A Quasi-polynomial Time Algorithm for the Extrapolated Dihedral Coset Problem over Power-of-Two Moduli*

Shi Bai[1], Hansraj Jangir[1], Elena Kirshanova[2], Tran Ngo[1], and William Youmans[1]

[1] Florida Atlantic University, United States.
shih.bai@gmail.com, hansrajnitw@gmail.com, ngotbtran@gmail.com, youmans.wj@gmail.com
[2] Technology Innovation Institute, United Arab Emirates.
elenakirshanova@gmail.com

Abstract. The Learning With Errors (LWE) problem, introduced by Regev (STOC'05), is one of the fundamental problems in lattice-based cryptography, believed to be hard even for quantum adversaries. Regev (FOCS'02) showed that LWE reduces to the quantum Dihedral Coset Problem (DCP). Later, Brakerski, Kirshanova, Stehlé and Wen (PKC'18) showed that LWE reduces to a generalization known as the Extrapolated Dihedral Coset Problem (EDCP). We present a *quasi-polynomial* time quantum algorithm for the EDCP problems over power-of-two moduli using a quasi-polynomial number of samples, which also applies to the SLWE problem defined by Chen, Liu, and Zhandry (Eurocrypt'22). Our EDCP algorithm can be viewed as a provable variant to the "Simon-meets-Kuperberg" algorithm introduced by Bonnetain and Naya-Plasencia (Asiacrypt'18), adapted to the EDCP setting. We stress that our algorithm does *not* affect the security of LWE with standard parameters, as the reduction from standard LWE to EDCP limits the number of samples to be polynomial.

Keywords: Lattices · Learning With Errors · Dihedral Coset Problem · Quantum algorithms

1 Introduction

Euclidean lattices have gained significant research attention due to their potential in constructing cryptographic schemes that are conjectured to be quantum-resistant. The advantages of lattice-based cryptography are evident, as three out of the four schemes selected in the recent NIST Post-Quantum Cryptography Standardization process – Dilithium [45], Falcon [54] and Kyber [61] – rely on the presumed intractability of lattice problems for their security. The security

* This research was funded in part by the U.S. National Science Foundation under Grant No. 2044855 & 2122229.

of lattice-based cryptography is fundamentally based on two well-established average-case problems: the Short Integer Solution (SIS) problem [1] and the Learning With Errors (LWE) problem [57]. Both problems have been shown to admit worst-case hardness guarantees for approximating certain worst-case lattice problems, through worst-to-average case reductions.

The security of lattice-based assumptions, such as LWE, can be analyzed using algorithms that address the approximate Shortest Vector Problem (SVP), particularly through lattice reduction techniques [60,35,33,26,48]. When assessing the security of cryptographic schemes against quantum attacks, NIST recommends evaluating attack complexity under circuit constraints. This has inspired substantial research into the quantum adaptations of classical lattice algorithms and their quantum complexity analysis [6,22,17,9,12,15]. Another approach is to directly reduce LWE to related problems that are defined quantumly, and then examine the quantum complexity for solving such problems. For instance, Brakerski, Kirshanova, Stehlé and Wen [19] showed that LWE can be reduced to a quantum computational problem known as the Extrapolated Dihedral Coset Problem (EDCP). Our work focuses on this perspective – more precisely – we present a quasi-polynomial time and sample quantum algorithm for the EDCP problem. This algorithm also applies to a closely related problem, namely S $|\text{LWE}\rangle$, proposed by Chen, Liu and Zhandry [25].

1.1 LWE and its quantum versions

We give a brief introduction of some problems. The Learning with Errors (LWE) problem [58] with parameters n, q and $\alpha \in (0,1)$ asks to find $\mathbf{s} \in \mathbb{Z}_q^n$ from samples $(\mathbf{a}, \langle \mathbf{a}, \mathbf{s} \rangle + e \pmod q)$, where $\mathbf{a}$ is uniformly sampled in $\mathbb{Z}_q^n$ and e is sampled from a discrete Gaussian distribution of standard deviation αq.

Regev [58] showed that LWE reduces to the Dihedral Coset Problem (DCP), a problem which is closely related to the Hidden Subgroup Problem (HSP) on dihedral groups [42]. A converse reduction, however, is not known. Subsequently, Brakerski, Kirshanova, Stehlé, and Wen [19] showed that LWE is polynomial-time equivalent to a higher-dimensional analogue of DCP, known as the Extrapolated Dihedral Coset Problem (EDCP), which appears to be more naturally connected to the LWE problem. The $\text{EDCP}^l_{n,q,\chi}$ problem asks to recover the secret $\mathbf{s} \in \mathbb{Z}_q^n$ given l quantum states of the form:

$$\sum_{j \in \mathbb{Z}_q} \chi(j) \, |j\rangle \, |\mathbf{x}_i + j \cdot \mathbf{s} \pmod q\rangle \tag{1}$$

for uniformly random $\mathbf{x}_i \in \mathbb{Z}_q^n$. The two most relevant choices for the amplitude $\chi(\cdot)$ are the discrete Gaussian on $\mathbb{Z}$ with standard deviation r and the discrete uniform distribution on a support of size M. We denote them as $\text{G-EDCP}^l_{n,q,r}$ and $\text{U-EDCP}^l_{n,q,M}$, respectively. The work [19] shows that LWE can be reduced to both G-EDCP and U-EDCP. Taking the latter as an example, LWE with parameter α reduces to U-EDCP with $M \approx 1/(\alpha \cdot \mathsf{poly}(n))$. Conversely, an U-EDCP state can be tightly reduced to an LWE sample with $\alpha \approx 1/M$.

Recently, Chen, Liu and Zhandry [25] introduced another quantum problem known as S$|$LWE$\rangle$, which is closely related to both LWE and EDCP. The problem consists of finding $\mathbf{s} \in \mathbb{Z}_q^n$ given quantum states of the form:

$$\sum_{e \in \mathbb{Z}_q} \chi(e) \, |\langle \mathbf{a}_i, \mathbf{s} \rangle + e \pmod{q}\rangle \,,$$

where $\mathbf{a}_i$ is uniformly random in $\mathbb{Z}_q^n$ and known classically. Similarly, the two common choices for the amplitude $\chi(\cdot)$ are the discrete Gaussian and uniform distributions. It has been shown that the EDCP problem reduces to S$|$LWE$\rangle$ [19,25].

Given the natural connections among LWE, EDCP and S$|$LWE$\rangle$, it is essential to investigate their hardness from the perspective of quantum algorithms.

1.2 Our contributions

We summarize our main contributions.

1. *Quasi-polynomial time algorithm for* EDCP. First, we present an algorithm for EDCP that runs in time $2^{\mathcal{O}(\log n \log q)}$, requiring the same asymptotic number of EDCP samples and $\mathsf{poly}(n)$ quantum space. For $q = \mathsf{poly}(n)$, this becomes quasi-polynomial. The algorithm works for U-EDCP with any support size M and for G-EDCP with any deviation r, but requires the modulus q to be a power-of-two. To the best of our knowledge, this is the first provable algorithm to break the subexponential complexity barrier [19,42,56] for EDCP. This algorithm is given in Section 4.
2. *Polynomial time algorithm for* U-EDCP *with* $M = q/c$. A variant of our algorithm has polynomial running time for U-EDCP when $M = q/c$ for a constant c. This improves upon the polynomial time algorithm for EDCP from [25], which is applicable to the regime $M = q - c$. Note that these regimes are also solvable in $\mathsf{poly}(n)$ time via the EDCP-to-LWE reduction by the Arora-Ge algorithm [10]. This algorithm is given in Section 5.
3. *Quasi-polynomial time algorithm for* $\mathrm{EDCP}_{n,q,f}$ *for certain* f. We give a polynomial time algorithm for EDCP where the amplitude function f is positive valued and non-negligible on $\mathbb{Z}_q$, similar to the conditions required in [25, Theorem 12]. We also leverage our main algorithm to extend this result to a quasi-polynomial time algorithm for EDCP requiring only that f is non-negligible for at least two (known) points. This is given in Section 5.
4. *Quasi-polynomial time algorithm for* S$|$LWE$\rangle$. Our algorithm for the EDCP extends naturally to S$|$LWE$\rangle$, through a reduction from S$|$LWE$\rangle$ to $\overline{\text{EDCP}}$, a variant of the EDCP problem. For comparison, the best currently known algorithm for S$|$LWE$\rangle$ with Gaussian amplitude runs in subexponential time and requires subexponential samples [23]. This result is given in Section 6.

To conclude, our results provide further evidence that the EDCP / S$|$LWE$\rangle$ problem may not be as hard as DCP in certain parameter regimes, as already conjectured in [19].

Technical overview. We give a high-level overview of our algorithm in Section 4. We consider U-EDCP$_{n,q,M}$ instances where $M = 2$. We first apply a quantum Fourier Transform on the second register of the input states from Equation (1) and then measure, which leads to a state of the form:

$$|0\rangle + \omega_q^{\langle \mathbf{y}_k, \mathbf{s} \rangle} |1\rangle , \tag{2}$$

where $\mathbf{y}_k$ is uniform known classically, and $\omega_q := e^{2\pi i/q}$.

In a single step of our (iterative) algorithm we tensor $(n+1)$ of the above states to get:

$$\sum_{\mathbf{j} \in \mathbb{Z}_2^{n+1}} \omega_q^{\langle \mathbf{Y} \cdot \mathbf{j}, \mathbf{s} \rangle} |\mathbf{j}\rangle \tag{3}$$

for some classically known uniform matrix $\mathbf{Y} \in \mathbb{Z}_q^{n \times (n+1)}$. In a new register we compute $\mathbf{Y} \cdot \mathbf{j} \pmod 2$ and measure it to collapse the state to the superposition:

$$\omega_q^{\langle \mathbf{Y} \cdot \mathbf{x}_0, \mathbf{s} \rangle} |\mathbf{x}_0\rangle + \omega_q^{\langle \mathbf{Y} \cdot \mathbf{x}_1, \mathbf{s} \rangle} |\mathbf{x}_1\rangle ,$$

where $\mathbf{x}_0, \mathbf{x}_1$ are the unique solutions to the linear system $\mathbf{Y} \cdot \mathbf{j} \equiv \mathbf{t} \pmod 2$, and can be computed efficiently via Gaussian elimination. Up to a global phase, this state is the same as:

$$|0\rangle + \omega_q^{\langle \mathbf{Y} \cdot (\mathbf{x}_1 - \mathbf{x}_0), \mathbf{s} \rangle} |1\rangle = |0\rangle + \omega_{q/2}^{\langle \mathbf{y}', \mathbf{s} \rangle} |1\rangle ,$$

where $\mathbf{Y} \cdot (\mathbf{x}_1 - \mathbf{x}_0) \pmod q$ is even and $\mathbf{y}'$ is defined by lifting $\mathbf{Y} \cdot (\mathbf{x}_1 - \mathbf{x}_0) \pmod q$ to $\mathbb{Z}^n$ and halving the entries. Thus, we obtained samples similar to those in Equation (2), but with the modulus halved. Iterating this process we collect n samples with modulus two, e.g.,

$$\left\{ |0\rangle + (-1)^{\langle \mathbf{y}_k, \mathbf{s} \rangle} |1\rangle \right\}_{k=1}^{n}.$$

Measuring such a state in the Hadamard basis gives a linear equation in $\mathbf{s} \pmod 2$, and hence, with n linearly independent equations, we can learn n bits of $\mathbf{s}$. We can obtain the next batch of n bits of $\mathbf{s}$ by 'clearing' out the already known bits via a linear transformation on input EDCP states. Since each step requires $\mathcal{O}(n)$ states and there are $\mathcal{O}(\log q)$ steps, overall we require $2^{\Omega(\log n \log q)}$ samples.

We noted that the idea of tensoring more than two quantum states has been used in the Regev's polynomial-space algorithm [56] for Kuperberg's sieve [42,43], as well as folklore in [39]. In contrast to previous approaches, our method processes vectors $\mathbf{y}$ in a breadth-first way, 'clearing' a single bit across all entries at each merge. This approach incurs only a polynomial cost for each merge step.

Bonnetain and Naya-Plasencia's algorithm. It should be noted that our algorithm for EDCP is not entirely new. In fact, the core merging step (Equation (3)) and the lifting procedure already appeared[3] in the "Simon-meets-Kuperberg" algorithm of Bonnetain and Naya-Plasencia [18,16]. While their work

[3] We thank one of the reviewers for bringing this algorithm to our attention.

focuses on symmetric constructions, our focus is on the EDCP problem, which is more naturally connected to the LWE problem. We provide a rigorous, provable analysis that applies to EDCP instances with arbitrary M, along with its implications for the hardness of related problems such as $\mathrm{S}\,|\mathrm{LWE}\rangle$.

1.3 Related work and comparison

We review prior and related work, and compare our results to them. We first summarize known algorithms for EDCP, including our contributions, in Figure 1. For clarity, we focus on the U-EDCP problem here, though our algorithm also applies to the G-EDCP problem via the G-EDCP-to-U-EDCP reductions.

