, combining with the reduction from LWE to EDCP (by adapting Regev's reduction [Reg02]), they show that EDCP, as a natural generalization of DCP, is equivalent to LWE.

Efficient quantum algorithms are known for variants of EDCP when the modulus q and the distribution D satisfy certain conditions [FIM+03, CvD07, IPS18]. Let us remark that EDCP with those parameter settings are not known to be as hard as worse-case SVP or LWE through the reductions of [Reg02] or [BKSW18].

In this paper we show polynomial-time quantum algorithms that solve EDCP with the following parameter settings.

Theorem 4. *Let* $n \in \mathbb{N}$ *and* $q \in \mathsf{poly}(n)$. *Let* $f : \mathbb{Z}_q \to \mathbb{R}$ *be such that the state* $\sum_{e \in \mathbb{Z}_q} f(e)|e\rangle$ *is efficiently constructible and* $\eta := \min_{z \in \mathbb{Z}_q} |\hat{f}(z)|$ *is non-negligible. Let* $m \in \Omega\left(n \cdot q/\eta^2\right) \subseteq \mathsf{poly}(n)$. *There is a polynomial time quantum algorithm that solves* $\mathsf{EDCP}_{n,m,q,\hat{f}}$

Theorem 5. *Let* $n \in \mathbb{N}$ *and* $q \in \mathsf{poly}(n)$. *Let* c *be a constant integer such that* $0 < c < q$. *Let* $m \in \Omega\left((q-c)^3 \cdot n^{c+1} \cdot q \cdot \log q\right) \subseteq \mathsf{poly}(n)$, *there is a quantum algorithm running in time* $\mathsf{poly}(n)$ *that solves* $\mathsf{EDCP}_{n,m,q,D}$ *where* D *is the uniform distribution over* $[0, q-c) \cap \mathbb{Z}$.

We note that EDCP with the parameters in Theorem 5 has already been solved in the work of Ivanyos et al. [IPS18] by a quantum algorithm with similar complexity. The parameters in Theorem 4 are not covered by the result in [IPS18], but the implication of such a parameter setting is unclear. Nevertheless, we include our result to demonstrate the wide applicability of our techniques. We will compare our algorithm with the one in [IPS18] in §1.3.

1.2 Solving the quantum versions of LWE via filtering

As mentioned, our main technical contribution is to solve $\mathsf{S}|\mathsf{LWE}\rangle$ and $\mathsf{C}|\mathsf{LWE}\rangle$ (the quantum versions of LWE we define) with interesting parameters using a

filtering technique. Let us first explain the basic idea of filtering, then extend it to the general case.

The basic idea of filtering. To illustrate the basic idea of filtering, let us focus on how to use it to solve $\mathsf{S|LWE}\rangle$, namely, learning the secret $u \in \mathbb{Z}_q^n$ given a uniformly random matrix $A \in \mathbb{Z}_q^{n\times m}$ and the following state:

$$|\phi_u\rangle := \bigotimes_{i=1}^{m} \sum_{e_i\in\mathbb{Z}_q} f(e_i)|\ (u^T A)_i + e_i \pmod q\ \rangle. \tag{1}$$

Let us remark that an efficient quantum algorithm for $\mathsf{S|LWE}\rangle_{n,m,q,f}$ does not necessarily imply an efficient quantum algorithm for $\mathsf{C|LWE}\rangle_{n,m,q,f}$, since the quantum algorithm for $\mathsf{S|LWE}\rangle_{n,m,q,f}$ may, for example, destroy the input state. However, the quantum algorithm we show for $\mathsf{S|LWE}\rangle_{n,m,q,f}$ directly works for $\mathsf{C|LWE}\rangle_{n,m,q,f}$, so we focus on $\mathsf{S|LWE}\rangle_{n,m,q,f}$.

Let us assume m can be an arbitrary polynomial of n, q is a constant prime. The readers can think of f as any distribution. For readers who would like to have a concrete example, you can think of f as the QFT of bounded uniform distribution, i.e., let $g(z) := 1/\sqrt{2\beta+1}$ for $z \in [-\beta, \beta] \cap \mathbb{Z}$ and 0 elsewhere, then set $f := \hat{g}$ (f is then the discrete sinc function, but in the analysis we will not use the expression of f at all, we will only use g). By solving $\mathsf{S|LWE}\rangle_{n,m,q,f}$ and $\mathsf{C|LWE}\rangle_{n,m,q,f}$ with such a choice of f, we can get a polynomial quantum algorithm for $\mathsf{SIS}^{\infty}_{n,m,q,\beta}$ with a constant prime q and any $\beta \in [1, q/2)$, which was not known before. All the details of the analysis will be given in §3. Here let us explain the basic idea of filtering using this example.

Let us define

$$\text{for } v \in \mathbb{Z}_q,\ |\psi_v\rangle := \sum_{e\in\mathbb{Z}_q} f(e)|(v+e) \pmod q\rangle.$$

Therefore the input state in Eqn (1) can also be written as

$$|\phi_u\rangle = \bigotimes_{i=1}^{m} |\psi_{(u^T A)_i}\rangle.$$

To learn u from $|\phi_u\rangle$, our algorithm proceeds in two stages: first we look at each coordinate $|\psi_{(u^T A)_i}\rangle$ for $i = 1, ..., m$ separately, with the goal of learning some classical information about each coordinate of $u^T A$. We then continue with a classical step, which uses the information obtained about each coordinate of $u^T A$ to learn u.

Warm-up 1: Orthogonal states. Suppose the vectors in the set $\{|\psi_v\rangle\}_{v\in\mathbb{Z}_q}$ were orthogonal. Then we could define a unitary U such that $U|\psi_v\rangle = |v\rangle$. We could then apply this unitary component-wise to $|\phi_u\rangle$ and measure the results in the computational basis, learning $u^T A$. Gaussian elimination then recovers u.

Warm-up 2: Filtering out a single value. Unfortunately, the $|\psi_v\rangle$ will typically not be orthogonal, so such a unitary as above will not exist. This means we cannot learn v with certainty from $|\psi_v\rangle$.

Nevertheless, we can learn *some* information about v from $|\psi_v\rangle$. Concretely, pick some value $y \in \mathbb{Z}_q$, and consider an arbitrary unitary U_y such that

$$U_y|\psi_y\rangle = |0\rangle.$$

Now imagine applying U_y to $|\psi_v\rangle$, and measuring in the computational basis. If $v = y$, then the measurement will always give 0. Unfortunately, since the $|\psi_v\rangle$ are not orthogonal, measuring $U_y|\psi_v\rangle$ for $v \neq y$ may also give 0. Therefore, while a 0 outcome gives us some prior on the value of v, it does not let us conclude anything for certain.

On the other hand, if a measurement gives a *non-zero* value, then we know for certain that $v \neq y$. This is the basic idea of our filtering approach: we filter out the case where $v = y$, learning an inequality constraint on v. This can be seen as a weak form of unambiguous state discrimination [Per88], where the measurement either gives unambiguous information about the unknowns or is thrown away. It turns out that, in some parameter regimes, learning such non-equality constraints will let us compute u.

Concretely, given an unknown state $|\phi_u\rangle$, we choose an independent random y_i for each coordinate, apply the unitary U_{y_i} to the ith coordinate, and measure. Any measurement result that gives 0, we throw away; for typical $|\psi_v\rangle$, few measurements will give 0. The remaining results yield inequality constraints of the form $(u^T A)_i \neq y_i$. We then apply the classical Arora-Ge algorithm [AG11] to these constraints. This algorithm works by viewing the inequality constraints as degree $q-1$ constraints and then re-linearizing them. This process converts the inequality constraints into equality constraints, but at the cost of blowing up the number of unknowns to $\approx n^{q-1}$. In the regime where q is a constant and m is a sufficiently large polynomial, the system can be solved in polynomial-time using Gaussian elimination.

Our algorithm: filtering out multiple values. Our algorithm so far is limited to filtering out a single value, which in turn limits us to a constant q, due to our use of Arora-Ge.

In order to get a polynomial-time algorithm for larger q, we must filter out more points; in fact, in order to use Arora-Ge, we need our constraints to have constant degree, which in turn means we must filter out all but a constant number of elements of $\mathbb{Z}_q$. Filtering out so many points requires care.

Consider the goal of filtering out two values. If there exists, for $y_0, y_1 \in \mathbb{Z}_q$, a unitary U_{y_0,y_1} such that

$$U_{y_0,y_1}|\psi_{y_b}\rangle = |b\rangle,$$

then we could apply U_{y_0,y_1} and measure in the computational basis. If the result is not equal to 0 or 1, then we know that $v \notin \{y_0, y_1\}$, thus filtering out two values.

In general such a unitary does not exist, as it would require $|\psi_{y_0}\rangle$ and $|\psi_{y_1}\rangle$ to be orthogonal. Instead what we do is to define a unitary U_{y_0,y_1} such that

$$U_{y_0,y_1}|\psi_{y_b}\rangle \in \mathrm{Span}(|0\rangle, |1\rangle) \ .$$

This method also naturally extends to filtering a larger number of y values. The limitation is that, as the number of y increases, the probability of getting a successful measurement (where "success" means, e.g. getting a result other than 0,1) decreases. For example, suppose the $|\psi_v\rangle$ all lie in the space of dimension $d \ll q$. Then after excluding d values, the probability of a successful measurement will be 0. Even if the vectors are technically linear independent but close to a d-dimensional subspace, the probability will be non-zero but negligible. This, for example, rules out an algorithm for the case where f is discrete Gaussian.

Therefore, whether or not the algorithm will succeed depends crucially on the "shape" of the states $|\psi_v\rangle$, and in particular, the distribution f. Our applications roughly follow the outline above, analyzing specific cases of $|\psi_v\rangle$. Our main observation is that, since all the vectors $|\psi_v\rangle$ are just shifts of a single fixed vector, we can construct a unitary operator by taking the normalized Gram-Schmidt orthogonalization of a circulant matrix M_f, defined by

$$M_f := [|\psi_v\rangle, |\psi_{v+1}\rangle, ..., |\psi_{v+q-1}\rangle]. \tag{2}$$

This allows us to relate the success probability of filtering out $q-1$ values to the length of the last Gram-Schmidt vector of M_f (before normalization). The length of the last Gram-Schmidt vector is related to the eigenvalues of the circulant matrix M_f, and it can be bounded in terms of the discrete Fourier transform $\hat{f}$ of f. Our calculation suggests that when $\hat{f}$ is *non-negligible* over all the values in $\mathbb{Z}_q$, the success probability of correctly guessing each coordinate is non-negligible. Therefore when m is a sufficiently large polynomial, we get a polynomial-time algorithm for $\mathsf{S|LWE}\rangle_{n,m,q,f}$. In Figure 1.2 we give four examples of error amplitudes. It shows that if the minimum of $\hat{f}$ is non-negligible, then the length of the last Gram-Schmidt of M_f is non-negligible.