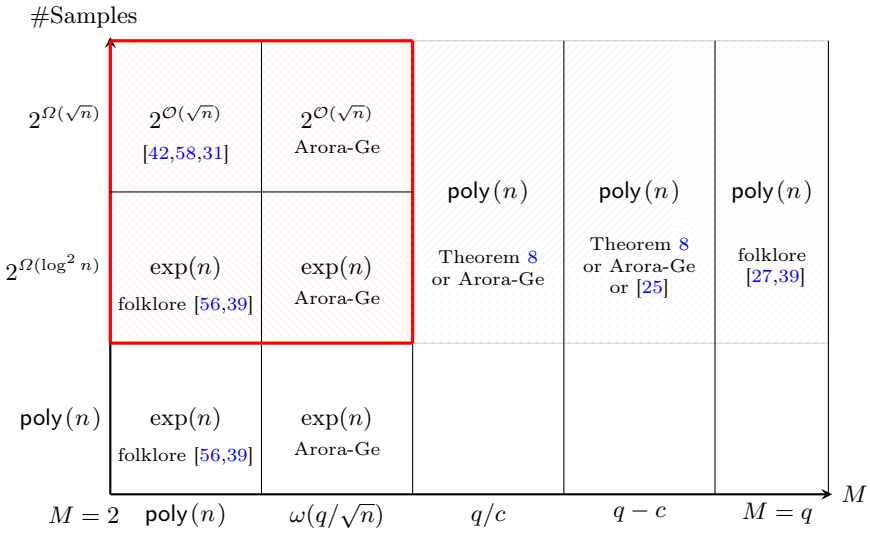

Fig. 1. Complexity of U-EDCP$_{n,q,M}$ as a function of M (x-axis) and the number of U-EDCP samples (y-axis). We fix a power-of-two $q = \mathsf{poly}(n)$. Our quasi-polynomial time algorithm outperforms the previously known subexponential (or exponential) algorithms in the top-left red-shaded region. It has the same polynomial time complexity in the region where $M > q/c$ for a constant c (region in blue), compared to the combination of the EDCP-to-LWE reduction [19] and Arora-Ge algorithm [10].

Kuperberg [42] introduced a subexponential time algorithm for DCP, known as Kuperberg's sieve, which also requires a subexponential number of samples. By combining this algorithm with Regev's reduction [58] from LWE to DCP, it appears plausible to obtain a subexponential algorithm for LWE *if* the number of DCP samples produced is subexponential. However, it is worth noting that the reduction in [58] is probabilistic, which only produces perfect DCP states with inverse polynomial probability. Furthermore, imperfect states are not efficiently detectable, so this does not lead to a subexponential algorithm for LWE.

The original algorithm of Kuperberg [42] also requires a subexponential quantum space. Regev [56] proposed an improvement to Kuperberg's sieve that uses only polynomial quantum space, at the cost of a slight increase in the running time. A second algorithm by Kuperberg [43] proposed further improvements based on this idea, which offer a better heuristic running time for general N and allow for balancing space and time. Asymptotically, all of these algorithms run in subexponential time and require a subexponential number of samples.

Subexponential algorithms for the U-EDCP problem are folklore, inheriting from the subexponential algorithms for DCP, as discussed above. First, (part of) the reduction from LWE to DCP by Regev [58] already implies a reduction from U-EDCP to DCP. Therefore, one can apply one of the subexponential-time algorithms discussed earlier to solve U-EDCP. Additionally, Doliskani [31] provides an algorithm that directly addresses U-EDCP with general support, based on Kuperberg's sieve and the algorithm of Childs and van Dam [27]. In summary, for the U-EDCP problem with general parameters, the best known methods run in subexponential time and require a subexponential number of U-EDCP samples. This corresponds to the square region for $M = \mathsf{poly}(n)$ and $\#\text{samples} = 2^{\Omega(\sqrt{n})}$ in Figure 1. Our quasi-polynomial algorithm in Section 4 outperforms these algorithms in this region.

Furthermore, our algorithm also covers the square region where $M = \mathsf{poly}(n)$, with the number of samples being $2^{\Omega(\log^2 n)}$ in Figure 1. In comparison, the best previous algorithm for U-EDCP, for such parameters, is the folklore algorithm from [56,39], which is achieved through a trade-off between sample complexity and running-time in Regev's algorithm [56] (see Subsection 3.3). It is also noted that one can transform U-EDCP instances into LWE instances using the reduction from [19], followed by lattice reduction. However, it is easy to see that lattice reduction requires exponential time, since both $1/\alpha$ and q are polynomial. Our quasi-polynomial algorithm outperforms these algorithms for this region as well.

$\text{U-EDCP}_{n,q,M}$ becomes easier as M increases. Leveraging the reduction from U-EDCP to LWE, one can show that U-EDCP with $M = \omega(q/\sqrt{n})$ reduces to a simpler version of LWE that can be solved in subexponential time using the Arora-Ge algorithm [10], when given a subexponential number of EDCP (hence, LWE) samples. This corresponds to the square region in Figure 1, where $M = \omega(q/\sqrt{n})$ and the number of samples is $2^{\Omega(\sqrt{n})}$. In comparison, our algorithm also runs in quasi-polynomial time in this region.

We further increase M: when $M = q/c$ for a constant c, we present a variant of our algorithm that runs in polynomial time for U-EDCP, as described in Section 5. This improves upon the polynomial-time algorithm for EDCP by Chen, Liu and Zhandry [25], which tackles the regime $M = q - c$. This corresponds to the square region in Figure 1, where $M = q/c$ (or $q - c$), with the number of samples polynomial. Setting $M = q$ makes EDCP solvable in polynomial time due to quantum Fourier transform, which is folklore [39].

In parallel, our quasi-polynomial time algorithm for EDCP can be extended to $\mathrm{S}\,|\text{LWE}\rangle$, via a reduction from $\mathrm{S}\,|\text{LWE}\rangle$ to a variant of EDCP. In comparison, the best algorithm [23] for $\mathrm{S}\,|\text{LWE}\rangle$ with Gaussian amplitude is a Kuperberg-

like algorithm that runs in subexponential time and also requires a subexponential number of samples. Our result also solves an open question from [23], namely: *"Given arbitrarily many samples, is* $\mathrm{S}\,|\mathrm{LWE}\rangle$ *with Gaussian amplitude still hard?"*

Finally, we *stress* that our quasi-polynomial algorithm does not compromise the security of LWE with standard parameters. Our algorithm requires a quasi-polynomial number of EDCP samples, whereas the existing reductions from standard LWE to EDCP [19,58] generate only a polynomial number of samples. Moreover, imperfect EDCP states cannot be efficiently detected. At present, we do not know how to either improve the reduction to yield more samples or enhance our algorithm. We attempted the latter without success and describe our failed attempts in Subsection 4.3. After introducing the necessary tools, we will give more details on the impacts and future directions in Section 7.

2 Preliminaries

Notations. We give the notations and definitions used in this paper. Let n be a positive integer. We let $[n] := \{0, \ldots, n-1\}$. Let q be a positive integer. We denote by $\mathbb{Z}_q$ the set of all integers modulo q. Lowercase bold letters such as $\mathbf{v}$ represent column vectors while uppercase bold letters such as $\mathbf{A}$ represent matrices. Given a vector $\mathbf{v}$, we denote by $\mathbf{v}^\top$ its transposed row vector. A vector $\mathbf{v}$ of length n has entries $(v_1, \ldots, v_n)^\top$. A zero vector is denoted as $\mathbf{0}$. The Hermitian inner product between two vectors $\mathbf{u}$ and $\mathbf{v}$ is denoted as $\langle \mathbf{u}, \mathbf{v} \rangle$. A matrix $\mathbf{B} = (\mathbf{b}_1, \mathbf{b}_2, \ldots, \mathbf{b}_n)$ is represented in a column-wise manner. For a vector $\mathbf{x}$ we denote its ℓ_2-norm by $\|\mathbf{x}\|$. We denote by $\log_b$ the logarithm of base b and denote $\log := \log_2$. We let $\lfloor a \rfloor$ denote the floor function. Let $\mathcal{B}_n(r)$ denote the n-dimensional ball of radius r centered at the origin. Let X be a countable set and $f : X \to \mathbb{R}$ be a function, we define $f(X) := \sum_{x \in X} f(x)$. Given two computational problems A and B, we denote $A \leq B$, if an algorithm for solving problem B could also be used as a subroutine to solve problem A.

We define the notations $\mathcal{O}(\cdot), \Omega(\cdot), \Theta(\cdot)$ and $\omega(\cdot)$ in the standard way. We let $\mathsf{poly}(\cdot)$ denote a polynomial function. These functions are typically defined with respect to the lattice rank n (or some security parameter κ). We say a positive function is quasi-polynomial if it is upper bounded by $2^{\mathcal{O}(\log^c n)}$ for some constant $c \geq 1$. Let $f : \mathbb{N} \to (0, 1]$ be a function. We say f is negligible if for all positive polynomials $p(\cdot)$ there exists a positive integer N such that $f(n) < \frac{1}{p(n)}, \forall n > N$. When f corresponds to a probability density function, we say f admits a negligible probability. Conversely, we say a function $g(n)$ is overwhelming if $1 - g(n)$ is negligible.

2.1 Euclidean lattices

A lattice $\mathcal{L}$ is an additive, finitely generated discrete subgroup of $\mathbb{Q}^m$. Let $\mathbf{B} = (\mathbf{b}_1, \cdots, \mathbf{b}_n) \in \mathbb{Q}^{m \times n}$ be a full rank matrix. The lattice $\mathcal{L}$ generated by $\mathbf{B}$ is defined as $\mathcal{L}(\mathbf{B}) = \{\mathbf{Bx} \mid \mathbf{x} \in \mathbb{Z}^n\}$. Let n denote the rank of the lattice $\mathcal{L}$. We

call the lattice $\mathcal{L}$ full rank when $m = n$. The matrix $\mathbf{B}$ is called a basis of $\mathcal{L}(\mathbf{B})$. Given a lattice $\mathcal{L}$, we let $\lambda_1(\mathcal{L})$ denote the Euclidean norm of a shortest non-zero vector in $\mathcal{L}$. Similarly, the i-th minimum $\lambda_i(\mathcal{L})$ is the radius of the smallest sphere centered at the origin that contains i linearly independent lattice points. A fundamental computational problem on lattices is the γ-approximate shortest vector problem (SVP_γ). On input a lattice basis $\mathbf{B}$, SVP_γ asks to find a non-zero lattice vector $\mathbf{v} \in \mathcal{L}(\mathbf{B})$ such that $\|\mathbf{v}\| \leq \gamma \cdot \lambda_1(\mathcal{L}(\mathbf{B}))$. When $\gamma = 1$, it is exact SVP. A closely-related problem is the shortest independent vectors problem (SIVP_γ): on input a lattice basis $\mathbf{B}$, it asks to find n linearly independent lattice vectors of length at most $\gamma \cdot \lambda_n(\mathcal{L}(\mathbf{B}))$.

In lattice-based cryptography, two q-ary lattices are common. Let $\mathbf{A} \in \mathbb{Z}_q^{m \times n}$ for integers $m > n, q \geq 2$. We define $\mathcal{L}_q(\mathbf{A}) := \{\mathbf{y} \in \mathbb{Z}^m \mid \mathbf{y} = \mathbf{A} \cdot \mathbf{s} \pmod q, \exists \mathbf{s} \in \mathbb{Z}^n\}$ and $\mathcal{L}_q^\perp(\mathbf{A}) := \{\mathbf{x} \in \mathbb{Z}^m \mid \mathbf{A}^\top \cdot \mathbf{x} = \mathbf{0} \pmod q\}$. The two lattices are dual to each other up to a scalar factor, as $\mathcal{L}_q(\mathbf{A}) = q \cdot (\mathcal{L}_q^\perp(\mathbf{A}))^*$. We will use the following lower bound, e.g., a generalization of [34, Lemma 5.3].

Lemma 1 ([23, Lemma 18]). *Let $q \geq 2$, $m \geq 2n \log q$, then for all but $q^{-0.16n}$-fractions of $\mathbf{A} \in \mathbb{Z}_q^{m \times n}$, we have $\lambda_1^\infty(\mathcal{L}_q(\mathbf{A})) \geq q/4$, where λ_1^∞ denotes the length of a shortest vector w.r.t. the infinity norm.*

2.2 Probability distribution

Given a distribution χ, we let $\mathrm{Supp}(\chi)$ denote its support. Let S be a finite set. We denote by $U(S)$ the uniform distribution on S. For example, $U(\mathbb{Z}_q)$ denotes the uniform distribution on the set $\mathbb{Z}_q$. Let χ be a distribution over S. We denote by $x \leftarrow\!\!\$\, \chi$ the process of sampling $x \in S$ according to the distribution χ. When the distribution χ is uniform, we use the shortcut notation $x \leftarrow\!\!\$\, S$ such as $x \leftarrow\!\!\$\, \mathbb{Z}_q$. We will use the following fact on the rank of uniform binary matrices.

Lemma 2. *Let matrices $\mathbf{A}$ be sampled uniformly from $\mathbb{Z}_2^{n \times (n+1)}$. Then $\mathbf{A}$ admits a full rank with some constant probability $C \geq 0.577$.*

Proof. The probability of $\mathbf{A}$ being full row-rank is $p(n) := \prod_{i=2}^{n+1}(1 - 2^{-i})$. The conclusion follows as $p(n)$ is decreasing in n and $\lim_{n \to \infty} p(n) \geq 0.577$. □

Discrete Gaussian. We will use several distributions in this work. For any vector $\boldsymbol{c} \in \mathbb{R}^n$ and any real deviation $r > 0$, the spherical Gaussian function is:

$$\rho_{r,\boldsymbol{c}}(\boldsymbol{x}) := \exp(-\pi \|\boldsymbol{x} - \boldsymbol{c}\|^2 / r^2).$$

The Gaussian distribution has density function $\upsilon_{r,\boldsymbol{c}}(\boldsymbol{x}) := \rho_{r,\boldsymbol{c}}(\boldsymbol{x})/r^n$. We may omit the subscript $\boldsymbol{c}$ when $\boldsymbol{c} = \mathbf{0}$. The spherical *discrete* Gaussian distribution over a lattice $\mathcal{L} \subseteq \mathbb{R}^n$ with deviation $r > 0$ and center $\boldsymbol{c}$ is defined as:

$$D_{\mathcal{L},r,\boldsymbol{c}} := \rho_{r,\boldsymbol{c}}(\boldsymbol{x})/\rho_{r,\boldsymbol{c}}(\mathcal{L}),\ \forall \boldsymbol{x} \in \mathcal{L}.$$

We give a few facts on the discrete Gaussian. First, we cite the tail bounds for Gaussian distributions on lattices.

Lemma 3 ([13, Lemma 1.5]). *For any $c > 1/\sqrt{2\pi}$, rank-n lattices $\mathcal{L}$ and vectors $\mathbf{u} \in \mathbb{R}^n$, we have:*

$$\rho(\mathcal{L} \setminus \mathcal{B}_n(c\sqrt{n})) < C^n \cdot \rho(\mathcal{L}), \quad \textit{and}$$

$$\rho((\mathcal{L} + \mathbf{u}) \setminus \mathcal{B}_n(c\sqrt{n})) < 2C^n \cdot \rho(\mathcal{L}),$$

where $C := c \cdot \sqrt{2\pi e} \cdot e^{-\pi c^2}$ for $C < 1$. We also have the bound:

$$\rho_r((\mathcal{L} + \mathbf{u}) \setminus \mathcal{B}_n(r\sqrt{n})) < 2^{-\Omega(n)} \cdot \rho_r(\mathcal{L}).$$

These bounds imply that truncating the tail has only a negligible impact on the distribution. For the lattice $\mathbb{Z}$, a similar bound can be established by connecting the radius to the security parameter.

Lemma 4 ([19, Lemma 1] and [44, Lemma 4.4]). *Let κ be the security parameter and $r > 0$, we have:*

$$\rho_r(\mathbb{Z} \setminus \mathcal{B}_1(\sqrt{\kappa}\, r)) < 2^{-\Omega(\kappa)} \cdot \rho_r(\mathbb{Z}).$$

This gives the following negligible event:

$$\Pr[|z| > \sqrt{\kappa}\, r, z \hookleftarrow_{\$} D_{\mathbb{Z},r}] < 2^{-\Omega(\kappa)}.$$

We will also use the following estimate on the total Gaussian mass on $\mathbb{Z}^n$.

Lemma 5 ([59, Claim 8.1]). *For $n \geq 1$ and $r > 0$, we have:*

$$r^n \cdot (1 + 2 \cdot e^{-\pi r^2})^n \leq \rho_r(\mathbb{Z}^n) \leq r^n \cdot (1 + (2 + 1/r) \cdot e^{-\pi r^2})^n.$$

2.3 Fourier transform

Let $f : \mathbb{R}^n \to \mathbb{R}$ be an absolutely integral function. The Fourier transform of f is defined as $\hat{f}(\mathbf{y}) := \int_{\mathbb{R}^n} f(\mathbf{x}) e^{-2\pi i \langle \mathbf{x}, \mathbf{y} \rangle} d\mathbf{x}$. The Fourier transform of the Gaussian function is $\hat{\rho}_r = r^n \cdot \rho_{1/r}$. We use the following form of Poisson summation.

Theorem 1 (Poisson summation). *Let $\mathcal{L}$ be a lattice and $\mathbf{c} \in \mathbb{R}^n$, we have:*

$$\rho_r(\mathcal{L} + \mathbf{c}) = r^n \cdot \det(\mathcal{L}^*) \cdot \sum_{\mathbf{x} \in \mathcal{L}^*} \exp(2\pi i \langle \mathbf{x}, \mathbf{c} \rangle) \cdot \rho_{1/r}(\mathbf{x}).$$

When the shift $\mathbf{c} = \mathbf{0}$, this simplifies to $\rho_r(\mathcal{L}) = \det(\mathcal{L}^*) \cdot \hat{\rho}_r(\mathcal{L}^*)$.

The discrete Fourier transform of a function f over $\mathbb{Z}_q^n$ is defined as

$$\tilde{f}(\mathbf{y}) := (1/q^{n/2}) \sum_{\mathbf{x} \in \mathbb{Z}_q^n} f(\mathbf{x})\, \omega_q^{\langle \mathbf{x}, \mathbf{y} \rangle}, \ \forall \mathbf{y} \in \mathbb{Z}_q^n.$$

2.4 Quantum computation

The state space of an n-qubit quantum system resides in the 2^n-dimensional complex space $\mathbb{C}^{2^n}$. A quantum state $|u\rangle$ can be expressed as a superposition with respect to some standard basis $\{|i\rangle\}_i$, as $|u\rangle = \sum_i a_i |i\rangle$, where the coefficients $a_i \in \mathbb{C}$ satisfy the normalization $\sum_i |a_i|^2 = 1$. The trace distance T between two quantum states ρ and σ is $T(\rho, \sigma) := (1 - |\langle\rho|\sigma\rangle|^2)^{1/2}$. The trace distance bounds the distinguishability of two quantum states in a similar way as the statistical distance between two distributions. In particular, trace-preserving quantum operations cannot increase the trace distance and unitary operations preserve the trace distance [50, Theorem 9.2]. Similarly, the induced distributions on the possible outcomes of measurements have equal or smaller statistical distance than the trace distance [50, Theorem 9.1]. We say two quantum states are close if their trace distance is inverse exponentially small.

We will present our algorithms in the quantum circuit model, where the operations performed on quantum states are unitary operators, known as "gates". As usual, we will implicitly assume a set of universal operators and simplify the discussions leveraging the Solovay-Kitaev theorem [40,29].

We will use the Quantum Fourier Transform (QFT), which can be done for arbitrary finite Abelian groups [41] – we mostly use the QFT over $\mathbb{Z}_q^n$. Let $q \geq 2$ be an integer and denote $\omega_q := e^{2\pi i/q}$. The QFT over $\mathbb{Z}_q^n$ of an input state $|\phi\rangle := \sum_{\mathbf{x}\in\mathbb{Z}_q^n} f(\mathbf{x}) |\mathbf{x}\rangle$ is

$$\mathrm{QFT}_{q^n} |\phi\rangle = (1/q^{n/2}) \sum_{\mathbf{y}\in\mathbb{Z}_q^n} \sum_{\mathbf{x}\in\mathbb{Z}_q^n} \omega_q^{\langle \mathbf{x}, \mathbf{y}\rangle} f(\mathbf{x}) |\mathbf{y}\rangle .$$

This can be instantiated efficiently with $\mathsf{poly}(n \log q)$ gates. We use the exact QFT, but there is also an approximate QFT [28].

To conclude, our algorithms consist of primitive components that can be efficiently instantiated or approximated using unitary operations such as modular addition/multiplication and QFT. The algorithms may also use a polynomial number of ancilla registers – initialized to the state $|0\rangle$ – to facilitate the computation. Hence, the algorithms may also include operations such as ancilla initialization, tracing out, and measurement in the computational basis. These operations are trace-preserving and therefore cannot increase the trace distance. For further background on quantum computation, we refer to [50,64].

We will use the following form of the quantum rejection sampling [51].

Lemma 6 (Adapted from [19, Lemma 7]). *Let $|\phi\rangle$ be a normalized input quantum state, where $|\phi\rangle := \sum_{j\in D_1} \pi_j |\xi_j\rangle |j\rangle$ for some (possibly unknown) $|\xi_j\rangle \in \mathbb{C}^d$ and $\pi_j \geq 0$. Let support $D_2 \subseteq D_1$. There is a quantum algorithm that, with $|\phi\rangle$ as input, outputs a quantum state $|\psi\rangle := (1/B) \sum_{j\in D_2} \sigma_j |\xi_j\rangle |j\rangle$ where $0 \leq \sigma_j \leq \pi_j$ and succeeds with probability $B^2 = \sum_{j\in D_2} \sigma_j^2$.*

2.5 Computational problems

Two fundamental average-case problems used in lattice-based cryptography are the short integer solution problem (SIS) [1] and the learning with errors problem (LWE) [57]. They are defined as follows.

Definition 1 (Search Learning with Errors [57]). *With input parameters $m, n \geq 1$, a modulus $q \geq 2$ and a distribution χ, the search version of $\mathrm{LWE}^m_{n,q,\chi}$ problem consists of m samples of the form $(\mathbf{a}, b) \in \mathbb{Z}_q^n \times \mathbb{Z}_q$, with $\mathbf{a} \hookleftarrow_{\$} \mathbb{Z}_q^n$, $b = \langle \boldsymbol{a}, \boldsymbol{s} \rangle + e \pmod q$ and $e \hookleftarrow_{\$} \chi$ for a fixed $\mathbf{s} \in \mathbb{Z}_q^n$. We say that an algorithm solves the search $\mathrm{LWE}^m_{n,q,\chi}$ problem if it finds $\mathbf{s}$ with non-negligible probability w.r.t the input $n \log q$.*

Typically, we take $m = \Omega(n \log q)$ for the problem to be well-defined. We also express $\mathrm{LWE}^m_{n,q,\alpha}$ samples in its matrix form $(\mathbf{A}, \mathbf{b}) \in \mathbb{Z}_q^{m \times n} \times \mathbb{Z}_q^m$ where $\mathbf{b} = \mathbf{As}+\mathbf{e} \pmod q$. If the number of samples m is not specified, we denote it as $\mathrm{LWE}_{n,q,\chi}$. In this work, χ is usually a discrete Gaussian $D_{\mathbb{Z},\alpha q}$ (in such case, we use the shortcut $\mathrm{LWE}_{n,q,\alpha}$) or bounded uniform distribution. Regev [57, Theorem 1.1] shows that: (informally) for $\alpha q = \Omega(\sqrt{n})$, an efficient algorithm for $\mathrm{LWE}_{n,q,\alpha}$ implies an efficient (quantum) algorithm for worst-case lattice problems such as SIVP_γ with $\gamma = \tilde{\mathcal{O}}(n/\alpha)$. For polynomial modulus q, this requires α to be at most polynomially small – we refer to such α, q as "standard LWE parameters". A dual problem to LWE is the short integer solution problem (SIS) [1].

Definition 2 (Short Integer Solution [1]). *Input parameters $m, n \geq 1$, a modulus $q \geq 2$ and a positive integer bound β, and let $\mathbf{A}^\top$ be uniformly sampled from $\mathbb{Z}_q^{n \times m}$. The $\mathrm{SIS}_{m,n,q,\beta}$ problem asks to find a non-zero $\boldsymbol{x} \in \mathbb{Z}^m$ such that $\mathbf{A}^\top \boldsymbol{x} \equiv \mathbf{0} \pmod q$ and $\|\boldsymbol{x}\|_2 \leq \beta$.*

Typically, we take $m = \Omega(n \log q)$. A variant of the problem, namely the inhomogeneous short integer solution problem (ISIS), asks to recover a non-zero $\boldsymbol{x} \in \mathbb{Z}^m$ given $\mathbf{A}^\top \boldsymbol{x} \equiv \boldsymbol{y} \pmod q$ with $\|\boldsymbol{x}\|_2 \leq \beta$, where $\mathbf{y}$ is given as input. Sometimes, the solution is bounded by the infinity norm, in such case, we let it denote as $\mathrm{SIS}^\infty_{m,n,q,\beta}$ (or $\mathrm{ISIS}^\infty_{m,n,q,\beta}$). The SIS problem has been shown [1,46] to be at least as hard as approximating several worst-case lattice problems.

We now introduce some problems that are defined quantumly. Regev [58] showed that LWE can be reduced to the Dihedral Coset Problem (DCP), which is closely related to the Hidden Subgroup Problem (HSP) on dihedral groups [42]. Concretely, DCP^l_N asks to recover the secret $s \in \mathbb{Z}_N$, given l states with uniform x_i, of the form:

$$\{|0\rangle |x_i\rangle + |1\rangle |x_i + s \bmod N\rangle\}_{i=1}^l.$$

The space $\mathbb{Z}_N$ hides the secret, which is usually exponential w.r.t the security parameter. A higher dimensional version of the DCP problem, which appears to be more naturally related to LWE, is defined in [19] as follows:

Definition 3 (Search Extrapolated Dihedral Coset Problem [19, Definition 4]). *On input a dimension parameter n, modulus $q \geq 2$ and a (discrete)*

distribution χ, the search Extrapolated Dihedral Coset Problem ($\mathrm{EDCP}^l_{n,q,\chi}$*) consists of l input states of the form*

$$\left\{ \sum_{j \in \mathrm{Supp}(\chi)} \chi(j) \left|j\right\rangle \left|\mathbf{x}_i + j \cdot \mathbf{s} \bmod q\right\rangle \right\}^l_{i=1},$$

where $\mathbf{x}_i \in \mathbb{Z}_q^n$ are sampled uniformly, and asks to recover the secret $\mathbf{s} \in \mathbb{Z}_q^n$.