1.3 The related work of Ivanyos et al.

Let us briefly compare our paper with the work of Ivanyos et al. [IPS18][1]. As mentioned in §1.1, EDCP with the parameters in Theorem 5 has already been solved in [IPS18] by a quantum algorithm with similar complexity. While we solve EDCP using the quantum reduction from EDCP to $\mathsf{S|LWE}\rangle$ with sinc error distribution, Ivanyos et al. used a reduction from EDCP to a problem called "learning from disequations" (LSF), defined as follows: the goal is to learn a secret $s \in Z_q^n$ by querying an oracle which outputs some $a \in \mathbb{Z}_q^n$ such that $\langle a, s\rangle \in A$, where A is a known subset of $\mathbb{Z}_q$. Given the set A and $m \in n^{O(|A|)}$

[1] In the initial version of our paper (August 25, 2021) we were not aware of the results in [IPS18]. We sincerely thank Gábor Ivanyos for telling us the results in [IPS18].

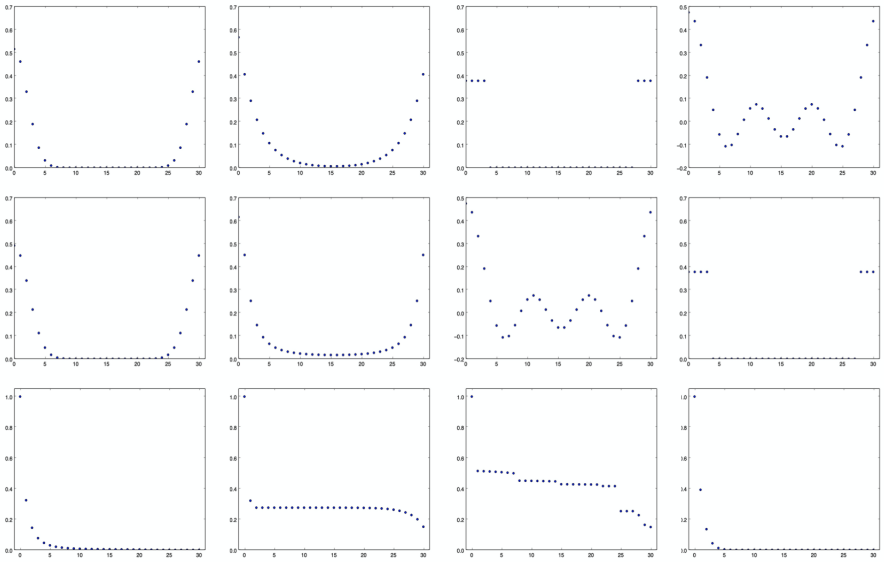

Fig. 2. Examples of error amplitude f (top), its DFT $\hat{f}$ (middle), and the length of the i^{th} Gram-Schmidt vector of M_f in Eqn. (2) for $0 \leq i < q$ (bottom). Let $q = 31$ for all examples. The error amplitude f is (from left to right): (1) Gaussian: $f(x) = \exp(-(x/3)^2)$; (2) Laplacian: $f(x) = \exp(-|x/3|)$; (3) Uniform over $[-3, 3] \cap \mathbb{Z}$; (4) The DFT of Uniform over $[-3, 3] \cap \mathbb{Z}$.

samples $a_1, ..., a_m$, they solve LSF in time $n^{O(|A|)}$ classically (using the Arora-Ge algorithm). This means when $|A|$ is a constant, the problem of LSF is solvable in $\mathsf{poly}(n)$ time.

In their algorithm they also used an idea similar to what we called "filtering". While we use filtering to solve $\mathsf{S}|\mathsf{LWE}\rangle$, they used the idea of filtering in the quantum reduction from EDCP to the LSF problem.

Overall, both papers use the idea of filtering to solve EDCP for the parameters settings in Theorem 5, but the intermediate problems we reduced to are different. It appears that solving $\mathsf{S}|\mathsf{LWE}\rangle$ allows us to obtain a richer variety of algorithms. In particular, it allows us to obtain a quantum algorithm for SIS^{∞}, which was not achieved in [IPS18]. Furthermore, our results give evidence that the $\mathsf{S}|\mathsf{LWE}\rangle$ and $\mathsf{C}|\mathsf{LWE}\rangle$ problems are quantumly easier to solve than the classical LWE problem, which shows another hope of solving the worst-case lattice problems. Let us elaborate on this point in the next section.

1.4 Future directions

Our results show polynomial time quantum algorithms for variants of average-case lattice problems. They do not appear to affect the security of any lattice-based cryptosystems in use. One may ask how far are we from solving standard

LWE or approximate SVP for all lattices? Here we discuss two potential approaches of extending our results towards those ultimate goals.

Our first observation is that in order to solve standard LWE, "all" what we need to do is to solve $\mathsf{C}|\mathsf{LWE}\rangle_{n,m,q,f}$ with a smaller m than what we have achieved in Theorem 2 or Corollary 1. For the simplicity of explanation, assume the parameters σ, B, q satisfy $\sigma < B \ll q \in \mathsf{poly}(n)$. To solve decisional $\mathsf{LWE}_{n,m,q,D}$ where the noise distribution D is uniform over $[-\sigma, \sigma] \cap \mathbb{Z}$, it suffices to solve $\mathsf{C}|\mathsf{LWE}\rangle_{n,m,q,f}$ where f is the uniform distribution over $[-B, B] \cap \mathbb{Z}$, and with $m \leq B/\sigma$. Currently, using our result in Corollary 1, we need $m \in \Omega\left(n \cdot q^4 \cdot (2B+1)\right)$, which is polynomial in n but way larger than B/σ.

The algorithm of breaking decisional LWE via solving $\mathsf{C}|\mathsf{LWE}\rangle$ is well-known and was implicitly used in the attempt of designing quantum algorithms for lattice problems in [ES16]. Let the decisional $\mathsf{LWE}_{n,m,q,D}$ instance be $(A \in \mathbb{Z}_q^{n \times m}, y \in \mathbb{Z}_q^m)$ where y is either an LWE sample or uniformly random. We solve $\mathsf{C}|\mathsf{LWE}\rangle_{n,m,q,f}$, i.e., construct a state

$$|\rho\rangle := \sum_{u \in \mathbb{Z}_q^n} \bigotimes_{i=1}^{m} \left(\sum_{e_i \in \mathbb{Z}_q} f(e_i) |a_i \cdot u + e_i \pmod q \rangle \right).$$

Let U_y denote a unitary operator that maps any $x \in \mathbb{Z}_q^m$ to $x + y$. Then we compute $\langle \rho | U_y | \rho \rangle$ by performing a Hadamard test. If y is an LWE sample, we expect the overlap between $|\rho\rangle$ and $|\rho + y\rangle$ to be at least $(1 - \sigma/B)^m$, whereas if y is uniform, we expect the overlap to be 0. Therefore, if we are able to solve $\mathsf{C}|\mathsf{LWE}\rangle_{n,m,q,f}$ with $m \leq B/\sigma$, then we can solve decisional $\mathsf{LWE}_{n,m,q,D}$. The distributions f and D in the example can be changed to other ones, but all of them require m to be relatively small in order to break standard LWE.

If we are not able to decrease the number of samples in our solutions of $\mathsf{C}|\mathsf{LWE}\rangle_{n,m,q,f}$ or $\mathsf{S}|\mathsf{LWE}\rangle_{n,m,q,f}$, another hope of solving worst-case approximate SVP is to modify Regev's reduction [Reg05]. Recall that Regev reduces worst-case approximate SVP to LWE with Gaussian noise and *arbitrarily polynomially* many classical samples. Suppose we can replace LWE with classical samples by its quantum variants $\mathsf{C}|\mathsf{LWE}\rangle$ or $\mathsf{S}|\mathsf{LWE}\rangle$, and replace Gaussian distribution by distributions with non-negligible DFT (like bounded uniform or Laplace distributions). Then approximate SVP can be solved using Theorem 2 without decreasing the number of samples. However, it is not clear to us whether modifying Regev's reduction is feasible or not.

2 Preliminaries

Notation and terminology. Let $\mathbb{R}, \mathbb{Z}, \mathbb{N}$ be the set of real numbers, integers and positive integers. For $q \in \mathbb{N}_{\geq 2}$, denote $\mathbb{Z}/q\mathbb{Z}$ by $\mathbb{Z}_q$. For $n \in \mathbb{N}$, $[n] := \{1, ..., n\}$. When a variable v is drawn uniformly at random from the set S, we denote by $v \leftarrow U(S)$.

A vector in $\mathbb{R}^n$ is represented in column form by default. For a vector v, the i^{th} component of v will be denoted by v_i. For a matrix A, the i^{th} column vector

of A is denoted a_i. We use A^T to denote the transpose of A, A^H to denote the conjugate transpose of A. The length of a vector is the ℓ_p-norm $\|v\|_p := (\sum v_i^p)^{1/p}$, or the infinity norm given by its largest entry $\|v\|_\infty := \max_i\{|v_i|\}$. The length of a matrix is the norm of its longest column: $\|A\|_p := \max_i \|a_i\|_p$. By default, we use ℓ_2-norm unless explicitly mentioned.

2.1 Quantum background

We assume the readers are familiar with the basic concepts of quantum computation. All the background we need in this paper is available in standard textbooks of quantum computation, e.g., [NC16]. When writing a quantum state such as $\sum_{x\in S} f(x)|x\rangle$, we typically omit the normalization factor except when needed. When a state can be approximately constructed within a negligible distance, we sometimes say the state is constructible and not mention the negligible distance.

Efficiently constructible unitary operators. In this paper we will use the fact that all the unitary matrices of polynomial dimension can be efficiently approximated within exponentially small distance.

Proposition 1 (Page 191 of [NC16]). *Any unitary matrix U on an n-qubit system can be written as a product of at most $2^{n-1}(2^n-1)$ two-level unitary matrices.*

Then, using Solovay-Kitaev Theorem, all the unitary matrices of $\mathsf{poly}(n)$ dimensions (therefore, applied on $O(\log n)$ qubits) can be approximated by $2^{O(\log n)} \in \mathsf{poly}(n)$ elementary quantum gates.

Proposition 2. *Let $\mathcal{G}$ denote set of unitary matrices that are universal for two-level gates. Given a unitary matrix $U \in \mathbb{C}^{d\times d}$, there is a classical algorithm that runs in time $\mathsf{poly}(d)$, outputs a sequence of two-level unitary matrices $U_1, ..., U_m \in \mathcal{G}$ such that $\prod_{i=1}^m U_i$ approximates U within distance negligible in d, and $m \in \mathsf{poly}(d)$.*

Looking ahead, the quantum algorithms in this work require quantum Fourier transform, superposition evaluations of classical circuits on quantum states and quantum gates that operate on $O(\log n)$ qubits. Thus, all quantum algorithms in the work can be efficiently approximated.