We say an algorithm solves $\mathrm{EDCP}^l_{n,q,\chi}$ efficiently if it finds $\mathbf{s}$ with advantage $\mathsf{poly}(1/(n \log q))$ in time $\mathsf{poly}(n \log q)$. Typically we require $l = \Omega(n \log q)$ for the problem to be well-defined [31], which is always the case in this work. The two most interesting amplitude functions for χ are discrete Gaussian and uniform (up to a normalization factor): when $\chi = D_{\mathbb{Z},r}$ for $r \geq 1$, we denote it as $\text{G-EDCP}^l_{n,q,r}$. When $\chi = U(\mathbb{Z}_M)$ for $2 \leq M \leq q$, we denote it as $\text{U-EDCP}^l_{n,q,M}$. We omit the number of samples l when it is not important.

The EDCP problem reduces to the LWE problem tightly, and conversely, LWE reduces to the EDCP problem where the number of samples produced are limited. We give both reductions.

Theorem 2 **(G-EDCP $\leq$ LWE [19, Theorem 4]).** *Let κ be the security parameter. Given $l = \Omega(n \log q)$ many $\text{G-EDCP}_{n,q,r}$ samples where $r = \Omega(\sqrt{\kappa})$, there exists a probabilistic polynomial-time quantum reduction from $\text{G-EDCP}^l_{n,q,r}$ to $\mathrm{LWE}^l_{n,q,\alpha}$ with $\alpha = 1/r$ that succeeds with probability $(1 - 2^{-\Omega(\kappa)})^l$.*

The theorem also assumes $q/r = \Omega(\sqrt{\kappa})$ implicitly to ensure that the support in Definition 3 is predominantly contained in $\mathbb{Z}_q$. It is worth noting that this reduction works for l, beyond polynomial size, as long as it is below exponential. The converse reduction, however, puts a restriction on the number of samples that can be obtained with respect to the given LWE parameters.

Theorem 3 **(LWE $\leq$ G-EDCP [19, Theorem 3]).** *Let κ be the security parameter. Given $m = \Omega(n \log q)$ many $\mathrm{LWE}_{n,q,\alpha}$ samples, there exists a probabilistic polynomial-time quantum reduction from $\mathrm{LWE}^m_{n,q,\alpha}$ to $\text{G-EDCP}^l_{n,q,r}$ with*

$$r \cdot l < 1/(6\sqrt{2\pi e m \kappa} \alpha q^{n/m}) \tag{4}$$

that succeeds with probability $(1 - 2^{-ml})^{ml} \geq 1/4$.

This reduction transforms m LWE samples into l G-EDCP samples subject to Equation (4). For standard LWE parameters (e.g., both q and $1/\alpha$ are polynomial), the condition restricts $m \cdot l$ to be a polynomial. This implies that the reduction can only produce polynomially many EDCP samples, given standard LWE parameters.

From a cryptanalysis point of view, it is more convenient to work with the U-EDCP problem. There exists reductions between the G-EDCP and U-EDCP problems using quantum rejection sampling.

Lemma 7 (G-EDCP $\leq$ U-EDCP **[19, Lemma 8]**). *Let κ be the security parameter. Let n, q, M, l be integers greater than 1 and r a positive real number. There exists a probabilistic polynomial-time quantum reduction from* $\text{G-EDCP}^{l}_{n,q,r}$ *to* $\text{U-EDCP}^{\mathcal{O}(l/\kappa)}_{n,q,M}$*, where $M = c \cdot r$ for some constant c, that succeeds with probability $1 - 2^{-\Omega(\kappa)}$.*

This reduction transforms l G-EDCP samples to $\mathcal{O}(l/\kappa)$ U-EDCP samples with overwhelming success probability. On the other hand, there exists a converse reduction from U-EDCP $\leq$ G-EDCP.

Lemma 8 (U-EDCP $\leq$ G-EDCP **[19, Lemma 9]**). *Let κ be the security parameter. Let n, q, M, l be integers greater than 1 and r a positive real number. There exists a probabilistic polynomial-time quantum reduction from* $\text{U-EDCP}^{l}_{n,q,M}$ *to* $\text{G-EDCP}^{\mathcal{O}(l/\kappa^{1.5})}_{n,q,r}$*, where $M = \sqrt{\kappa} \cdot r$, that succeeds with a probability $1 - 2^{-\Omega(\kappa)}$.*

The EDCP problem also admits self-reductions [19,31] that convert a larger support into a smaller one. We use the following self-reduction.

Lemma 9 (U-EDCP **self-reduction [31, Lemma 10]**). *Let κ be the security parameter. Let n, q, l, M, M' be integers greater than 1 and $M \geq M'$. There exists a probabilistic polynomial-time quantum reduction from* $\text{U-EDCP}^{l}_{n,q,M}$ *to* $\text{U-EDCP}^{\Theta(l)}_{n,q,M'}$ *that succeeds with a constant success probability.*

Two closely related quantum problems are $\text{S}\,|\text{LWE}\rangle$ and $\text{C}\,|\text{LWE}\rangle$. $\text{S}\,|\text{LWE}\rangle$ is similar to LWE but defined on quantum state.

Definition 4 (Solve $\text{S}\,|\text{LWE}\rangle$ **[25,23]).** *On input parameters m, n, q and a function $\chi : \mathbb{Z}_q \to \mathbb{C}$, the* $\text{S}\,|\text{LWE}\rangle_{n,m,q,\chi}$ *problem consists of m input states:*

$$\Big\{ \sum_{e_i \in \mathbb{Z}_q} \chi(e_i)\,|\langle \mathbf{a}_i, \mathbf{s}\rangle + e_i \bmod q\rangle \Big\}_{i=1}^{m},$$

where $\mathbf{a}_i \in \mathbb{Z}_q^n$ are uniformly distributed and known classically, and asks to recover the secret $\mathbf{s} \in \mathbb{Z}_q^n$.

In our context, the function χ often comes from a probability density function which is related to the error distribution of some LWE instances. When χ is a discrete Gaussian, we denote it as $\text{G-S}\,|\text{LWE}\rangle$; when χ is uniform, we denote it as $\text{U-S}\,|\text{LWE}\rangle$. The $\text{G-S}\,|\text{LWE}\rangle$ problem has been used implicitly in the reduction from EDCP to LWE in the proof of Theorem 2. In addition, there is an efficient quantum reduction from $\text{S}\,|\text{LWE}\rangle_{n,q,\bar{\chi}}$ to $\text{EDCP}_{n,q,\chi}$ in [19, Theorem 4] and [25, Lemma 47]. Note that this reduction uses the QFT, which converts the input distribution into its DFT in the output.

A related problem, namely $\text{C}\,|\text{LWE}\rangle$, is defined as the task of constructing a certain quantum state.

Definition 5 (Construct $\text{C}\,|\text{LWE}\rangle$ **[25,23]).** *Input parameters m, n, q and a function $\chi : \mathbb{Z}_q \to \mathbb{C}$, the* $\text{C}\,|\text{LWE}\rangle_{n,m,q,\chi}$ *problem asks to construct a quantum*

state of the form:

$$\sum_{\mathbf{s}\in\mathbb{Z}_q^n}\bigotimes_{i=1}^{m}\Big(\sum_{e_i\in\mathbb{Z}_q}\chi(e_i)\,|\langle\mathbf{a}_i,\,\mathbf{s}\rangle+e_i \bmod q\rangle\Big),$$

where the $\mathbf{a}_i\in\mathbb{Z}_q^n$ *are uniformly distributed and given as inputs.*

When χ is a discrete Gaussian, we denote G-C $|\mathrm{LWE}\rangle^m_{n,q,r}$; when χ is uniform on a support size M, we denote it as U-C $|\mathrm{LWE}\rangle^m_{n,q,M}$. It is noted by [25] that if there is a quantum algorithm that solves the S $|\mathrm{LWE}\rangle$ problem without collapsing the input quantum states, then there is a quantum algorithm that solves the C $|\mathrm{LWE}\rangle$ problem for the same parameters. Looking ahead, our quasi-polynomial algorithms work for S $|\mathrm{LWE}\rangle$, however, do destroy the input states. The C $|\mathrm{LWE}\rangle$ problem has important applications. For example, a recent work by Debris-Alazard, Fallahpour and Stehlé [30] showed that solving C $|\mathrm{LWE}\rangle$ for appropriate amplitudes implies a quantum oblivious LWE sampler. In addition, both S $|\mathrm{LWE}\rangle$ and C $|\mathrm{LWE}\rangle$ naturally reduce to the LWE problem. The converse, however, is less trivial – we refer the details to [25,23].

Finally, there is an efficient reduction from SIS to LWE implicitly stated in [62] and discussed in [25]. Let $\mathbf{A}^\top\in\mathbb{Z}_q^{n\times m}$ be the given SIS matrix. The reduction aims to construct a quantum state of the form:

$$\mathrm{C}\,|\mathrm{SIS}\rangle:=\sum_{\mathbf{x}\in\mathbb{Z}_q^m,\,\mathbf{A}^\top\mathbf{x}\equiv 0 \pmod q}\chi(\mathbf{x})\,|\mathbf{x}\rangle\,.$$

It is sufficient to construct a C $|\mathrm{LWE}\rangle$ state of the form

$$\sum_{\mathbf{s}\in\mathbb{Z}_q^n}\sum_{\mathbf{e}\in\mathbb{Z}_q^m}\tilde{\chi}(\mathbf{e})\,|\mathbf{A}\mathbf{s}+\mathbf{e} \bmod q\rangle\,,$$

which is the inverse QFT of C $|\mathrm{SIS}\rangle$. Observe that the classical SIS problem reduces to C $|\mathrm{SIS}\rangle$ through a measurement (in the C $|\mathrm{SIS}\rangle$ problem, the length of the secret follows closely from the distribution χ), which in turn reduces to C $|\mathrm{LWE}\rangle$, and then further reduces to LWE naturally.

3 Algorithms for LWE, (E)DCP and S $|\mathrm{LWE}\rangle$

In this section, we discuss previously known algorithms for the LWE, (E)DCP, S $|\mathrm{LWE}\rangle$ problems (and more). We will also present some folklore algorithms that we believe are known within the community. All previous known algorithms, to the best of our knowledge, have at least subexponential time/sample complexity, or deal with specific marginal parameters.

3.1 Subexponential algorithms for DCP, EDCP and S $|\mathrm{LWE}\rangle$

We discuss subexponential algorithms – both in terms of time complexity and the number of samples required – for the DCP, EDCP and S $|\mathrm{LWE}\rangle$ problems.

Subexponential algorithms for DCP**.** No quantum polynomial time algorithm is known for the DCP problem. Kuperberg proposed several sieving-like quantum algorithms [42,43] that run in subexponential time and require a subexponential number of samples. These remain the best-known algorithms for solving DCP asymptotically. We give a brief introduction of the algorithm and refer to Appendix B for more details.

On input we have DCP states of the form $|0\rangle|x\rangle+|1\rangle|x+s \pmod N\rangle$, where $x \leftarrow_{\$} \mathbb{Z}_N$. A quantum Fourier transform and measurement result in the state:

$$|\psi_y\rangle \propto |0\rangle + \omega_N^{ys}|1\rangle, \tag{5}$$

where $y \leftarrow_{\$} \mathbb{Z}_N$ is classically known. Tensor two such states $|\psi_{y_1}\rangle, |\psi_{y_2}\rangle$, apply a CNOT gate, and then measure the second qubit, which gives:

$$|\psi_{y_1 \pm y_2}\rangle \propto |0\rangle + \omega_N^{(y_1 \pm y_2)s}|1\rangle,$$

up to a global phase, and each outcome occurs with probability 1/2. Now, note that the classical values of y_1 and y_2 are known, allowing us to target merging those y_1 and y_2 that share certain least significant bits (LSBs). This results in the state $|\psi_{y_1-y_2}\rangle$ with its LSBs zeroed.This combinatorial approach is central to all Kuperberg-like algorithms, which can be considered as a k-List algorithm [63]. The following theorem gives the time and samples complexity for the basic algorithm when N is a power of two. A variant algorithm [42] works for general N.

Theorem 4 (Kuperberg's sieve [42, Theorem 3.1]). *Let N be a power of two and $n = \log_2 N$. Kuperberg's sieve solves the* DCP_N *problem in time $\tilde{\mathcal{O}}(2^{3\sqrt{n}})$ that requires $\mathcal{O}(2^{3\sqrt{n}})$ samples.*

The original Kuperberg's algorithm [42] also requires subexponential quantum space. Regev [56] proposed an improvement that only uses a polynomial quantum space, at the cost of a slight increase in the running-time and the number of required samples. The core idea is to tensor a larger number of states at each step and actively solve for the solution, rather than relying solely on the randomness of samples to share the same LSBs. We describe its main idea. Let k, l be integer parameters where $l = \mathcal{O}(\sqrt{n \log n})$, $k = \mathcal{O}(\sqrt{n/\log n})$ and $n = k \cdot l + 1$. We tensor $(l+4)$ such states $\{|\psi_{y_i}\rangle\}_{i=1}^{l+4}$ from Equation (5) which produces a new state of the form $\sum_{\mathbf{j} \in \mathbb{Z}_2^{l+4}} \omega_N^{s \cdot \langle \mathbf{y}, \mathbf{j} \rangle} |\mathbf{j}\rangle$, where the vector $\mathbf{y} := (y_1, \cdots, y_{l+4})$. Since we know $\mathbf{y}$ classically, we can construct the following state:

$$\sum_{\mathbf{j} \in \mathbb{Z}_2^{l+4}} \omega_N^{s \cdot \langle \mathbf{y}, \mathbf{j} \rangle} |\mathbf{j}\rangle \, |\langle \mathbf{y}, \mathbf{j} \rangle \bmod 2^l\rangle.$$

We will then measure the second register and obtain some value t such that $\langle \mathbf{y}, \mathbf{j} \rangle \equiv t \pmod{2^l}$. One can find the solutions $\mathbf{j}$ by a classical brute-force enumeration, which takes $\mathcal{O}(2^l)$ time. Regev [56] showed that there is a small constant number of solutions, and we project to two such solutions $|\mathbf{j}_1\rangle, |\mathbf{j}_2\rangle$ via a

projective measurement. Thus we have the following states,

$$|0\rangle + \omega_N^{s\cdot\langle \mathbf{y},\, \mathbf{j}_2-\mathbf{j}_1\rangle}\,|1\rangle\,, \tag{6}$$

up to a global phase. Now it can be seen that $\langle \mathbf{y},\, \mathbf{j}_2-\mathbf{j}_1\rangle \equiv \mathbf{0} \pmod{2^l}$ and hence the l LSBs are "cleared" in the new state. The main difference, compared to [42], lies in solving a knapsack-like problem during the merging process. Consequently, this approach requires only a polynomial number of quantum samples to clear $\mathcal{O}(\sqrt{n\log n})$ bits at each merging step. The remainder of the algorithm follows similarly, proceeding through $\mathcal{O}(\sqrt{n/\log n})$ iterations. To conclude, Regev's algorithm runs in time $2^{\mathcal{O}(\sqrt{n\log n})}$, requires the same asymptotic number of samples, and uses $\mathcal{O}(n)$ quantum space.

Another work by Kuperberg [43] proposed improved algorithms based on this idea, which admit a better heuristic running time for general N and allow for balancing space and time. Asymptotically, these algorithms also run in subexponential time.

Subexponential algorithms for EDCP**.** Subexponential algorithms for EDCP problems are folklore. We focus on moduli q that are polynomial in n. Also, due to Lemma 7 & 8, we do not distinguish G-EDCP from U-EDCP. For example, one can combine the following G-EDCP $\leq$ DCP reduction with one of the aforementioned algorithms for DCP.

Lemma 10 (One-dimension G-EDCP $\leq$ DCP **[19, Lemma 11]).** *Let κ be security parameter. Let N, l be integers, and r be a positive real number. There exists a polynomial-time quantum reduction from* $\text{G-EDCP}^{l}_{1,N,r}$ *to* $\text{DCP}_N^{\mathcal{O}(l/(\log r\cdot\kappa^2))}$ *if $r \geq 3\log N$; and from* $\text{G-EDCP}^{l}_{1,N,r}$ *to* $\text{DCP}_N^{\mathcal{O}(l/(\log r\cdot\kappa))}$ *otherwise.*

One can solve 1-dim EDCP problem in subexponential time with respect to $\log N$ (we can do better when r is large, as discussed in Subsection 3.3). It is noted that one can reduce 1-dim G-EDCP to 1-dim LWE by Theorem 2 and then to n-dim LWE by the modulus-dimension switching [20, Theorem 3.1]. However, the reduction from n-dim LWE to n-dim G-EDCP in Theorem 3 only produces a polynomial number of samples.

Alternatively, one could use (part) of the reduction from the n-dim LWE to the DCP problem by Regev [58]. More precisely, Regev defines a Two-Point Problem ($\text{2PP}_{n,q}$) with some failure parameter δ. When $\delta \to \infty$, the $\text{2PP}_{n,q}$ problem converges to the $\text{U-EDCP}_{n,q,2}$ problem. Regev [58, Lemma 3.2] showed that the perfect $\text{2PP}_{n,q}$ problem (and essentially $\text{U-EDCP}_{n,q,2}$) reduces to the $\text{DCP}_{(2q)^n}$ problem. It is plausible to think that the n-dim secret in $\text{U-EDCP}_{n,q,2}$ admits similar entropy to the 1-dim secret (in a larger space) in $\text{DCP}_{(2q)^n}$.

Doliskani [31] gives an algorithm that works directly for U-EDCP with general support. Starting from a subexponential number of $\text{U-EDCP}_{n,q,M}$ samples, one can use Lemma 9 to get a subexponential number of $\text{U-EDCP}_{n,q,2}$ samples. Doliskani showed that one can use a Kuperberg-like procedure to produce DCP_q samples. The modulus q is polynomial, hence one can use the algorithm in [27] to solve it efficiently. The dominating part of the algorithm lies in the Kuperberg sieve step which is subexponential.

Subexponential algorithms for S$|$LWE$\rangle$**.** There also exists subexponential algorithms for G-S$|$LWE$\rangle$ – the S$|$LWE$\rangle$ problem with Gaussian distribution [23]. The algorithms use quantum rejection sampling to convert G-S$|$LWE$\rangle$ samples to samples of the form seen in Equation (6) (e.g., the QFT of U-EDCP$_{n,q,2}$ states), then use Kuperberg's sieve to solve it. We cite the main result.

Lemma 11 (Subexp. alg. for G-S$|$LWE$\rangle$ **[23, Corollary 34]).** *Let m, n, q be LWE parameters and c be a given (known) real number. Given $2^{\Theta(\sqrt{n \log q})}$ samples of the form:*

$$\sum_{e \in \mathbb{Z}} \rho_r(e) \exp(2\pi i c e/q) \, |\langle \mathbf{a}, \mathbf{s} \rangle + e \bmod q\rangle \,, \tag{7}$$

where the deviation $r = \Omega(\sqrt{n})$ and $r \leq q/\sqrt{n}$, there exists a quantum algorithm that finds $\mathbf{s} \in \mathbb{Z}_q^n$ in time $2^{\Theta(\sqrt{n \log q})}$.

In fact, the samples in Equation (7) are slightly more general than S$|$LWE$\rangle$ samples, since there is an additional parameter c (which is known). Recall that [25] shows that C$|$LWE$\rangle$ reduces to S$|$LWE$\rangle$ provided the oracle for the latter does not collapse the states. This reduction is incompatible with the above lemma so we do not know whether the same idea works for the C$|$LWE$\rangle$ problem. It is worth noting that all the algorithms discussed require perfect samples.

3.2 Polynomial algorithms for EDCP, S$|$LWE$\rangle$ and C$|$LWE$\rangle$ under specific parameters

Efficient quantum algorithms also exist for the EDCP, S$|$LWE$\rangle$ and C$|$LWE$\rangle$ problems, when the parameters or the distributions satisfy specific conditions.

Childs and van Dam [27] give a quantum algorithm for the 1-dimensional U-EDCP$_{1,2^n,M}$ problem. Their algorithm runs in polynomial-time, when $M = \Omega(2^{n/c})$ for any constant $c \geq 3$, given $l \geq c$ samples. Furthermore, the EDCP $\leq$ LWE reduction in Theorem 2 reduces EDCP$_{1,2^n,M}$ to LWE$_{1,2^n,1/M}$, which reduces to LWE$_{\sqrt{n},2^{\sqrt{n}},1/M}$ via modulus-dimension switching [20]. Using lattice reduction, the latter problem can be solved efficiently when $M = 2^{\Omega(\sqrt{n})}$.

Chen, Liu, and Zhandry present a quantum filtering algorithm [25] which relies on a quantum instantiation of the Arora-Ge algorithm [10]. Consequently, similar to the Arora-Ge algorithm, these algorithms run in polynomial-time for specific parameters. Concretely, they give several polynomial algorithms for EDCP, S$|$LWE$\rangle$ and C$|$LWE$\rangle$, under different sets of parameters.

Theorem 5 (Polynomial alg. [25, Theorems 32, 45]). *Let $q = \mathsf{poly}(n)$ and $\chi : \mathbb{Z}_q \to \mathbb{R}$ be the amplitude of the state $\sum_{e \in \mathbb{Z}_q} \chi(e) \, |e\rangle$ which can be efficiently constructed. If $\eta := \min_{y \in \mathbb{Z}_q} |\tilde{\chi}(y)|$ is non-negligible, then there exist polynomial time quantum algorithms that solve* S$|$LWE$\rangle_{n,q,\chi}^m$, C$|$LWE$\rangle_{n,q,\chi}^m$ *and* EDCP$_{n,q,\tilde{\chi}}^m$ *problems for $m \in \Omega(nq/\eta^2)$.*

The theorem does not cover the usual discrete Gaussian parameters, e.g., when $\chi = D_{\mathbb{Z},r}$ where $\sqrt{\kappa} \leq r \leq q/\sqrt{\kappa}$, since its DFT contains negligible density at the tail. However, it applies when χ is a bounded uniform distribution for the $\mathrm{S}\,|\mathrm{LWE}\rangle$ and $\mathrm{C}\,|\mathrm{LWE}\rangle$ problems, stated as follows:

Corollary 1 (Polynomial alg. for uniform dist. [25, Corollary 36]). *Let $q = \mathsf{poly}(n)$ and integer $B < (q-1)/2$ such that $\gcd(2B+1, q) = 1$. Let U be the discrete uniform distribution on $[-B, B] \cap \mathbb{Z}$. There exist polynomial time quantum algorithms that solve the $\mathrm{S}\,|\mathrm{LWE}\rangle^m_{n,q,U}$ and $\mathrm{C}\,|\mathrm{LWE}\rangle^m_{n,q,U}$ problems when $m \in \Omega(nq^4(2B+1))$.*

This corollary shows that the U-$\mathrm{S}\,|\mathrm{LWE}\rangle$ and U-$\mathrm{C}\,|\mathrm{LWE}\rangle$ problems are somewhat easy, conditioned on the parameters. One may also notice that the above lemma does not cover the EDCP problem. The following theorem works for the U-EDCP problem when the distribution has a sufficiently large support, e.g., almost as wide as q.

Theorem 6 (Polynomial alg. for wide uniform dist. [25, Theorems 39, 46]). *Let prime $q = \mathsf{poly}(n)$ and integer B such that $q - (2B+1) = c$ is a constant. Let U be the discrete uniform distribution on $[-B, B] \cap \mathbb{Z}$. There exist polynomial time quantum algorithms that solve $\mathrm{EDCP}^m_{n,q,U}$, $\mathrm{S}\,|\mathrm{LWE}\rangle^m_{n,q,\tilde{U}}$ and $\mathrm{C}\,|\mathrm{LWE}\rangle^m_{n,q,\tilde{U}}$ problems when $m \in \Omega((q-c)^3 n^{c+1} q \log q)$.*

It is noted that the EDCP parameters in this theorem have already been addressed by Ivanyos, Prakash, and Santha [37] using a quantum algorithm with similar complexity, which also relies on the Arora-Ge algorithm. The above theorem also applies to the $\mathrm{S}\,|\mathrm{LWE}\rangle$ and $\mathrm{C}\,|\mathrm{LWE}\rangle$ problems, with the distribution being $\tilde{U}$. Furthermore, these algorithms can be used to solve the $\mathrm{SIS}^\infty_{m,n,q,\beta}$ where $\beta = (q-c)/2$ using the SIS to $\mathrm{C}\,|\mathrm{LWE}\rangle$ reduction of Subsection 2.5.

3.3 Folklore algorithms

We describe several algorithms that, while not explicitly documented in the literature, are believed to be known folklore and may be of independent interest.

Polynomial samples algorithms for $(\mathrm{E})\,\mathrm{DCP}$**.** All algorithms discussed in Subsection 3.1 require at least a subexponential number of states as input. However, the number of required samples can be reduced to *polynomial*, by trading off running-time to *exponential*. Notice that Regev's algorithm (in Subsection 3.1) requires $l^{\mathcal{O}(k)}$ DCP states. One can choose the parameters l, k such that the algorithm requires only $\mathsf{poly}(n)$ DCP states. Concretely, setting $l = \Theta(n \log(n/\log n)/\log n)$ and $k = n/l = \Theta(\log n/\log(n/\log n))$, leads us to an algorithm that uses $\mathsf{poly}(n)$ DCP states and runs in time $\mathcal{O}(2^l) = \mathcal{O}(2^n)$, where the most expensive step is (classically) enumerating the solutions $\mathbf{j}$ to the system $\langle \mathbf{y}, \mathbf{j} \rangle \equiv t \pmod{2^l}$. Since an algorithm solving DCP [58,19] can also solve LWE via the existing reductions, a single-exponential algorithm for DCP is not surprising. For example, for LWE with $(1/\alpha, q) \in \mathsf{poly}(n)$, the best-known

algorithms are also single-exponential, using lattice reduction in Appendix A. Similarly, one can solve EDCP in exponential time with polynomial samples by considering (part) of the reduction from the n-dim LWE to the DCP problem in Regev [58] – which shows that U-EDCP$_{n,q,2}$ reduces to the DCP$_{(2q)^n}$, as discussed in Subsection 3.1.