Quantum Fourier Transform. For any integer $q \geq 2$, let $\omega_q = e^{2\pi i/q}$ denote a primitive q-th root of unity. Define a unitary matrix $F_q \in \mathbb{C}^{q\times q}$ where $(F_q)_{i,j} := \frac{1}{\sqrt{q}} \cdot \omega_n^{ij}$, for $i, j \in \mathbb{Z}_q$.

Theorem 6 (QFT). *The unitary operator $\mathsf{QFT}_q := F_q$ can be implemented by $\mathsf{poly}(\log q)$ elementary quantum gates. When QFT_q is applied on a quantum state $|\phi\rangle := \sum_{x\in\mathbb{Z}_q} f(x)|x\rangle$, we have*

$$\mathsf{QFT}_q|\phi\rangle = \sum_{y\in\mathbb{Z}_q} \hat{f}(y)|y\rangle := \sum_{y\in\mathbb{Z}_q}\sum_{x\in\mathbb{Z}_q} \frac{1}{\sqrt{q}} \cdot \omega_q^{xy} \cdot f(x)|y\rangle.$$

2.2 Arora-Ge algorithm for solving LWE

We have defined the SIS, DCP, and LWE problems in the introduction. Here let us mention the Arora-Ge algorithm for solving LWE when the support of the error distribution is small. The following theorem is implicitly proven in [AG11, Section 3].

Theorem 7. *Let q be a prime, n be an integer. Let $\mathcal{D}_{\mathsf{noise}}$ be an error distribution which satisfies:*

1. *The support of $\mathcal{D}_{\mathsf{noise}}$ is of size $D < q$.*
2. $\Pr[e = 0,\ e \leftarrow \mathcal{D}_{\mathsf{noise}}] = \frac{1}{\delta}$ *for some* $\delta > 1$.

Then, let N be $\binom{n+D}{D}$ and C be a sufficiently large constant. Let $m := CN\delta q \log q$. There is a classical algorithm that solves $\mathsf{LWE}_{n,m,q,\mathcal{D}_{\mathsf{noise}}}$ in time $\mathsf{poly}(m)$ and succeeds with probability $1 - q^{-N}$.

Note that the probability is only taken over the randomness of samples. The algorithm is deterministic.

Suppose the error distribution $\mathcal{D}_{\mathsf{noise}}$ is known (which is always the case in our application). We can remove the second condition in Theorem 7 by shifting the error distribution such that the probability of getting 0 is maximized. More precisely, suppose $\mathcal{D}_{\mathsf{noise}}$ outputs some $e' \in \mathbb{Z}_q$ with the highest probability; we can always change an LWE sample (a_i, y_i) to $(a_i, y_i - e')$, and apply Arora-Ge on the shifted samples. Thus, we can shift the error distribution so that the probability of getting zero error is at least $1/q$. This transformation gives the following simple corollary.

Corollary 2. *Let q be a prime, n be an integer. Let $\mathcal{D}_{\mathsf{noise}}$ be an error distribution whose support is of size $D < q$ and known to the algorithm. Let $m := C \cdot n^D q^2 \log q$ where C is a sufficiently large constant. There is a classical algorithm that solves $\mathsf{LWE}_{n,m,q,\mathcal{D}_{\mathsf{noise}}}$ in time $\mathsf{poly}(m)$ and succeeds with probability $1 - O(q^{-n^D})$.*

3 The Idea of Filtering and a Mini Result for SIS^∞

In this section we give more details of the basic idea of *filtering*. Using the basic filtering technique, we obtain a polynomial-time quantum algorithm for $\mathsf{SIS}^\infty_{n,m,q,\beta}$ with q being a constant prime, m being as large as $\Omega(n^{q-1})$, and $1 \leq \beta < q/2$. Quantum polynomial-time algorithms for SIS^∞ with such parameter settings have not been given before.

Theorem 8. *Let n be an integer, q be a constant prime modulus. There is a quantum algorithm running in time $\mathsf{poly}(n)$ that solves $\mathsf{SIS}^\infty_{n,m,q,\beta}$ with $m \in \Omega\left(\frac{n^{q-1}q^2 \log q}{0.9 - 1/(2\beta+1)}\right) \subseteq \mathsf{poly}(n)$ and any $\beta \in \mathbb{Z}$ such that $1 \leq \beta < q/2$.*

Note that in the above theorem, $0.9 - 1/(2\beta+1)$ is at least 0.56 for $\beta \geq 1$. Thus, m is in the order of $n^{q-1}q^2 \log q$.

Let us first recall the existing quantum reduction from SIS to the problem of constructing certain LWE states presented implicitly in [SSTX09], then show the filtering technique and explain how to construct the required LWE states.

3.1 Recalling the quantum reduction from SIS to LWE

To give a quantum algorithm for solving $\mathsf{SIS}^{\infty}_{n,m,q,\beta}$ w.r.t. a uniformly random matrix $A \in \mathbb{Z}_q^{n\times m}$, it suffices to produce a state $\sum_{z\in([-\beta,\beta]\cap\mathbb{Z})^m \text{ s.t. } Az=0 \pmod q} |z\rangle$. As long as the set $([-\beta,\beta]\cap\mathbb{Z})^m$ contains a non-zero solution for $Az = 0 \pmod q$, we can solve $\mathsf{SIS}^{\infty}_{n,m,q,\beta}$ with probability $\geq 1/2$ by simply measuring the state.

The following is a quantum reduction from SIS to LWE where the distribution of z is general. Let $f : \mathbb{Z}_q \to \mathbb{R}$ be a function (in the example above, f is the uniform distribution over $[-\beta,\beta]\cap\mathbb{Z}$). We abuse the notation to let $f : \mathbb{Z}_q^m \to R$ be defined as $f(x) = \prod_{i=1}^m f(x_i)$ (we will clearly state the domain when using f).

Proposition 3. *To construct an* SIS state *of the form*

$$|\phi_{\mathsf{SIS}}\rangle := \sum_{z\in\mathbb{Z}_q^m \ s.t. \ Az=0 \pmod q} f(z)|z\rangle.$$

It suffices to construct an LWE state *of the following form*

$$|\phi_{\mathsf{LWE}}\rangle := \sum_{u\in\mathbb{Z}_q^n}\sum_{e\in\mathbb{Z}_q^m} \hat{f}(e)|u^T A + e^T \pmod q\rangle,$$

where $\hat{f}(e_i) = \sum_{x_i\in\mathbb{Z}_q} \frac{1}{\sqrt{q}}\cdot\omega_q^{e_i x_i} f(x_i)$, *for* $i = 1, ..., m$, *and* $\hat{f}(e) = \prod_{i=1,...,m} \hat{f}(e_i) = \sum_{x\in\mathbb{Z}_q^m} \frac{1}{\sqrt{q^m}} \cdot \omega_q^{\langle e,x\rangle} f(x)$.

Proof. $\mathsf{QFT}_q^m|\phi_{\mathsf{LWE}}\rangle = |\phi_{\mathsf{SIS}}\rangle$. □

The following lemma is immediate from Proposition 3.

Lemma 1. *Let* n, m, q *be any integers such that* $m \in \Omega(n\log q) \subseteq \mathsf{poly}(n)$. *Let* $0 < \beta < q/2$. *Let* f *be the uniform distribution over* $([-\beta,\beta]\cap\mathbb{Z})^m$. *Let* A *be a matrix in* $\mathbb{Z}_q^{n\times m}$. *If there is a polynomial-time quantum algorithm that generates a state negligibly close to* $\sum_{u\in\mathbb{Z}_q^n}\sum_{e\in\mathbb{Z}_q^m} \hat{f}(e)|u^T A + e\rangle$, *then there is a polynomial-time quantum algorithm that solves* $\mathsf{SIS}^{\infty}_{n,m,q,\beta}$ *for* A.

3.2 Constructing the LWE state via filtering

Now let us describe an algorithm for $\mathsf{C}|\mathsf{LWE}\rangle$.

1. The algorithm first prepares the following state:

$$\sum_{x\in\mathbb{Z}_q^m} f(x)|x\rangle \otimes \sum_{u\in\mathbb{Z}_q^n} |u\rangle,$$

   where we assume we work with a function f such that $\sum_{x\in\mathbb{Z}_q^m} f(x)|x\rangle$ can be efficiently generated. If so, then the whole state can be efficiently generated.
2. It then applies QFT_q^m on the x registers and gets:

$$\left(\mathsf{QFT}_q^m \sum_{x\in\mathbb{Z}_q^m} f(x)|x\rangle\right) \otimes \sum_{u\in\mathbb{Z}_q^n} |u\rangle = \left(\sum_{e\in\mathbb{Z}_q^m} \hat{f}(e)|e\rangle\right) \otimes \left(\sum_u |u\rangle\right).$$

3. It then adds $u^T A$ to the e registers in superposition, the state becomes:

$$\sum_{u\in\mathbb{Z}_q^n} \sum_{e\in\mathbb{Z}_q^m} \hat{f}(e)|u^T A + e\rangle \otimes |u\rangle. \tag{3}$$

4. Suppose there is a quantum algorithm that takes a state $\sum_{u\in\mathbb{Z}_q^n} \sum_{e\in\mathbb{Z}_q^m} \hat{f}(e)|u^T A + e\rangle \otimes |u\rangle$, outputs a state that is negligibly close to

$$\sum_{u\in\mathbb{Z}_q^n} \sum_{e\in\mathbb{Z}_q^m} \hat{f}(e)|u^T A + e\rangle \otimes |0\rangle, \tag{4}$$

   then we are done.

Let us now explain how to learn the secret u from the following state

$$|\phi_u\rangle := \sum_{e\in\mathbb{Z}_q^m} \hat{f}(e)|u^T A + e\rangle. \tag{5}$$

For convenience, although $|\phi_u\rangle$ depends on A, we ignore the subscript since A will be clear from the context. We focus on the case where q is a constant prime, and for $i = 1, ..., m$, $f(e_i) = 1/\sqrt{2\beta+1}$ for $e_i \in [-\beta, \beta] \cap \mathbb{Z}$ and 0 elsewhere. At the end of this subsection we will prove Theorem 8.

For the convenience of the rest of the presentation, let us also define

$$\text{for } v \in \mathbb{Z}_q,\ |\psi_v\rangle := \sum_{e\in\mathbb{Z}_q} \hat{f}(e)|(v+e) \bmod q\rangle. \tag{6}$$

Therefore Eqn (3) can also be written as

$$\sum_{u\in\mathbb{Z}_q^n} (|\phi_u\rangle \otimes |u\rangle) = \sum_{u\in\mathbb{Z}_q^n} \left(\bigotimes_{i=1,\ldots,m} |\psi_{(u^T A)_i}\rangle \otimes |u\rangle\right). \tag{7}$$

Now let us fix a vector $u \in \mathbb{Z}_q^n$. To learn u from $|\phi_u\rangle$, we look at each coordinate $|\psi_{(u^T A)_i}\rangle$ for $i = 1, ..., m$ separately. Let us (classically) pick a uniformly

random $y_i \in \mathbb{Z}_q$, then define a q-dimensional unitary matrix that always maps $|\psi_{y_i}\rangle$ to $|0\rangle$; more precisely,

$$U_{y_i} := \sum_{j=0}^{q-1} |j\rangle\langle\alpha_{i,j}|,$$

where $|\alpha_{i,0}\rangle := |\psi_{y_i}\rangle$ and the rest of the vectors $\{|\alpha_{i,j}\rangle\}_{j=1}^{q-1}$ are picked arbitrarily as long as U_{y_i} is unitary.

Looking ahead, we will apply U_{y_i} on $|\psi_{(u^T A)_i}\rangle$. Suppose we measure $U_{y_i}|\psi_{(u^T A)_i}\rangle$ and get an outcome in $\{0, 1, ..., q-1\}$. If the outcome is not 0, then we are 100% sure that $(u^T A)_i \neq y_i$. This is the basic idea of *filtering*, namely, we will filter out the case where $(u^T A)_i = y_i$ for a randomly chosen $y_i \in \mathbb{Z}_q$. Then we will handle the rest of the $q-1$ possibilities of $(u^T A)_i$ using the Arora-Ge algorithm (recall that we assume q is a constant in this subsection).

To explain why filtering works, consider for any $x, y \in \mathbb{Z}_q$. Let $U_y := \sum_{j=0}^{q-1} |j\rangle\langle\alpha_j|$ where $|\alpha_0\rangle := |\psi_y\rangle$ and the rest of the vectors $\{|\alpha_j\rangle\}_{j=1}^{q-1}$ span the rest of the space which are orthogonal to $|\alpha_0\rangle$. Then for any x, $|\psi_x\rangle$ can be written as a linear combination of basis $\{|\alpha_j\rangle\}_{j=0}^{q-1}$ (which contains $|\psi_y\rangle$), i.e,

$$|\psi_x\rangle = \sum_{j=0}^{q-1} \langle\alpha_j|\psi_x\rangle \cdot |\alpha_j\rangle.$$

Imagine if we apply U_y on $|\psi_x\rangle$ and measure, we will get q different possible outcomes.

- If the outcome is 0, we know that both $x = y$ and $x \neq y$ can happen.
  - If $x = y$, the outcome is 0 with probability 1;
  - Otherwise, the outcome is 0 with probability $|\langle\psi_x|\psi_y\rangle|^2$, which can still be non-zero.
- If the outcome is not 0, we know that it can only be the case: $x \neq y$. Because when $x = y$, the measurement will always give outcome 0.

In the next lemma, we show that if we choose y uniformly at random, the above measurement will give a non-zero outcome with "good" probability.

Lemma 2. *Let y be a uniformly random value in $\mathbb{Z}_q$. Then for any $x \in \mathbb{Z}_q$, the probability of measuring $U_y|\psi_x\rangle$ and getting an outcome not equal to 0 is at least $1 - 1/(2\beta + 1)$:*

$$\forall x, \Pr_{y \in \mathbb{Z}_q}[s \neq 0 \wedge s \leftarrow M_{\mathsf{st}} \circ U_y|\psi_x\rangle] \geq 1 - \frac{1}{2\beta + 1},$$

where M_{st} is the measurement operator in the computational basis.

Proof. Fixing y, the probability of getting outcome 0 is $|\langle\psi_y|\psi_x\rangle|^2$. The probability of getting a non-zero outcome (when y is chosen uniformly at random) is: $\frac{1}{q}\sum_{y=0}^{q-1}\left(1 - |\langle\psi_y|\psi_x\rangle|^2\right)$.

To bound the probability, we define $|\hat{\psi}_a\rangle = \sum_{x=0}^{q-1} f(x)\omega_q^{-xa}|x\rangle$, we have $|\psi_a\rangle = \mathsf{QFT}_q|\hat{\psi}_a\rangle$. For any x, y, the inner product $\langle\psi_x|\psi_y\rangle = \langle\hat{\psi}_x|\hat{\psi}_y\rangle$. The probability we want to bound is,

$$\begin{aligned}1 - \frac{1}{q}\sum_{y=0}^{q-1}\langle\psi_y|\psi_x\rangle\langle\psi_x|\psi_y\rangle = \; & 1 - \frac{1}{q}\sum_{y=0}^{q-1}\mathrm{Tr}\left[|\hat{\psi}_x\rangle\langle\hat{\psi}_x|\,|\hat{\psi}_y\rangle\langle\hat{\psi}_y|\right] \\ = \; & 1 - \frac{1}{q}\cdot\mathrm{Tr}\left[|\hat{\psi}_x\rangle\langle\hat{\psi}_x|\left(\sum_{y=0}^{q-1}|\hat{\psi}_y\rangle\langle\hat{\psi}_y|\right)\right].\end{aligned}$$

Let $\Psi := \sum_{y=0}^{q-1}|\hat{\psi}_y\rangle\langle\hat{\psi}_y|$. It can be simplified as follows:

$$\begin{aligned}\Psi = \sum_{y=0}^{q-1}|\hat{\psi}_y\rangle\langle\hat{\psi}_y| &= \sum_{y=0}^{q-1}\sum_{x\in\mathbb{Z}_q, x'\in\mathbb{Z}_q} f(x)f(x')\omega_q^{(x'-x)y}|x\rangle\langle x'| \\ &= \sum_{x\in\mathbb{Z}_q, x'\in\mathbb{Z}_q} f(x)f(x')\sum_{y=0}^{q-1}\omega_q^{(x'-x)y}|x\rangle\langle x'| \\ &= q\cdot\sum_{x=0}^{q-1} f(x)^2|x\rangle\langle x| \\ &= \frac{q}{2\beta+1}\sum_{x=-\beta}^{\beta}|x\rangle\langle x|.\end{aligned}$$

Here, we use the fact that $f(x) = 1/\sqrt{2\beta+1}$ for any $x \in [-\beta, \beta]\cap\mathbb{Z}$ and $f(x) = 0$ otherwise.Therefore,

$$1 - \frac{1}{q}\mathrm{Tr}\left[|\hat{\psi}_x\rangle\langle\hat{\psi}_x|\left(\sum_{y=0}^{q-1}|\hat{\psi}_y\rangle\langle\hat{\psi}_y|\right)\right] = 1 - \frac{1}{2\beta+1}\mathrm{Tr}\left[|\hat{\psi}_x\rangle\langle\hat{\psi}_x|\left(\sum_{x=-\beta}^{\beta}|x\rangle\langle x|\right)\right],$$

which is at least $1 - \frac{1}{2\beta+1}$. This follows from $\mathrm{Tr}\left[|\hat{\psi}_x\rangle\langle\hat{\psi}_x|\left(\sum_{x=-\beta}^{\beta}|x\rangle\langle x|\right)\right] \leq 1$. □

Thus, if we measure the superposition $|\phi_u\rangle$ entry-by-entry, with overwhelming probability, we will get at least $(1 - 1/(2\beta+1) - \varepsilon)m$ outcomes which are not 0 and at most $(1/(2\beta+1) + \varepsilon)m$ outcomes are 0 (for any constant $\varepsilon > 0$). Here we choose $\varepsilon = 0.1$.

Lemma 3. *For any fixed $x_1, \cdots, x_m \in \mathbb{Z}_q$, uniformly random $y_1, \cdots, y_m \in \mathbb{Z}_q$, the probability of measuring $U_{y_i}|\psi_{x_i}\rangle$ for all $i \in [m]$ and at least $(0.9 - 1/(2\beta+1))m$ outcomes being non-zero is at least $1 - O(e^{-m})$. Namely, for any fixed $x_1, \cdots, x_m \in \mathbb{Z}_q$,*

$$\Pr_{y_1,\cdots,y_m\in\mathbb{Z}_q}\left[z \geq \left(0.9 - \frac{1}{2\beta+1}\right)\cdot m \wedge \forall i, s_i \leftarrow M_{\mathsf{st}}\circ U_{y_i}|\psi_{x_i}\rangle\right] \geq 1 - O(e^{-m}),$$

where z is defined as the number of non-zero outcomes among all $s_1, \cdots, s_m$ and M_{st} is the measurement operator in the computational basis.

Proof. This is a direct consequence of Lemma 2 and Chernoff bound. □

By union bound, it can also be shown that, with probability at least $1 - O(q^n e^{-m})$, when the measurements on each bit are chosen uniformly at random, we will get at least $(0.9 - 1/(2\beta+1)) \cdot m$ non-zeros for all $u \in \mathbb{Z}_q^n$.

Corollary 3. *For any fixed A in $\mathbb{Z}_q^{n\times m}$, the probability that for all $u \in \mathbb{Z}_q^n$, measuring $U_{y_i}|\psi_{(u^TA)_i}\rangle$ for all $i \in [m]$ and at least $(0.9-1/(2\beta+1))m$ outcomes being non-zero is at least $1 - O(q^n e^{-m})$. Namely,*

$$\Pr_{y_1,\cdots,y_m\in\mathbb{Z}_q}\left[\forall u \in \mathbb{Z}_q^n, z_u \geq \left(0.9 - \frac{1}{2\beta+1}\right) m\right] \geq 1 - O(q^n e^{-m}),$$

where z_u is defined as the number of non-zero outcomes among all $s_{u,1}, \cdots, s_{u,m}$, each $s_{u,i}$ is defined as the measurement outcome of $U_{y_i}|\psi_{(u^TA)_i}\rangle$.

The above corollary implies that, for an overwhelming fraction (at least $1 - O(q^n e^{-m})$) of $y_1, \cdots, y_m$, the following event happens with probability at least $1-O(q^n e^{-m})$: for all $u \in \mathbb{Z}_q^n$, measuring $U_{y_i}|\psi_{(u^TA))_i}\rangle$ for all $i \in [m]$ and getting at least $(0.9 - 1/(2\beta+1))m$ outcomes being non-zero.

We are now ready to state the main theorem.