Polynomial time algorithm for EDCP ***when*** $M = q$**.** When the U-EDCP$_{n,q,M}$ problem is defined with a uniform distribution over the full support $\mathbb{Z}_q$, there is a folklore $\mathsf{poly}(n)$-time algorithm that solves it [39]. Indeed, input a state $\sum_{j\in\mathbb{Z}_q} |j\rangle |\mathbf{x} + j \cdot \mathbf{s} \bmod q\rangle$ and then apply a QFT over $\mathbb{Z}_q \times \mathbb{Z}_q^n$ on both registers. The result is $\sum_{z\in\mathbb{Z}_q} \sum_{\mathbf{y}\in\mathbb{Z}_q^n} \omega_q^{\langle \mathbf{x}, \mathbf{y}\rangle} \sum_{j\in\mathbb{Z}_q} \omega_q^{j\cdot(\langle \mathbf{s}, \mathbf{y}\rangle + z)} |z\rangle |\mathbf{y}\rangle$, which is equivalent to $\sum_{\mathbf{y}\in\mathbb{Z}_q^n} |-\langle \mathbf{s}, \mathbf{y}\rangle\rangle |\mathbf{y}\rangle$. Measuring the states gives $(-\langle \mathbf{s}, \mathbf{y}\rangle, \mathbf{y})$. Repeating this process for $\Omega(n)$ times reveals $\mathbf{s}$ by Gaussian elimination.

Hardness of EDCP ***via reduction to*** LWE**.** The reduction given in [19] reduce n-dimensional G-EDCP$^l_{n,q,r}$ to LWE$^l_{n,q,\alpha=1/r}$, hence, under lattice reduction algorithms, G-EDCP$^l_{n,q,r}$ (and U-EDCP through the reduction from Lemma 8) has complexity $\exp(\mathcal{O}(n \log q \log n / \log^2(r)))$ for a sufficiently large l.

4 A quasi-polynomial time algorithm for EDCP

In this section, we present a quantum algorithm for solving U-EDCP with any support size M in quasi-polynomial time using a quasi-polynomial number of U-EDCP samples. To the best of our knowledge, this is the first quantum algorithm that runs below the subexponential regime, e.g., compared to the algorithms discussed in Subsection 3.1.

Theorem 7. *Let positive integers n, q, M and l be the parameters in* U-EDCP$^l_{n,q,M}$*, where q is a power-of-two. There exists a quantum algorithm that solves* U-EDCP$^l_{n,q,M}$ *in time $2^{\mathcal{O}(\log n \log q)}$ using $\mathsf{poly}(n)$ quantum space, when $l = 2^{\Omega(\log n \log q)}$.*

When q is polynomial (or at most quasi-polynomial) in n, which is typically the case, the running time and the number of samples required are both quasi-polynomial. In addition, this theorem naturally extends to G-EDCP via Lemma 7.

Corollary 2. *Let positive integers n, q, r and l be the parameters in* G-EDCP$^l_{n,q,r}$*, where q is a power-of-two. There exists a quantum algorithm that solves* G-EDCP$^l_{n,q,r}$ *in time $2^{\mathcal{O}(\log n \log q)}$ using $\mathsf{poly}(n)$ quantum space, when $l = 2^{\Omega(\log n \log q)}$.*

Our algorithm for U-EDCP directly addresses a variant problem called $\overline{\text{U-EDCP}}$, which is almost the QFT of U-EDCP. We first introduce the $\overline{\text{U-EDCP}}$ problem.

Definition 6 ($\overline{\text{U-EDCP}}$ problem). *The* $\overline{\text{U-EDCP}}^{l}_{n,q,M}$ *problem consists of* l *samples* $|\psi_k\rangle$ *of the form:*

$$|\psi_k\rangle \propto \sum_{j=0}^{M-1} \omega_q^{j\langle \mathbf{y}_k, \mathbf{s}\rangle} |j\rangle ,$$

where $k \in [l]$ *and* $\mathbf{y}_k \leftarrow\!\!\$\, \mathbb{Z}_q^n$ *are known classically, and asks to recover the secret* $\mathbf{s}$*.*

One can obtain $\overline{\text{U-EDCP}}$ samples by applying a QFT on the U-EDCP samples. Let $|\phi\rangle \propto \sum_{j=0}^{M-1} |j\rangle |\mathbf{x} + j \cdot \mathbf{s}\rangle$ be an input. Applying a QFT on the second register gives, up to a global phase, the following state:

$$(\mathbb{1} \otimes \text{QFT}_q^n) |\phi\rangle \propto \sum_{\mathbf{y} \in \mathbb{Z}_q^n} \sum_{j=0}^{M-1} \omega_q^{j\langle \mathbf{y}, \mathbf{s}\rangle} |j\rangle |\mathbf{y}\rangle .$$

Then measuring the second register gives an $\overline{\text{U-EDCP}}$ sample, where $\mathbf{y} \leftarrow\!\!\$\, \mathbb{Z}_q^n$ is known classically. Therefore, U-EDCP naturally reduces to $\overline{\text{U-EDCP}}$, and we will focus on algorithms for the latter.

4.1 Algorithm for U-EDCP with $M = 2$

We begin by describing our algorithm for $\overline{\text{U-EDCP}}_{n,q,2}$ (and hence $\text{U-EDCP}_{n,q,2}$) and then rely on a self-reduction to handle the general case of $\overline{\text{U-EDCP}}_{n,q,M}$. Our proof relies on an iterative application of the following merging step. In this work, we assume the underlying field operations are efficient and we focus on the arithmetic complexity.

Lemma 12. *Let* $p \geq 4$ *be a power-of-two. Given* $(n+1)$ *samples* $\{|\psi_k\rangle\}_{k=1}^{n+1}$ *from* $\overline{\text{U-EDCP}}_{n,p,2}$ *with known* $\mathbf{y}_k \leftarrow\!\!\$\, \mathbb{Z}_p^n$*, there exists a quantum polynomial-time algorithm that outputs one sample from* $\overline{\text{U-EDCP}}_{n,p/2,2}$ *with the same secret* $\mathbf{s}$*, succeeding with constant probability* $C \geq 0.577$*.*

Proof. By tensoring our initial $n+1$ $\overline{\text{U-EDCP}}$ samples, we can create the state:

$$\bigotimes_{k=1}^{n+1} |\psi_k\rangle \propto \sum_{\mathbf{j} \in \mathbb{Z}_2^{n+1}} \omega_p^{\langle \mathbf{Y}\cdot\mathbf{j}, \mathbf{s}\rangle} |\mathbf{j}\rangle , \tag{8}$$

where $\mathbf{Y} := (\mathbf{y}_1, \ldots, \mathbf{y}_{n+1}) \in \mathbb{Z}_p^{n \times (n+1)}$ is a column matrix formed by the $\mathbf{y}_k$'s. In a new register, we compute:

$$\sum_{\mathbf{j} \in \mathbb{Z}_2^{n+1}} \omega_p^{\langle \mathbf{Y}\cdot\mathbf{j}, \mathbf{s}\rangle} |\mathbf{j}\rangle |\mathbf{Y} \cdot \mathbf{j} \bmod 2\rangle ,$$

and then measure it to get some $\mathbf{t} \in \mathbb{Z}_2^n$ such that $\mathbf{Y} \cdot \mathbf{j} \equiv \mathbf{t} \pmod 2$. By definition, the input $\mathbf{Y}$ is uniform in $\mathbb{Z}_p$. Since $2 \mid p$, we see that $\mathbf{Y} \pmod 2$ is

uniform. By Lemma 2 it has full rank modulo 2 with constant probability. As $\mathbf{Y}$ is an $n \times (n+1)$ matrix, there are exactly two solutions $\{\mathbf{x}_0, \mathbf{x}_1\} \subseteq \mathbb{Z}_2^{n+1}$ for $\mathbf{Y}\mathbf{x} \equiv \mathbf{t} \pmod 2$. Classically, we can compute the two solutions using Gaussian elimination in time $\mathcal{O}(n^3)$. Therefore, we are left with the state:

$$\omega_p^{\langle \mathbf{Y}\cdot\mathbf{x}_0,\, \mathbf{s}\rangle} |\mathbf{x}_0\rangle + \omega_p^{\langle \mathbf{Y}\cdot\mathbf{x}_1,\, \mathbf{s}\rangle} |\mathbf{x}_1\rangle .$$

By factoring out the global phase and renaming, we obtain the state:

$$|0\rangle + \omega_p^{\langle \mathbf{y}',\, \mathbf{s}\rangle} |1\rangle$$

where the new vector $\mathbf{y}' := \mathbf{Y} \cdot (\mathbf{x}_1 - \mathbf{x}_0) \pmod p$ is even. Choose any $\mathbf{y}''$ such that $2\mathbf{y}'' = \mathbf{y}' \bmod p$. Since 2 divides p there is no unique $\mathbf{y}'' \bmod p$, but we can define it as the result of lifting $\mathbf{y}' \bmod p$ to $\mathbb{Z}^n$ and dividing each entry by 2.

Now since $\omega_{p/2}^{\langle \mathbf{y}'',\, \mathbf{s}\rangle} = \omega_{p/2}^{\langle \mathbf{y}'',\, \mathbf{s}\rangle \bmod p/2}$ and $\mathbf{y}''$ *is* unique modulo $p/2$, we have the state

$$|0\rangle + \omega_p^{2\langle \mathbf{y}'',\, \mathbf{s}\rangle} |1\rangle = |0\rangle + \omega_{p/2}^{\langle \mathbf{y}'',\, \mathbf{s}\rangle} |1\rangle$$

for a known vector $\mathbf{y}''$. By Lemma 13 it is uniform in $\mathbb{Z}_{p/2}^n$, so the above is a valid $\overline{\text{U-EDCP}}_{n,p/2,2}$ sample, with the original secret. □

Abusing notation, we will denote this lemma $\overline{\text{U-EDCP}}_{n,p,2}^{n+1} \to \overline{\text{U-EDCP}}_{n,p/2,2}^{1}$. This notation can be interpreted to mean that there exists an efficient quantum procedure that converts $(n+1)$ samples of the former problem into a single sample of the latter problem, sharing the same secret. Note that the idea of tensoring more than two samples has been used in Regev's polynomial-space algorithm [56] for Kuperberg's sieve, as well as in [39]. Unlike the previous approach, we process the vector in a breadth-first manner, handling a single bit for all coordinates, which results in only a polynomial cost for each merge. The following lemma shows the uniformity of the merged vectors in the above procedure.

Lemma 13. *Let $p \geq 4$ be a power of two. We consider the following procedure, which simulates the merging step in Equation (8) classically.*

- *sample $\mathbf{Y} \leftarrow_\$ \mathbb{Z}_p^{n\times(n+1)}$ until it has full rank over $\mathbb{Z}_2$ and $\mathbf{x}_0 \leftarrow_\$ \mathbb{Z}_2^{n+1}$ independently;*
- *let $\mathbf{x}_1 \in \mathbb{Z}_2^{n+1} \setminus \{\mathbf{x}_0\}$ denote the other solution to $\mathbf{Y}\mathbf{x}_0 = \mathbf{Y}\mathbf{x}_1 \pmod 2$;*
- *output $\mathbf{y}' := \mathbf{Y} \cdot (\mathbf{x}_1 - \mathbf{x}_0) \pmod p$, which is even.*

Then the distribution of $\mathbf{y}'$, induced from the input randomness on $\mathbf{Y}$ and $\mathbf{x}_0$, is uniform in $2\mathbb{Z}_{p/2}^n$.

Proof. We use a counting argument to prove this. It suffices to show that for any fixed $\mathbf{y}'$ and $\mathbf{x}_0$, the number of matrices $\mathbf{Y} \in \mathbb{Z}_p^{n\times(n+1)}$ with full rank modulo 2 that produce the given $\mathbf{y}'$, conditioned on the given $\mathbf{x}_0$, is a constant. Noting that the two sources of randomness, $\mathbf{Y}$ and $\mathbf{x}_0$, are uniformly independent, the proof is then completed by an averaging argument over $\mathbf{x}_0$. Now we let $\mathbf{y}' \in 2\mathbb{Z}_{p/2}^n$ and $\mathbf{x}_0 \in \mathbb{Z}_2^n$ be fixed. We count the total number of $\mathbf{Y}$ that produces the given $\mathbf{y}'$.

Given $\mathbf{y}'$ and $\mathbf{x}_0$, we sum over all possible $\mathbf{x}_1 \in \mathbb{Z}_2^{n+1} \setminus \{\mathbf{x}_0\}$ and count the number of $\mathbf{Y}$'s. Denote $\mathbf{x}' := (\mathbf{x}_1 - \mathbf{x}_0)$ which is non-zero and has entries in $\{-1, 0, 1\}$. We claim that, for any non-zero vector $\mathbf{x}' \in \{-1, 0, 1\}^n$ and any given even $\mathbf{y}'$, the number of $\mathbf{Y}$ in $\mathbb{Z}_p^{n\times(n+1)}$ with full rank modulo 2 such that $\mathbf{Y} \cdot \mathbf{x}' \equiv \mathbf{y}' \pmod p$ is precisely $|\mathrm{GL}_n(\mathbb{Z}_p)|$.

To see this, assume w.l.o.g that the last coordinate of $\mathbf{x}'$ is 1 (or -1). We enumerate over all possible left $n \times n$ sub-matrices $\mathbf{Y}_l$ of $\mathbf{Y}$, where we denote $\mathbf{Y} = [\mathbf{Y}_l \mid \mathbf{y}_r]$. For each such fixed $\mathbf{Y}_l$, the right-hand side vector $\mathbf{y}_r$ in $\mathbf{Y}$ is uniquely determined since the last coordinate of $\mathbf{x}'$ is invertible in $\mathbb{Z}_2$. Furthermore, we claim it is sufficient to count only *invertible* square sub-matrices $\mathbf{Y}_l$. Assume $\mathbf{Y}_l$ is non-invertible. Then it will have determinant a multiple of 2, so $\mathbf{Y}_l$ modulo 2 is also non-invertible and does not have full rank. Then the last row of its reduced echelon form is zero. It follows the last entry of $\mathbf{y}_r$ must also be 0 modulo 2, since $\mathbf{y}'$ is even and the last entry of $\mathbf{x}'$ is 1 or -1. Thus $\mathbf{Y}$ can not have full rank modulo 2 and the total number of such $\mathbf{Y}$ is $|\mathrm{GL}_n(\mathbb{Z}_p)|$.