Theorem 9. *Let n be an integer, q be a constant prime, C be a sufficiently large constant. Let $m \geq (0.9 - 1/(2\beta+1))^{-1} \cdot C \cdot n^{q-1}q^2 \log q$. Then there exists a QPT algorithm that with overwhelming probability, given a random $A \in \mathbb{Z}_q^{n\times m}$ and $\sum_{u\in\mathbb{Z}_q^n} |\phi_u\rangle \otimes |u\rangle$, outputs a state negligibly close to $\sum_{u\in\mathbb{Z}_q^n} |\phi_u\rangle$. Here $|\phi_u\rangle$ is defined in Eqn. (5).*

Proof. Our algorithm works as follows on state $\sum_{u\in\mathbb{Z}_q^n} |\phi_u\rangle|u\rangle = \sum_u \bigotimes_{i=1}^m |\psi_{(u^TA)_i}\rangle|u\rangle$:

1. Pick m uniformly random values $y_1, ..., y_m \in \mathbb{Z}_q$. For each $i \in [m]$, construct a unitary $U_i := \sum_{j=0}^{q-1} |j\rangle\langle\alpha_{i,j}|$ where for $j = 0, ..., q-1$,

$$|\alpha_{i,j}\rangle := \begin{cases} |\psi_{y_i}\rangle, & \text{when } j = 0; \\ \text{Arbitrary } q\text{-dim unit vector orthogonal to } \{|\alpha_{i,k}\rangle\}_{k=0}^{j-1}, & \text{for } j \geq 1; \end{cases} \tag{8}$$

2. For $i = 1, ..., m$, apply U_i to the i^{th} register, we get

$$\begin{aligned} U_i|\psi_{(u^TA)_i}\rangle &= U_i\left(\sum_{j=0}^{q-1}\langle\alpha_{i,j}|\psi_{(u^TA)_i}\rangle \cdot |\alpha_{i,j}\rangle\right) \\ &= \left(\sum_{j=0}^{q-1}\langle\alpha_{i,j}|\psi_{(u^TA)_i}\rangle \cdot |j\rangle\right) =: \sum_{s_{u,i}\in\mathbb{Z}_q} w_{s_{u,i}}|s_{u,i}\rangle. \end{aligned}$$

Here, $s_{u,i}$ denotes the 'measurement outcome' of $U_i|\psi_{(u^TA)_i}\rangle$, but we do not physically measure the register $s_{u,i}$. We denote the vector $(s_{u,1}, \cdots, s_{u,m})$ by s_u.

3. Then we apply the quantum unitary implementation of the classical algorithm in [AG11] to $\sum_u \sum_{s_u \in \mathbb{Z}_q^m} w_{s_u} |s_u\rangle \otimes |u\rangle := \sum_u \bigotimes_{i=1}^m \sum_{s_{u,i} \in \mathbb{Z}_q} w_{s_{u,i}} |s_{u,i}\rangle \otimes |u\rangle$. Let the algorithm $D_{y_1, y_2, \cdots, y_m}$ be the following in Fig 1:

Algorithm 1 Learning u from $u^T A$

procedure $D_{y_1, y_2, \cdots, y_m}(s_u)$
  for each $i = 1, 2, \cdots, m$ **do**
    if $s_{u,i} \neq 0$ (meaning that $(u^T A)_i \neq y_i$) **then**
      Let a_i and $y_i \pmod q$ be a sample of LWE
    end if
  end for
  If there are more than $m' = (0.9 - 1/(2\beta + 1))m$ samples, it runs Arora-Ge algorithm over those samples to learn u and outputs u.
end procedure

For any $y_1, \cdots, y_m \in Z_q$ and $u \in \mathbb{Z}_q^n$, if $s_{u,i} \neq 0$, then $(u^T A)_i \neq y_i$; moreover, the LWE sample (a_i, y_i) has an error distribution with support $\{1, ..., q-1\}$, so Corollary 2 applies here.

We apply $D_{y_1, y_2, \cdots, y_m}$ in superposition to $\sum_u \sum_{s_u \in \mathbb{Z}_q^m} w_{s_u} |s_u\rangle \otimes |u\rangle$. For every fixed $u \in \mathbb{Z}_q^n$, let Bad_u be the set such that if all $s_u \in \mathsf{Bad}_u$, when we apply this algorithm to s_u, it does not compute u correctly.

By Corollaries 3 and 2, for an overwhelming fraction $(1 - O(q^n e^{-m}))$ of $y_1, \cdots, y_m$, for every u, $\sum_{s_u \in \mathsf{Bad}_u} |w_{s_u}|^2 \leq \mathsf{negl}(n)$. This is because:

(a) By Corollary 3, for an overwhelming fraction of $y_1, \cdots, y_m$, for every u, s_u provides at least $(0.9 - 1/(2\beta + 1))m = C \cdot n^{q-1} q^2 \log q$ samples with probability more than $1 - O(q^n e^{-m})$. Since $m \gg n$, it happens with overwhelming probability.

(b) By Corollary 2, as long as there are more than $C \cdot n^{q-1} q^2 \log q$ random samples, Arora-Ge algorithm succeeds with probability more than $1 - O(q^{-n^{q-1}})$. Note that the probability is taken over these random samples; in our case, the probability is taken over $A, y_1, \cdots, y_m$.

Thus, for an overwhelming fraction of $A, y_1, \cdots, y_m$, for every u, the weight $\sum_{s_u \in \mathsf{Bad}_u} |w_{s_u}|^2 \leq O(q^n e^{-m} + q^{-n^{q-1}}) = \mathsf{negl}(n)$.

Therefore, for an overwhelming fraction of $A, y_1, \cdots, y_m$, the resulting state is:

$$
\begin{aligned}
|\phi\rangle &:= D_{y_1,y_2,\cdots,y_m} \cdot q^{-n/2} \sum_{u\in\mathbb{Z}_q^n} \sum_{s_u\in\mathbb{Z}_q^m} w_{s_u}|s_u,u\rangle \\
&= q^{-n/2} \sum_{u\in\mathbb{Z}_q^n} \left(\sum_{s_u\notin \mathsf{Bad}_u} w_{s_u}|s_u,0\rangle + \sum_{s_u\in\mathsf{Bad}_u} w_{s_u}|s_u, D_{y_1,\cdots,y_m}(s_u)\rangle \right) \\
&= q^{-n/2} \sum_{u\in\mathbb{Z}_q^n} \left(\sum_{s_u} w_{s_u}|s_u,0\rangle + \mathsf{negl}_u(n)|\mathsf{err}_u\rangle \right) \\
&= q^{-n/2} \sum_{u\in\mathbb{Z}_q^n} \sum_{s_u} w_{s_u}|s_u,0\rangle + \mathsf{negl}(n)|\mathsf{err}\rangle.
\end{aligned}
$$

Here $\mathsf{negl}_u(n)$ is a complex number whose norm is negligible in n, $|\mathsf{err}_u\rangle$ is some unit vector. Similarly, it is the case for $\mathsf{negl}(n)$ and $|\mathsf{err}\rangle$.

4. Finally, we apply $\bigotimes_{i=1}^m U_i^{-1}$ to uncompute the projections and get

$$
\bigotimes_{i=1}^m U_i^{-1}|\phi\rangle = \sum_{u\in\mathbb{Z}_q^n} \bigotimes_{i=1,\ldots,m} |\psi_{(u^TA)_i}\rangle\otimes|0\rangle + \mathsf{negl}(n)|\mathsf{err}'\rangle = \sum_{u\in\mathbb{Z}_q^n} |\phi_u\rangle\otimes|0\rangle + \mathsf{negl}(n)|\mathsf{err}'\rangle.
$$

Therefore we get a state negligibly close to $\sum_{u\in\mathbb{Z}_q^n}|\phi_u\rangle = \sum_{u\in\mathbb{Z}_q^n}\sum_{e\in\mathbb{Z}_q^m}\hat{f}(e)|u^TA+e\rangle$. This completes the proof of Theorem 9. □

Finally, by Lemma 1 and Theorem 9, we complete the proof of Theorem 8.

4 Gram-Schmidt for Circulant Matrices

The general filtering algorithms used later in this paper construct unitary matrices obtained from applying the normalized Gram-Schmidt orthogonalization (GSO) on circulant matrices. The success probabilities of the general filtering algorithms are related to the norm of the columns in the matrices obtained from GSO. Thus, let us provide some related mathematical background in this section.

Given an ordered set of $k \le n$ linearly independent vectors $\{b_0, ..., b_{k-1}\}$ in $\mathbb{R}^n$, let $B := \left(b_0, ..., b_{k-1}\right) \in \mathbb{R}^{n\times k}$. For convenience, we sometimes denote b_i by B_i. Recall the Gram-Schmidt orthogonalization process.

Definition 6 (GSO). *The Gram-Schmidt orthogonalization of B, denoted as $\mathsf{GS}(B) = \left(\mathsf{GS}(b_0), \cdots, \mathsf{GS}(b_{k-1})\right)$, is defined iteratively for $i = 0, ..., k-1$ as*

$$
\mathsf{GS}(b_i) = b_i - \sum_{j=0}^{i-1} \frac{\langle b_i, \mathsf{GS}(b_j)\rangle}{\langle \mathsf{GS}(b_j), \mathsf{GS}(b_j)\rangle} \cdot \mathsf{GS}(b_j).
$$

Let us also define the normalized version of Gram-Schmidt orthogonalization.

Definition 7. *Given an ordered set of $k \leq n$ linearly independent vectors $\{b_0, ..., b_{k-1}\}$ in $\mathbb{R}^n$, let $B := \left(b_0, ..., b_{k-1}\right) \in \mathbb{R}^{n \times k}$. The normalized Gram-Schmidt orthogonalization of B, denoted as $\mathsf{NGS}(B) = \left(\mathsf{NGS}(b_0), \cdots, \mathsf{NGS}(b_{k-1})\right)$, is defined for $i = 0, ..., k-1$ as $\mathsf{NGS}(b_i) := \mathsf{GS}(b_i) / \|\mathsf{GS}(b_i)\|_2$ where $\mathsf{GS}(b_i)$ is defined in Definition 6.*

The following lemma is helpful for bounding the length of GSO vectors.

Lemma 4 (Derived from Corollary 14 of [Mic12]). *Let $D = \left(d_0, ..., d_{k-1}\right) := B \cdot (B^T \cdot B)^{-1}$. Then we have $\|\mathsf{GS}(b_{k-1})\|_2 = 1/\|d_{k-1}\|_2$.*

GSO of circulant matrices. Let $C \in \mathbb{R}^{n \times n}$ be a real circulant matrix, defined as

$$C := \begin{pmatrix} c_0 & c_1 & c_2 & ... & c_{n-1} \\ c_{n-1} & c_0 & c_1 & ... & c_{n-2} \\ c_{n-2} & c_{n-1} & c_0 & ... & c_{n-3} \\ ... & ... & ... & ... & ... \\ c_1 & c_2 & c_3 & ... & c_0 \end{pmatrix}. \tag{9}$$

Fact 10. *The QFT basis is an eigenbasis of a circulant matrix, namely,*

$$C = F_n^{-1} \cdot \Lambda \cdot F_n, \tag{10}$$

where $(F_n)_{i,j} := \frac{1}{\sqrt{n}} \cdot \omega_n^{ij}$, for $0 \leq i, j \leq n-1$; $\Lambda := \mathrm{diag}(\lambda_0, ..., \lambda_{n-1})$, where $\lambda_i := \sum_{j=0}^{n-1} c_j \cdot \omega_n^{ij}$. In other words, the eigenvalues of C are the QFT of the first row of C.

In our application, we need to compute the lower bound of the length of the k^{th} column of $\mathsf{GS}(C)$, for some $1 \leq k \leq n$ such that the first k columns of C are linearly independent. Below we present a lemma for general parameter settings. For simplicity, the readers can assume we are interested in the range of parameters where n is a polynomial, and k is either equal to n or $n - c$ where c is a constant.

Lemma 5. *Let $C = F_n^{-1} \cdot \Lambda \cdot F_n$ be a real circulant matrix where $\Lambda := \mathrm{diag}(\lambda_0, ..., \lambda_{n-1})$, $\lambda_i := \sum_{j=0}^{n-1} c_j \cdot \omega_n^{ij}$. Suppose $\lambda_0, ..., \lambda_{k-1}$ are non-zero and $\lambda_k, ..., \lambda_{n-1}$ are zero. Then the length of the k^{th} column of $\mathsf{GS}(C)$, i.e., $\|\mathsf{GS}(C)_{k-1}\|_2$, is lower-bounded by*

1. *If $k = n$, then $\|\mathsf{GS}(C)_{n-1}\|_2 \geq \frac{1}{\sqrt{n}} \cdot \min_{i=0,...,n-1} |\lambda_i|$.*
2. *If $k < n$, then $\|\mathsf{GS}(C)_{k-1}\|_2 \geq \frac{\sqrt{n}}{k \cdot 2^{n-k}} \cdot \min_{i=0,...,k-1} |\lambda_i|$.*

Please refer to the full version for the proof.

5 Quantum Algorithm for Solving the LWE State Problems

Recall in our mini result, every time a "measurement" (we do not physically implement the measurement) gives a non-zero result; it provides us with an

inequality $\langle u, a_i \rangle \neq y_i$. The algorithm, therefore, collects enough inequalities and then runs Arora-Ge to learn the secret vector u. There are two bottlenecks in the previous algorithm: (1) we are only able to filter out one value for $\langle u, a_i \rangle$; (2) to run Arora-Ge, one needs to collect many samples (up to roughly n^{q-1}). Therefore, it is only possible to provide quantum polynomial-time algorithms for $\mathsf{S|LWE}\rangle$, $\mathsf{C|LWE}\rangle$, and SIS^{∞} for a constantly large modulus q.

In this section, we generalize the filtering algorithm in a way that allows us to filter out $q-c$ many possible values of $\langle u, a_i \rangle$ for some constant c even when q is a polynomially large modulus. In the best possible case, the filtering algorithm can filter out $q-1$ possibilities and get the exact value of $\langle u, a_i \rangle$. Therefore, to learn the secret vector $u \in \mathbb{Z}_q^n$, one can collect roughly n samples and run Gaussian elimination. However, the probability of filtering out $q-1$ or $q-c$ (for some constant c) values depends on the concrete f and is typically very small. We will precisely show when such a probability is non-negligible.

We now provide quantum algorithms for $\mathsf{C|LWE}\rangle_{n,m,q,f}$ (cf. Def. 3) and $\mathsf{S|LWE}\rangle_{n,m,q,f}$ (cf. Def. 4). Let us first present the algorithms for a general error amplitude f, then state corollaries for some functions f of special interest. Looking ahead, the results in the full version use a slight modification of the algorithms presented in this section. Namely, in this section we will only show algorithms for functions f which allow us to filter out $q-1$ possible values then use Gaussian elimination, whereas the results in in the full version require us to deal with a function f that allows us to filter out $q-c$ possible values then use Arora-Ge.

5.1 Overview of the general filtering algorithm

Let q be a polynomially large modulus, f be an arbitrary noise amplitude. Define $|\psi_v\rangle := \sum_{e \in \mathbb{Z}_q} f(e)|v+e \bmod q\rangle$ for every $v \in \mathbb{Z}_q$. Following the basic notations and ideas in §3.2, let us now explain how to filter out two possible values for $(u^T A)_i$, say we are filtering out $(u^T A)_i = y_i$ and $(u^T A)_i = y_i + 1$ where y_i is a random value in $\mathbb{Z}_q$. To do so, let us define a basis $\{|\alpha_{i,j}\rangle\}_{j=0}^{q-1}$ where $|\alpha_{i,0}\rangle = |\psi_{y_i}\rangle$ and $|\alpha_{i,1}\rangle = \mathsf{NGS}\left(|\psi_{y_i+1}\rangle\right)$. The rest of the vectors in the basis are picked arbitrarily as long as they are orthogonal to $|\alpha_{i,0}\rangle$ and $|\alpha_{i,1}\rangle$.

Define $U_{y_i} := \sum_{j=0}^{q-1} |j\rangle\langle\alpha_{i,j}|$. Suppose we "measure" $U_{y_i}|\psi_{(u^T A)_i}\rangle$ and get an outcome in $\{0, 1, ..., q-1\}$:

1. If the outcome is 0, then $(u^T A)_i$ can be any values in $\mathbb{Z}_q$;
2. If the outcome is 1, then we are 100% sure that $(u^T A)_i \neq y_i$, since if $(u^T A)_i = y_i$, then the measurement outcome must be 0.
3. If the outcome is ≥ 2, then we are 100% sure that $(u^T A)_i$ does not equal to y_i or $y_i + 1$.

The idea can be further generalized by continuing to do normalized Gram-Schmidt orthogonalization. Suppose for a moment that $|\psi_{y_i+j}\rangle$, for $j = 0, ..., q-$

1, are linearly independent. Then we define unitary matrices

$$U_{y_i} := \sum_{j=0}^{q-1} |j\rangle\langle\alpha_{i,j}|, \text{ where } |\alpha_{i,j}\rangle = \mathsf{NGS}\left(|\psi_{y_i+j}\rangle\right).$$

Following the previous logic, if we "measure" $U_{y_i}|\psi_{(u^TA)_i}\rangle$, only the outcome "$q-1$" gives us a definitive answer of $(u^TA)_i$, that is, $(u^TA)_i = y_i + q - 1 \pmod q$.

The probability of filtering out $q-1$ values. It remains to understand the probability of getting the measurement outcome $q-1$.

$$\begin{aligned} &\Pr_{y_i\in\mathbb{Z}_q}\left[(u^TA)_i = y_i + q - 1 \pmod q \ \wedge\ q-1 \leftarrow M_{\mathsf{st}} \circ U_{y_i}|\psi_{(u^TA)_i}\rangle\right] \\ =\ &\frac{1}{q}\cdot\sum_{j\in\mathbb{Z}_q}|\langle\alpha_{i,q-1}|\psi_{y_i+j}\rangle|^2 = \frac{1}{q}\cdot|\langle\alpha_{i,q-1}|\psi_{y_i+q-1}\rangle|^2, \end{aligned}$$

where the second equality follows from the fact that $|\alpha_{i,q-1}\rangle$ is defined to be orthogonal to all the states except $|\psi_{y_i+q-1}\rangle$. Furthermore,

$$|\langle\alpha_{i,q-1}|\psi_{y_i+q-1}\rangle| = \left|\mathsf{NGS}\left(|\psi_{y_i+q-1}\rangle\right)^\dagger|\psi_{y_i+q-1}\rangle\right| = \|\mathsf{GS}\left(|\psi_{y_i+q-1}\rangle\right)\|_2,$$

i.e., it is exactly the norm of the Gram-Schmidt of $|\psi_{y_i+q-1}\rangle$. This quantity has been shown in Lemma 5 to be related to the minimum of $\hat{f}$ over $\mathbb{Z}_q$, namely,

$$\|\mathsf{GS}\left(|\psi_{y_i+q-1}\rangle\right)\|_2 \geq \min_{x=0,\ldots,q-1}|\hat{f}(x)|.$$

Therefore, we are able to use the general filtering technique to achieve polynomial-time quantum algorithms for $\mathsf{S}|\mathsf{LWE}\rangle_{n,m,q,f}$ and $\mathsf{C}|\mathsf{LWE}\rangle_{n,m,q,f}$ where q is polynomially large and f is a function such that the minimum of $\hat{f}$ over $\mathbb{Z}_q$ is non-negligible.

5.2 Quantum algorithm for generating LWE states with general error

Theorem 11. *Let q be a polynomially large modulus. Let $f : \mathbb{Z}_q \to \mathbb{R}$ be the amplitude for the error state such that the state $\sum_{e\in\mathbb{Z}_q} f(e)|e\rangle$ is efficiently constructible and $\eta := \min_{z\in\mathbb{Z}_q}|\hat{f}(z)|$ is non-negligible. Let $m \in \Omega\left(n\cdot q/\eta^2\right) \subseteq \mathsf{poly}(n)$, there exist polynomial-time quantum algorithms that solve $\mathsf{C}|\mathsf{LWE}\rangle_{n,m,q,f}$ and $\mathsf{S}|\mathsf{LWE}\rangle_{n,m,q,f}$.*

Proof. We will describe an algorithm for $\mathsf{C}|\mathsf{LWE}\rangle_{n,m,q,f}$. The algorithm for $\mathsf{S}|\mathsf{LWE}\rangle_{n,m,q,f}$ appears as a subroutine in the algorithm for $\mathsf{C}|\mathsf{LWE}\rangle_{n,m,q,f}$.

1. The algorithm first prepares the following state:

$$\bigotimes_{i=1}^{m}\left(\sum_{e_i\in\mathbb{Z}_q} f(e_i)|e_i\rangle\right)\otimes\sum_{u\in\mathbb{Z}_q^n}|u\rangle.$$

We abuse the notation of f to let $f(e) := \prod_{i=1}^{m} f(e_i)$ for $e := (e_1, ..., e_m)$. Then the state above can be written as $\sum_{e\in\mathbb{Z}_q^m} f(e)|e\rangle\otimes\sum_{u\in\mathbb{Z}_q^n}|u\rangle$.