Now we let $\mathbf{x}_1$ vary over $\mathbb{Z}_2^{n+1} \setminus \{\mathbf{x}_0\}$. A simple computation shows there is no overlap between the $\mathbf{Y}$'s associated with two distinct $\mathbf{x}_1$'s. Hence for a fixed $(\mathbf{y}', \mathbf{x}_0)$ there are $|\,\mathrm{GL}_n(\mathbb{Z}_p)| \cdot (2^{n+1} - 1)$ full rank matrices $\mathbf{Y}$ such that $\mathbf{Y} \cdot (\mathbf{x}_0 - \mathbf{x}_1) = \mathbf{y}' \pmod p$. This number is independent of the choice of $(\mathbf{y}', \mathbf{x}_0)$. It follows that $\mathbf{y}'$ is uniform in $2\mathbb{Z}_{p/2}^n$.

□

Remark 1. As discussed in Subsection 1.2, our merging and lifting procedures used in the proof of Lemma 12 essentially follow the "Simon-meets-Kuperberg" algorithm of Bonnetain and Naya-Plasencia [18,16], originally proposed in the context of symmetric cryptanalysis. In Lemma 13 and Subsection 4.2, we provide a rigorous analysis that applies to EDCP instances with general M.

Now, we are ready to prove the main theorem. We first present the proof of Theorem 7 for the case $M = 2$.

Proof of Theorem 7 (for $\mathbf{M = 2}$). Let $q = 2^t$ be the modulus of the given $\overline{\text{U-EDCP}}_{n,q,2}$ problem. Define $q_i = q/2^i$ for $i \in [t]$. Following Lemma 12 we consider the following chain of combinations:

$$\overline{\text{U-EDCP}}^{l_0}_{n,q_0=q,2} \to \cdots \to \overline{\text{U-EDCP}}^{l_i}_{n,q_i,2} \to \cdots \to \overline{\text{U-EDCP}}^{l_{t-1}=n}_{n,q_{t-1}=2,2}. \tag{9}$$

Looking ahead, assume that n samples have been obtained for the final $\overline{\text{U-EDCP}}_{n,2,2}$ problem. The resulting samples are of the form:

$$\left\{ |0\rangle + (-1)^{\langle \mathbf{y}_k,\, \mathbf{s}\rangle} |1\rangle \right\}_{k=1}^{n}.$$

By measuring each in the Hadamard basis we learn $\langle \mathbf{y}_k,\, \mathbf{s}\rangle \pmod 2, \forall k \in [n]$. This allows us to recover $\bar{\mathbf{s}} = \mathbf{s} \pmod 2$ as soon as the linear system is non-singular. By Lemma 12 the vectors $\mathbf{y}_k$ are uniform in $\mathbb{Z}_2^n$, so by Lemma 2 this happens with constant probability. This procedure recovers the first n bits of $\mathbf{s}$. To recover the next n bits, we process fresh $\text{EDCP}_{n,q,2}$ samples by mapping

$$|j\rangle\, |\mathbf{x} + j \cdot \mathbf{s} \pmod q\rangle \to |j\rangle\, |\mathbf{x} + j \cdot \mathbf{s} - j \cdot \bar{\mathbf{s}} \pmod q\rangle\,.$$

After doing a QFT over $\mathbb{Z}_q^n$ on the second register and measuring, we obtain $\overline{\text{U-EDCP}}_{n,q/2,2}$ states of the form:

$$|0\rangle + \omega_q^{\langle \mathbf{y}_k,\, \mathbf{s}-\bar{\mathbf{s}}\rangle}\,|1\rangle = |0\rangle + \omega_{q/2}^{\langle \mathbf{y}_k,\, (\mathbf{s}-\bar{\mathbf{s}})/2\rangle}\,|1\rangle\,.$$

Note that $\mathbf{s}-\bar{\mathbf{s}} \pmod q$ is even, so the equality is justified by the same argument used in the proof of Lemma 12. One can now apply the above chain (partially) in a similar manner to learn the next n bits of $\mathbf{s}$. Proceeding in this way, $\log q$ iterations of this procedure are sufficient to recover the entire secret.

Now we analyze the complexity. It is sufficient to focus on recovering the first n bits of $\mathbf{s}$. Consider a single step $\overline{\text{U-EDCP}}^{l_i}_{n,q_i,2} \to \overline{\text{U-EDCP}}^{l_{i+1}}_{n,q_{i+1},2}$ of the chain in Equation (9). By a Chernoff bound argument, on input $\Omega(l_i)$ samples this step outputs at least $l_i/(n+1)$ samples with overwhelming probability. Since the chain of merges in Equation (9) has length $\log q - 1$, on input $\Omega(q \cdot n^{\log q}) = 2^{\Omega(\log n \log q)}$ samples the procedure outputs n samples with overwhelming probability. Finally, the algorithm processes $\mathcal{O}(n)$ samples at a time so it requires $\mathsf{poly}(n)$ quantum space, following the "pipeline" analogy of Regev [56]. □

4.2 Algorithm for U-EDCP with general M

For general M, we will rely on a self-reduction of $\overline{\text{U-EDCP}}$. Two self-reductions for the U-EDCP problem already exist in the literature. The first, given in [19, Lemma 10], relies on the quantum rejection sampling from [51, Sec.4]. Qualitatively, reducing $\text{U-EDCP}_{n,q,M}$ to $\text{U-EDCP}_{n,q,M'}$ for some $M' < M$ with this approach reduces the number of samples by a factor of $O(M/M')$ and requires the ratio M/M' to be polynomially bounded. The second self-reduction is Lemma 9, which loses half of the samples on average and places no restriction on the ratio M/M'. For our application, we use a simpler reduction specific to $\overline{\text{U-EDCP}}$, which may also be of independent interest. Concretely, we reduce $\overline{\text{U-EDCP}}_{n,q,M}$ to $\overline{\text{U-EDCP}}_{n,q,M'}$ which is lossless when M' divides M, and loses half of the samples on average otherwise. The proof is similar in nature to that of Lemma 9.

Lemma 14 ($\overline{\text{U-EDCP}}$ self-reduction). *Let n, q, l, l' be integers greater than 1 and M, M' integers at most q such that $M' < M$. Then there is a probabilistic polynomial time reduction from $\overline{\text{U-EDCP}}^{l}_{n,q,M}$ to $\overline{\text{U-EDCP}}^{l'}_{n,q,M'}$ with $l' = \Theta(l)$. Furthermore, if M' divides M then $l' = l$.*

Proof. Given an $\overline{\text{U-EDCP}}_{n,q,M}$ state:

$$\sum_{j=0}^{M-1} \omega_q^{j\langle \mathbf{y},\, \mathbf{s}\rangle}\,|j\rangle$$

compute $\lfloor j/M' \rfloor$ in a new register and measure it as k. Now let $k := \lfloor j/M' \rfloor$ be fixed. Note that $\lfloor j/M' \rfloor = k$ if and only if $j \in [M'k, M'(k+1))$, so we are left

with the state:

$$\sum_{j\in[0,M)\cap[M'k,M'(k+1))} \omega_q^{j\langle \mathbf{y},\,\mathbf{s}\rangle}\,|j\rangle\,.$$

A simple calculation shows that if $j < M'\lfloor M/M'\rfloor$ then $M'(k+1) \le M$. With probability $(M'\lfloor M/M'\rfloor)/M$, we are left with the state:

$$\sum_{j=M'k}^{M'(k+1)-1} \omega_q^{j\langle \mathbf{y},\,\mathbf{s}\rangle}\,|j\rangle\,.$$

Now factor out the common phase $\omega_q^{(M'k)\,\langle \mathbf{y},\,\mathbf{s}\rangle}$ and map $|j\rangle \mapsto |j - M'k\rangle$ to get:

$$\sum_{j=M'k}^{M'(k+1)-1} \omega_q^{(j-M'k)\langle \mathbf{y},\,\mathbf{s}\rangle}\,|j - M'k\rangle = \sum_{j=0}^{M'-1} \omega_q^{j\langle \mathbf{y},\,\mathbf{s}\rangle}\,|j\rangle\,.$$

Now we consider the success probability $(M'\lfloor M/M'\rfloor)/M$. If M' divides M, this is exactly 1. Otherwise, we consider $M' > M/2$ and $M' < M/2$ separately. When $M > M' > M/2$, we have $1 < M/M' < 2$, so $\lfloor M/M'\rfloor = 1$ and the success probability is at least $1/2$. When $M' < M/2$ then we have $(M'\lfloor M/M'\rfloor)/M > 1 - M'/M$ which is also at least $1/2$. □

Now we complete the proof of Theorem 7.

Proof of Theorem 7 (for general* M*). Follows immediately from Lemma 14 and Theorem 7 with $M = 2$. □

4.3 A failed generalization of our algorithm

A natural generalization would be an algorithm that operates directly on $\overline{\text{U-EDCP}}$ samples modulo $M > 2$, bypassing the aforementioned self-reduction. If possible, this approach could asymptotically reduce both the length of the reduction chain and the number of required samples. In particular, we could set $M = \mathsf{poly}(n)$, where $t = \log_M q$ is a constant. This would lead to a $\mathsf{poly}(n)$-time algorithm for $\overline{\text{U-EDCP}}$ (e.g., this would break standard LWE quantumly). We discuss why such generalization *fails*.

Concretely, let us try to generalize the proof of Lemma 12 by expressing $q = M^t$ for $M > 2$. We start with $\overline{\text{U-EDCP}}_{n,q,M}$ states as in Definition 6. Following the proof Lemma 12, we tensor $(n+1)$ states to have:

$$\bigotimes_{k=1}^{n+1} |\psi_k\rangle = \sum_{\mathbf{j}\in\mathbb{Z}_M^{n+1}} \omega_q^{\langle \mathbf{Y}\cdot\mathbf{j},\,\mathbf{s}\rangle}\,|\mathbf{j}\rangle\,,$$

where $\mathbf{Y} = (\mathbf{y}_1, \ldots, \mathbf{y}_{n+1}) \in \mathbb{Z}_q^{n\times(n+1)}$. In a new register compute $\mathbf{Y}\cdot\mathbf{j} \pmod M$ and measure it as $\mathbf{t} \in \mathbb{Z}_M^n$. Classically, in polynomial time we compute the set $X = \{\mathbf{x} \in \mathbb{Z}_M^{n+1} \mid \mathbf{Y}\cdot\mathbf{x} = \mathbf{t} \pmod M\}$. Assuming $\mathbf{Y}$ has full rank, the set of

all solutions $\mathbf{x}_i \in X$ can be written as $\mathbf{x}_i \equiv \mathbf{x}_0 + i \cdot \mathbf{u} \pmod{M}$ for $0 \leq i < M$, where $\mathbf{x}_0$ is any fixed solution and $\mathbf{u} \in \mathrm{Ker}(\mathbf{Y})$ is any fixed non-zero element in the kernel. Such an expression, seemingly, reproduces an EDCP-like sample. However, notice that[4] $x_i \equiv \mathbf{x}_0 + i \cdot \mathbf{u}$ holds $\pmod{M}$ but not necessarily $\pmod{q}$. In fact, they are of the form $\mathbf{x}_0 + i \cdot \mathbf{u} + M \cdot \mathbf{v}_i$ for some $\mathbf{v}_i$. Because of these different $\mathbf{v}_i$'s, factoring out the "global" phase fails:

$$\sum_{i=0}^{M-1} \omega_q^{\langle \mathbf{Y}\cdot\mathbf{x}_i,\, \mathbf{s}\rangle} |\mathbf{x}_i\rangle = \sum_{i=0}^{M-1} \omega_q^{\langle \mathbf{Y}\cdot(\mathbf{x}_0 + i\cdot\mathbf{u} + M\cdot\mathbf{v}_i),\, \mathbf{s}\rangle} |\mathbf{x}_i\rangle \not\propto \sum_{i=0}^{M-1} \omega_q^{i\langle \mathbf{Y}\cdot\mathbf{u},\, \mathbf{s}\rangle} |\mathbf{x}_i\rangle .$$

Therefore, it does not seem to be possible to "craft" the set X such that the above procedure gives us the desired state $\sum_{i=0}^{M-1} \omega_q^{i\langle \mathbf{Y}\cdot\mathbf{u},\, \mathbf{s}\rangle} |\mathbf{x}_i\rangle = \sum_{i=0}^{M-1} \omega_{q/M}^{j\langle \mathbf{y}',\, \mathbf{s}\rangle} |\mathbf{x}_i\rangle$.

5 Algorithms for EDCP with specific parameters

In Section 3 we discussed Theorem 5 and Corollary 1, polynomial time algorithms for EDCP, $\mathrm{S}\,|\mathrm{LWE}\rangle$ and $\mathrm{C}\,|\mathrm{LWE}\rangle$ with specific parameters presented by [25]. Both of these results rely on a technique they call quantum filtering, which uses a quantum version of the Arora-Ge algorithm. Corollary 1 gives a polynomial time algorithm for $\text{U-EDCP}_{n,q,M}$ when $M = q - c$, and Theorem 5 allows to solve $\mathrm{EDCP}_{n,q,f}$ in polynomial time when $\min_{z\in\mathbb{Z}_q} |f(z)|$ is non-negligble. Both algorithms require a polynomial number of samples.