2. It then adds $u^T A$ to the e registers in superposition, the state is:

$$\sum_{u\in\mathbb{Z}_q^n}\sum_{e\in\mathbb{Z}_q^m} f(e)|u^T A+e\rangle\otimes|u\rangle \tag{11}$$

Similarly, let us define

$$\text{for } v\in\mathbb{Z}_q,\ |\psi_v\rangle := \sum_{e\in\mathbb{Z}_q} f(e)|(v+e) \bmod q\rangle. \tag{12}$$

Therefore Eqn (11) can also be written as

$$\sum_{u\in\mathbb{Z}_q^n}\bigotimes_{i=1,\ldots,m}|\psi_{(u^T A)_i}\rangle\otimes|u\rangle. \tag{13}$$

3. Pick m uniformly random values $y_1, ..., y_m \in \mathbb{Z}_q$. Construct unitary matrices

$$\text{For } 1\le i\le m,\quad U_i := \sum_{j=0}^{q-1}|j\rangle\langle\alpha_{i,j}|,\ \text{where } |\alpha_{i,j}\rangle := \mathsf{NGS}\left(|\psi_{y_i+j}\rangle\right). \tag{14}$$

4. For $i = 1, ..., m$, apply U_i to the i^{th} register, we get

$$\begin{aligned} U_i|\psi_{(u^T A)_i}\rangle &= U_i\left(\sum_{j=0}^{q-1}\langle\alpha_{i,j}|\psi_{(u^T A)_i}\rangle\cdot|\alpha_{i,j}\rangle\right)\\ &= \left(\sum_{j=0}^{q-1}\langle\alpha_{i,j}|\psi_{(u^T A)_i}\rangle\cdot|j\rangle\right) =: \sum_{s_{u,i}\in\mathbb{Z}_q} w_{s_{u,i}}|s_{u,i}\rangle. \end{aligned}$$

5. Then we apply the quantum unitary implementation of Gaussian elimination to the superposition $\sum_u\sum_{s_u\in\mathbb{Z}_q^m} w_{s_u}|s_u\rangle := \sum_u\bigotimes_{i=1}^{m}\sum_{s_{u,i}\in\mathbb{Z}_q} w_{s_{u,i}}|s_{u,i}\rangle$. The algorithm $D_{y_1,y_2,\cdots,y_m}$ is described in Algorithm 2. In Lemma 8, we prove our parameters guarantee that with overwhelming probability,
   (a) There exists a set Bad_u such that for all $s_u \in \mathsf{Bad}_u$, when we apply $D_{y_1,y_2,\cdots,y_m}$ to s_u, it does not compute u correctly;
   (b) For an overwhelming choice of $y_1, \cdots, y_m$, for all u, $\sum_{s_u\in\mathsf{Bad}_u}|w_{s_u}|^2 = O(q^n e^{-m} + q^{-n}) = \mathsf{negl}(n)$. Here q^{-n} is the probability that the linear system is not full rank with $2n$ samples.

Algorithm 2 Learning u from $u^T A$

procedure $D_{y_1,y_2,\cdots,y_m}(\{s_{u,i}\}_{1\le i\le m})$
  for each $i = 1, 2, \cdots, m$ **do**
    if If $s_{u,i} = q-1$ (meaning that $(u^T A)_i = y_i + q - 1 \pmod q$) **then**
      Let a_i and $y_i - 1 \bmod q$ be one sample of the linear system
    end if
  end for
  With overwhelming probability, there are $\ge 2\cdot n$ random samples (to make sure the linear system is full rank)
  Run the Gaussian elimination algorithm to learn u and return u
end procedure

Therefore, for an overwhelming fraction of $A, y_1, \cdots, y_m$, the resulting state is:

$$\begin{aligned}|\phi\rangle &:= q^{-n/2}\cdot D_{y_1,y_2,\cdots,y_m}\sum_{u\in\mathbb{Z}_q^n}\sum_{s_u\in\mathbb{Z}_q^m} w_{s_u}|s_u,u\rangle\\ &= q^{-n/2}\sum_{u\in\mathbb{Z}_q^n}\left(\sum_{s_u\notin \mathsf{Bad}_u} w_{s_u}|s_u,0\rangle + \sum_{s_u\in\mathsf{Bad}_u} w_{s_u}|s_u, D_{y_1,\cdots,y_m}(s_u)\rangle\right)\\ &= q^{-n/2}\sum_{u\in\mathbb{Z}_q^n}\sum_{s_u} w_{s_u}|s_u,0\rangle + \mathsf{negl}(n)|\mathsf{err}\rangle.\end{aligned}$$

Here $\mathsf{negl}(n)$ returns a complex number whose norm is negligible in n, $|\mathsf{err}\rangle$ is some unit vector.

6. Finally, we just apply $\bigotimes_{i=1}^m U_i^{-1}$ to uncompute the projections and get

$$\bigotimes_{i=1}^m U_i^{-1}|\phi\rangle = \sum_{u\in\mathbb{Z}_q^n}\bigotimes_{i=1,\ldots,m}|\psi_{(u^TA)_i}\rangle\otimes|0\rangle + \mathsf{negl}(n)|\mathsf{err}'\rangle.$$

Thus, with overwhelming probability, we get a state close to $\sum_{u\in\mathbb{Z}_q^n}\sum_{e\in\mathbb{Z}_q^m} f(e)|u^TA+e\rangle$. It completes the description of our algorithm.

The analysis. Let us begin with an explanation of the properties of the unitary matrices U_i defined in Eqn (14). Recall from Eqn. (12) that $|\psi_v\rangle = \sum_{e\in\mathbb{Z}_q} f(e)|(v+e) \pmod q\rangle$.

Let $W_i := \sum_{j=0}^{q-1}|j\rangle\langle\psi_{y_i+j}|$. In other words, $W_i^T = \left(|\psi_{y_i}\rangle,\cdots,|\psi_{y_i+q-1}\rangle\right)$. Then $U_i^T = \mathsf{NGS}\left(W_i^T\right)$. We would like to show that the length of the GSO of $|\psi_{y_i+q-1}\rangle$, i.e., the length of the last column of $\mathsf{GS}\left(W_i^T\right)$, is non-negligible.

Lemma 6. $\|\mathsf{GS}\left(|\psi_{y_i+q-1}\rangle\right)\|_2 \ge \min_{z\in\mathbb{Z}_q}|\hat{f}(z)| = \eta$.

Proof. Note that W_i^T is a circulant matrix. The eigenvalues of W_i^T are $\left\{\sqrt{q}\cdot\hat{f}(z)\right\}_{z\in\mathbb{Z}_q}$ (see Fact 10). Therefore, by applying Lemma 5, we have $\|\mathsf{GS}\left(|\psi_{y_i+q-1}\rangle\right)\|_2 \ge \min_{z\in\mathbb{Z}_q}|\hat{f}(z)| = \eta$. □

Next, we relate the GSO of $|\psi_{y_i+q-1}\rangle$ to the probability of getting desirable samples in Algorithm 2.

Lemma 7. *For any fixed $x_1, \cdots, x_m \in \mathbb{Z}_q$,*

$$\Pr_{y_1,\cdots,y_m\in\mathbb{Z}_q}\left[z \geq \Omega\left(m\cdot(\eta^2/q)\right) \wedge \forall i, s_i \leftarrow M_{\mathsf{st}}\circ U_{y_i}|\psi_{x_i}\rangle\right] \geq 1 - O(e^{-m}),$$

where z is defined as the number of outcomes such that $s_i = q-1$ among all $s_1, \cdots, s_m$ and M_{st} is a measurement operator in the computational basis.

Proof. For $i = 1, ..., m$, we have

$$\begin{aligned}&\Pr_{y_i}[y_i + q - 1 = x_i] \cdot \Pr\left[s_i = q-1 \wedge s_i \leftarrow M_{\mathsf{st}}\circ U_{y_i}|\psi_{x_i}\rangle \mid y_i + q - 1 = x_i\right]\\ &= \frac{1}{q}\cdot|\langle\alpha_{i,q-1}|\psi_{y_i+q-1}\rangle|^2 = \frac{1}{q}\cdot\|\mathsf{GS}\left(|\psi_{y_i+q-1}\rangle\right)\|_2^2 \geq \frac{\eta^2}{q}.\end{aligned}$$

The lemma then follows Chernoff bound. □

Lemma 8. *When $m \in \Omega\left(n\cdot q/\eta^2\right) \subseteq \mathsf{poly}(n)$, for an overwhelming fraction of all possible $A, y_1, \cdots, y_m$, we have: for all u, $\sum_{s_u\in\mathsf{Bad}_u}|w_{s_u}|^2 \leq \mathsf{negl}(n)$.*

Proof. It follows from Lemma 7 that when $m \in \Omega\left(n\cdot q/\eta^2\right)$, we have $\geq 2\cdot n$ samples where $(u^T A)_i = y_i - 1 \pmod q$ with overwhelming probability. Thus, we can use Gaussian elimination to compute u. Therefore $\sum_{s_u\in\mathsf{Bad}_u}|w_{s_u}|^2 \leq \mathsf{negl}(n)$. □

This completes the proof of Theorem 11. □

5.3 Examples of error distributions of special interest

We give some examples of error amplitude f where $\min_{y\in\mathbb{Z}_q}|\hat{f}(y)|$ is non-negligible and q is polynomially large. The first example is where f is the bounded uniform distribution.

Corollary 4. *Let q be a polynomially large modulus. Let $B \in \mathbb{Z}$ such that $0 < 2B+1 < q$ and $\gcd(2B+1, q) = 1$. Let $f : \mathbb{Z}_q \to \mathbb{R}$ be $f(x) := 1/\sqrt{2B+1}$ where $x \in [-B, B]\cap\mathbb{Z}$ and 0 elsewhere. Let $m \in \Omega\left(n\cdot q^4\cdot(2B+1)\right) \subseteq \mathsf{poly}(n)$, there exist polynomial-time quantum algorithms that solve $\mathsf{C|LWE}\rangle_{n,m,q,f}$ and $\mathsf{S|LWE}\rangle_{n,m,q,f}$.*

Proof. The QFT of f is

$$\forall y\in\mathbb{Z}_q,\ \ \hat{f}(y) := \sqrt{\frac{1}{q\cdot(2B+1)}}\cdot\sum_{x=-B}^{B}\omega_q^{xy} = \sqrt{\frac{1}{q\cdot(2B+1)}}\cdot\frac{\sin\left(\frac{2\pi}{q}\cdot\frac{2B+1}{2}\cdot y\right)}{\sin\left(\frac{2\pi}{q}\cdot\frac{y}{2}\right)}. \tag{15}$$

Here we use the identity: $1 + 2\cos x + \cdots + 2\cos nx = \sin\left((n+\frac{1}{2})x\right)/\sin\left(\frac{x}{2}\right)$.

Note that when $y = 0$, $\hat{f}(y) = \sqrt{\frac{2B+1}{q}}$. When $y \in \{1, ..., q-1\}$, the denominator satisfies $0 < \sin\left(\frac{2\pi}{q} \cdot \frac{y}{2}\right) \leq 1$; since $\gcd(2B+1, q) = 1$, we have $\frac{(2B+1)y}{q} \notin \mathbb{Z}$ for any $y \in \{1, ..., q-1\}$, the numerator satisfies $\left|\sin\left(\frac{2\pi}{q} \cdot \frac{2B+1}{2} \cdot y\right)\right| \geq \left|\sin\left(\frac{\pi}{q}\right)\right| > \frac{1}{q}$.

Therefore $\eta = \min_{y \in \mathbb{Z}_q} |\hat{f}(y)| \geq \sqrt{\frac{1}{q \cdot (2B+1)}} \cdot \frac{1}{q}$. The corollary follows by plugging $\eta \geq \sqrt{\frac{1}{q \cdot (2B+1)}} \cdot \frac{1}{q}$ in Theorem 11. □

Remark 3. When $\gcd(2B + 1, q) = v$ for some $v > 1$, we have $\frac{(2B+1)y}{q} \in \mathbb{Z}$ for $q/v - 1$ values of $y \in \{1, ..., q-1\}$. Therefore $\hat{f}(y)$ defined in Eqn. (15) is 0 on $q/v - 1$ values. It is not clear to us how to extend our algorithm to the case where $\gcd(2B + 1, q) > 1$.

Other examples of f where $\min_{y \in \mathbb{Z}_q} |\hat{f}(y)|$ is non-negligible and q is polynomially large include Laplace and super-Gaussian functions. Their q-DFT is easier to express by first taking the continuous Fourier transform (CFT) of f, denoted as g, then discretize to obtain the DFT. Namely, for $y \in \mathbb{Z}_q$, $\hat{f}(y) = \frac{\sum_{z \in y + q\mathbb{Z}} g(z/q)}{\sum_{z \in \mathbb{Z}} g(z/q)}$. Let $0 < B < q/n^c$ for some $c > 0$.

1. Laplace: $f(x) = e^{-|x/B|}$, the CFT of f is $g(y) \propto \frac{2}{1+4(\pi By)^2}$.
2. Super-Gaussian: For $0 < p < 2$, $f(x) = e^{-|x/B|^p}$, the CFT of f is asymptotic to $g(y) \propto -\frac{\pi^{-p-\frac{1}{2}} |By|^{-p-1} \Gamma(\frac{p+1}{2})}{\Gamma(-\frac{p}{2})}$ (see, for example, [MS19]).

Acknowledgement

We sincerely thank Gábor Ivanyos for telling us the results in [IPS18]. We would also like to thank Luowen Qian, Léo Ducas, and the anonymous reviewers for their helpful comments. Y.C. is supported by Tsinghua University start-up funding and Shanghai Qi Zhi Institute. Q.L. is supported by the Simons Institute for the Theory of Computing, through a Quantum Postdoctoral Fellowship. M.Z. is supported in part by NSF.

References

AG11. Sanjeev Arora and Rong Ge. New algorithms for learning in presence of errors. In *Automata, Languages and Programming - 38th International Colloquium, ICALP 2011, Zurich, Switzerland, July 4-8, 2011, Proceedings, Part I*, pages 403–415, 2011. 6, 9, 14, 20

Ajt96. Miklós Ajtai. Generating hard instances of lattice problems (extended abstract). In *STOC*, pages 99–108, 1996. 1, 2, 3

BDK+18. Joppe W. Bos, Léo Ducas, Eike Kiltz, Tancrède Lepoint, Vadim Lyubashevsky, John M. Schanck, Peter Schwabe, Gregor Seiler, and Damien Stehlé. CRYSTALS - kyber: A cca-secure module-lattice-based KEM. In *EuroS&P*, pages 353–367. IEEE, 2018. 2

BKSW18. Zvika Brakerski, Elena Kirshanova, Damien Stehlé, and Weiqiang Wen. Learning with errors and extrapolated dihedral cosets. In *Public Key Cryptography (2)*, volume 10770 of *Lecture Notes in Computer Science*, pages 702–727. Springer, 2018. 3, 4, 5, 6, 7

BS16. Jean-François Biasse and Fang Song. Efficient quantum algorithms for computing class groups and solving the principal ideal problem in arbitrary degree number fields. In *Proceedings of the Twenty-Seventh Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 893–902. SIAM, 2016. 2

BV11. Zvika Brakerski and Vinod Vaikuntanathan. Efficient fully homomorphic encryption from (standard) LWE. In *IEEE 52nd Annual Symposium on Foundations of Computer Science, FOCS 2011, Palm Springs, CA, USA, October 22-25, 2011*, pages 97–106, 2011. 2

BV15. Zvika Brakerski and Vinod Vaikuntanathan. Constrained key-homomorphic prfs from standard lattice assumptions. In *Theory of Cryptography*, pages 1–30. Springer, 2015. 3

CDPR16. Ronald Cramer, Léo Ducas, Chris Peikert, and Oded Regev. Recovering short generators of principal ideals in cyclotomic rings. In *EUROCRYPT (2)*, volume 9666 of *Lecture Notes in Computer Science*, pages 559–585. Springer, 2016. 2

CDW17. Ronald Cramer, Léo Ducas, and Benjamin Wesolowski. Short stickelberger class relations and application to ideal-svp. In *EUROCRYPT (1)*, volume 10210 of *Lecture Notes in Computer Science*, pages 324–348, 2017. 2

CGS14. Peter Campbell, Michael Groves, and Dan Shepherd. Soliloquy: A cautionary tale, 2014. 2

CN97. Jin-yi Cai and Ajay Nerurkar. An improved worst-case to average-case connection for lattice problems. In *FOCS*, pages 468–477. IEEE Computer Society, 1997. 3

CvD07. Andrew M. Childs and Wim van Dam. Quantum algorithm for a generalized hidden shift problem. In *SODA*, pages 1225–1232. SIAM, 2007. 7

DKL+18. Léo Ducas, Eike Kiltz, Tancrède Lepoint, Vadim Lyubashevsky, Peter Schwabe, Gregor Seiler, and Damien Stehlé. Crystals-dilithium: A lattice-based digital signature scheme. *IACR Trans. Cryptogr. Hardw. Embed. Syst.*, 2018(1):238–268, 2018. 2, 3, 4

DKRV18. Jan-Pieter D'Anvers, Angshuman Karmakar, Sujoy Sinha Roy, and Frederik Vercauteren. Saber: Module-lwr based key exchange, cpa-secure encryption and cca-secure KEM. In *AFRICACRYPT*, volume 10831 of *Lecture Notes in Computer Science*, pages 282–305. Springer, 2018. 2

DM13. Nico Döttling and Jörn Müller-Quade. Lossy codes and a new variant of the learning-with-errors problem. In *EUROCRYPT*, volume 7881 of *Lecture Notes in Computer Science*, pages 18–34. Springer, 2013. 6

EHKS14. Kirsten Eisenträger, Sean Hallgren, Alexei Y. Kitaev, and Fang Song. A quantum algorithm for computing the unit group of an arbitrary degree number field. In *STOC*, pages 293–302. ACM, 2014. 2

ES16. Lior Eldar and Peter W. Shor. An efficient quantum algorithm for a variant of the closest lattice-vector problem, 2016. 12

FIM+03. Katalin Friedl, Gábor Ivanyos, Frédéric Magniez, Miklos Santha, and Pranab Sen. Hidden translation and orbit coset in quantum computing. In *STOC*, pages 1–9. ACM, 2003. 7

GGH96. Oded Goldreich, Shafi Goldwasser, and Shai Halevi. Collision-free hashing from lattice problems. *Electronic Colloquium on Computational Complexity (ECCC)*, 3(42), 1996. 2

GKZ19. Alex B Grilo, Iordanis Kerenidis, and Timo Zijlstra. Learning-with-errors problem is easy with quantum samples. *Physical Review A*, 99(3):032314, 2019. 5

GPV08. Craig Gentry, Chris Peikert, and Vinod Vaikuntanathan. Trapdoors for hard lattices and new cryptographic constructions. In *STOC*, pages 197–206, 2008. 2, 3

GVW13. Sergey Gorbunov, Vinod Vaikuntanathan, and Hoeteck Wee. Attribute-based encryption for circuits. In *STOC*, pages 545–554. ACM, 2013. 2

IPS18. Gábor Ivanyos, Anupam Prakash, and Miklos Santha. On learning linear functions from subset and its applications in quantum computing. In *ESA*, volume 112 of *LIPIcs*, pages 66:1–66:14. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2018. 7, 10, 11, 28

Kup05. Greg Kuperberg. A subexponential-time quantum algorithm for the dihedral hidden subgroup problem. *SIAM J. Comput.*, 35(1):170–188, 2005. 2

Mah18. Urmila Mahadev. Classical homomorphic encryption for quantum circuits. In *FOCS*, pages 332–338. IEEE Computer Society, 2018. 2

Mic02. Daniele Micciancio. Improved cryptographic hash functions with worst-case/average-case connection. In *STOC*, pages 609–618. ACM, 2002. 3

Mic12. Daniele Micciancio. CSE 206A: Lattice Algorithms and Applications. Lecture 2: The dual lattice, 2012. 22

MP13. Daniele Micciancio and Chris Peikert. Hardness of sis and lwe with small parameters. In *Advances in Cryptology–CRYPTO 2013*, pages 21–39. Springer, 2013. 3

MR07. Daniele Micciancio and Oded Regev. Worst-case to average-case reductions based on Gaussian measure. *SIAM Journal on Computing*, 37(1):267–302, 2007. 3

MS19. Stephen D Miller and Noah Stephens-Davidowitz. Kissing numbers and transference theorems from generalized tail bounds. *SIAM Journal on Discrete Mathematics*, 33(3):1313–1325, 2019. 28

NC16. Michael A. Nielsen and Isaac L. Chuang. *Quantum Computation and Quantum Information (10th Anniversary edition)*. Cambridge University Press, 2016. 13

Per88. A. Peres. How to differentiate between non-orthogonal states. *Physics Letters A*, 128:19–19, 1988. 9

Reg02. Oded Regev. Quantum computation and lattice problems. In *FOCS*, pages 520–529. IEEE Computer Society, 2002. 1, 2, 3, 7

Reg05. Oded Regev. On lattices, learning with errors, random linear codes, and cryptography. In *STOC*, pages 84–93. ACM, 2005. 2, 3, 4, 5, 12

SE94. Claus-Peter Schnorr and Martin Euchner. Lattice basis reduction: Improved practical algorithms and solving subset sum problems. *Mathematical programming*, 66(1):181–199, 1994. 3

SSTX09. Damien Stehlé, Ron Steinfeld, Keisuke Tanaka, and Keita Xagawa. Efficient public key encryption based on ideal lattices. In *Advances in Cryptology - ASIACRYPT 2009, 15th International Conference on the Theory and Application of Cryptology and Information Security, Tokyo, Japan, December 6-10, 2009. Proceedings*, pages 617–635, 2009. 3, 4, 5, 15