First, we show that a modified version of our algorithm allows to solve the harder instance of U-EDCP for $M = q/c$ in polynomial time, providing an asymptotic improvement over Corollary 1. Note that we do not require q to be a power of two here, only that c is a power of two.

Theorem 8. *Let $n \in \mathbb{N}$ and $q = \mathsf{poly}(n)$. Let c be a constant power of two dividing q. Let $M = q/c$ and $l = \Omega(n^{\log c} \log q)$. There is a polynomial time algorithm for* $\text{U-EDCP}^{l}_{n,q,M}$.

We sketch the proof ideas: When M divides q, Lemma 12 (and Lemma 13) can be modified to transform $(n+1)$ samples of $\overline{\text{U-EDCP}}_{n,q,M}$ to a single sample of $\overline{\text{U-EDCP}}_{n,q/M=c,2}$. This can be done using the ideas discussed in Section 4.3 followed by a projective measurement onto two states, which succeeds with probability $1/q$.

Assume we have $\Omega(n^{\log c})$ samples of $\overline{\text{U-EDCP}}_{n,c,2}$. Then we can apply Theorem 7 with modulus c to recover n bits of the secret in time $\mathsf{poly}(n)$. By a Chernoff bound, $\Omega(n^{\log c})$ samples are sufficient to succeed with constant probability. Thus to recover all $n \log q$ bits of the secret we need $l = \Omega(n^{\log c} \log q)$.

Next we give a simple proof of a variation of Theorem 5 which assumes that f takes on only positive values (which is the case whenever f is a distribution). Using our algorithm as well as rejection sampling (Lemma 6) we then extend

[4] We thank Damien Stehlé for pointing this out.

this result to a quasi-polynomial time algorithm for $\mathrm{EDCP}_{n,q,f}$ when f is non-negligible for at least two points. Neither of these results place any restriction on the divisors of q. In the following we will assume that samples of $\mathrm{EDCP}_{n,q,f}$ are normalized such that $\sum_j f(j)^2 = 1$.

Theorem 9. *Let $n \in \mathbb{N}$ and $q = \mathsf{poly}(n)$. Let $f : \mathbb{Z}_q \to \mathbb{R}^+$ be such that $\eta := \min_{z \in \mathbb{Z}_q} f(z)$ is non-negligible. Let $l = \Omega(n/(q\eta^2))$. There is a polynomial time algorithm for $\mathrm{EDCP}^l_{n,q,f}$.*

Proof. We use quantum rejection sampling as stated in Lemma 6 to reduce the problem to U-EDCP with $M = q$. Using the notation of Lemma 6 we set $\pi_j := f(j)$ and $\sigma_j := \eta$, so that $\sigma_j \leq \pi_j$. Then rejection sampling produces:

$$(1/B) \sum_{j \in \mathbb{Z}_q} \eta \, |j\rangle \, |\mathbf{x} + j \cdot \mathbf{s}\rangle$$

with probability $B^2 = q\eta^2$. This is proportional to a U-EDCP state with $M = q$. As discussed in Section 3.3 only $\Omega(n)$ such states are needed to recover the secret, so $l = \Omega(n/(q\eta^2))$.

□

Theorem 10. *Let $n \in \mathbb{N}$ and a power-of-two modulus $q = \mathsf{poly}(n)$. Let $f : \mathbb{Z}_q \to \mathbb{R}^+$ be such that f is non-negligible for at least two points $z_0, z_1 \in \mathbb{Z}_q$ which are known. Set $\eta := \min\{f(z_0), f(z_1)\}$ and let $l = 2^{\Omega(\log n \log q)}$. There is a quasi-polynomial time algorithm for $\mathrm{EDCP}^l_{n,q,f}$.*

Proof. For each input state we use rejection sampling as in Lemma 6 with $\pi_{z_i} := f(z_i)$ for $i \in [q]$, $\sigma_{z_i} := \eta$ for $i \in \{0, 1\}$ and $\sigma_{z_i} := 0$ otherwise. This results in states of the form $|z_0\rangle \, |\mathbf{x}_k + z_0 \cdot \mathbf{s}\rangle + |z_1\rangle \, |\mathbf{x}_k + z_1 \cdot \mathbf{s}\rangle$ with non-negligible probability $2\eta^2$. Doing a QFT over $\mathbb{Z}_q^n$ on the second register and relabeling yields a $\overline{\mathrm{U\text{-}EDCP}}_{n,q,2}$ state. With $l = 2^{\Omega(\log n \log q)}$ such states we can recover the secret by Theorem 7. □

Alternative EDCP algorithm by Arora-Ge. As alluded to in Section 1 it is possible to solve $\mathrm{EDCP}^l_{n,q,M}$ by first reducing the problem to $\mathrm{LWE}^l_{n,q,1/M}$ (e.g., Theorem 2) and applying the Arora-Ge algorithm, which has complexity $2^{\tilde{O}((\alpha q)^2)}$ and requires $2^{\tilde{O}((\alpha q)^2)}$ samples. When $\alpha q = q/M = c$ this complexity is $\mathsf{poly}(n)$, so Arora-Ge can also be used to solve $\mathrm{EDCP}^l_{n,q,q/c}$ in polynomial time when $l = \mathsf{poly}(n)$. We consider this approach with other parameters when comparing algorithms for EDCP in Figure 1.

6 A quasi-polynomial time algorithm for S $|\mathrm{LWE}\rangle$

In this section, we present a quasi-polynomial time quantum algorithm for solving the Gaussian-S $|\mathrm{LWE}\rangle$ problem, leveraging our EDCP algorithm from Section 4. To the best of our knowledge, this appears to be the first algorithm that runs below the subexponential regime. By comparison, the previously best-known algorithm for G-S $|\mathrm{LWE}\rangle$ is the Kuperberg-like algorithm in Lemma 11.

Theorem 11. *Let κ be the security parameter and $n, q = \mathsf{poly}(\kappa)$ be integers, where q is a power-of-two. Let $r = \Omega(\sqrt{\kappa})$ and $q/r = \Omega(\sqrt{\kappa})$. There exists a quantum algorithm that solves $\mathrm{S}\,|\mathrm{LWE}\rangle^{l}_{n,q,r}$ in time $2^{\mathcal{O}(\log^2 n)}$ using $\mathsf{poly}(n)$ quantum space, when $l = 2^{\Omega(\log^2 n)}$.*

Proof. The proof is a combination of Lemma 15, 16 and Theorem 7. □

Our algorithm will use a variant problem called $\overline{\text{G-EDCP}}$, defined similarly to $\overline{\text{U-EDCP}}$ in Definition 6.

Definition 7 ($\overline{\text{G-EDCP}}$ problem). *The $\overline{\text{G-EDCP}}^{l}_{n,q,r}$ problem consists of l samples $|\psi_k\rangle$ of the form:*

$$|\psi_k\rangle \propto \sum_{j\in\mathbb{Z}_q} \rho_r(j)\, \omega_q^{j\cdot\langle \mathbf{y}_k,\, \mathbf{s}\rangle}\, |j\rangle\,,$$

where $k \in [l]$ and $\mathbf{y}_k \leftarrow\!\!\$\, \mathbb{Z}_q^n$ are known classically, and asks to recover the secret $\mathbf{s}$.

One can obtain $\overline{\text{G-EDCP}}$ samples by applying a QFT on the G-EDCP samples and then measure. This is similar to the $\overline{\text{U-EDCP}}$ of Definition 6 so we omit the details. We usually take $r \leq q/\sqrt{\kappa}$ to ensure that the support in the definition is predominantly contained in $\mathbb{Z}_q$. To prove the theorem, we first show a reduction $\text{G-S}\,|\text{LWE}\rangle \leq \overline{\text{G-EDCP}}$ and then a reduction $\overline{\text{G-EDCP}} \leq \overline{\text{U-EDCP}}$. Therefore, our algorithm for $\overline{\text{U-EDCP}}$ implies an algorithm for $\text{G-S}\,|\text{LWE}\rangle$.

6.1 G-S $|$LWE$\rangle \leq \overline{\text{G-EDCP}}$

We begin by describing a polynomial-time reduction from $\text{G-S}\,|\text{LWE}\rangle$ to $\overline{\text{G-EDCP}}$. The proof is conceptually similar to a part of the reduction from EDCP to LWE of [19, Theorem 4], but reversed.

Lemma 15 ($\text{G-S}\,|\text{LWE}\rangle < \overline{\text{G-EDCP}}$). *Let κ be the security parameter and $n, q = \mathsf{poly}(\kappa)$ be integers. Let $r = \Omega(\sqrt{\kappa})$ and $q/r = \Omega(\sqrt{\kappa})$. There exists a quantum polynomial-time reduction from $\mathrm{S}\,|\mathrm{LWE}\rangle^{l}_{n,q,r}$ to $\overline{\text{G-EDCP}}^{l}_{n,q,\sigma}$, where $\sigma = q/r$, admitting the same secret $\mathbf{s}$.*

Proof. We are given $\text{G-S}\,|\text{LWE}\rangle$ samples of the form $\sum_{e\in\mathbb{Z}_q} \rho_r(e)\,|\langle \mathbf{a},\, \mathbf{s}\rangle + e \pmod q\rangle$ which is close to:

$$\sum_{e\in\mathbb{Z}} \rho_{r/q}(e/q)\,|\langle \mathbf{a},\, \mathbf{s}\rangle + e \pmod q\rangle$$

using the tail bound from Lemma 4. Then we change the variable $b := \langle \mathbf{a},\, \mathbf{s}\rangle + e$ and split the summation of $b \in \mathbb{Z}$ into a double summation of $z \in \mathbb{Z}_q$ and $j \in \mathbb{Z}$. This leads to:

$$\sum_{z\in\mathbb{Z}_q} \Big(\sum_{j\in\mathbb{Z}} \rho_{r/q}(j + (z - \langle \mathbf{a},\, \mathbf{s}\rangle)/q)\Big)\,|z\rangle\,.$$

Using the Poisson summation formula of Theorem 1, the above state is proportional to:

$$\sum_{z\in\mathbb{Z}_q}\Big(\sum_{j\in\mathbb{Z}}\omega_q^{j\cdot(-\langle \mathbf{a},\,\mathbf{s}\rangle+z)}\rho_{q/r}(j)\Big)\,|z\rangle\,.$$

Since $q=\mathsf{poly}(\kappa)$ and $r=\Omega(\sqrt{\kappa})$, the above state is close to:

$$\sum_{z\in\mathbb{Z}_q}\Big(\sum_{j\in\mathbb{Z}_q}\omega_q^{j\cdot(-\langle \mathbf{a},\,\mathbf{s}\rangle+z)}\rho_{q/r}(j)\Big)|z\rangle\,.$$

Now we do an inverse QFT over $\mathbb{Z}_q$ and denote $\mathbf{a}'=-\mathbf{a}$, which gives:

$$\sum_{j\in\mathbb{Z}_q}\omega_q^{j\cdot\langle \mathbf{a}',\,\mathbf{s}\rangle}\rho_{q/r}(j)\,|j\rangle\,.$$

□

Note that $\mathbf{a}'$ is known classically and is uniform in $\mathbb{Z}_q^n$. Let A be the normalization scalar of the final state, e.g., $A^2\cdot\sum_{j\in\mathbb{Z}}\rho_{q/r}^2(j)=1$. Using Lemma 5, we get $A\approx 2^{1/4}\cdot\sqrt{r/q}$ up to an error factor of $2^{-\Omega(\kappa)}$.

6.2 $\overline{\text{G-EDCP}}$ to $\overline{\text{U-EDCP}}$ reduction

We give a $\overline{\text{G-EDCP}}$ to $\overline{\text{U-EDCP}}$ reduction using quantum rejection sampling. This proof is analogous to Lemma 7. Another related work [24, Theorem 32] provides a similar reduction, but focuses on a support of size two, with an inverse polynomial success probability. The following reduction converts $\overline{\text{G-EDCP}}$ samples to $\overline{\text{U-EDCP}}$ samples with constant success probability.

Lemma 16 ($\overline{\text{G-EDCP}}\le\overline{\text{U-EDCP}}$)**.** *Let κ be the security parameter. Let n,q,M and $l=\omega(\kappa)$ be positive integers, and σ be a positive real number such that $\sigma=\Omega(\sqrt{\kappa})$ and $q/\sigma=\Omega(\sqrt{\kappa})$. Then quantum rejection sampling converts $\overline{\text{G-EDCP}}^{l}_{n,q,\sigma}$ to $\overline{\text{U-EDCP}}^{\mathcal{O}(l)}_{n,q,M}$, where $M=c\cdot\sigma$ for some constant c, and succeeds with overwhelming probability.*

Proof. We use the quantum rejection sampling of Lemma 6 for $d=1$, e.g., when the $|\xi_k\rangle$'s degenerate to unknown complex phases. Concretely, the input $\overline{\text{G-EDCP}}_{n,q,\sigma}$ state (normalized) has the form:

$$\sum_{j\in\mathbb{Z}_q}\rho_\sigma(j)/(2^{1/4}\sqrt{\sigma})\,\omega_q^{j\cdot\langle \mathbf{a},\,\mathbf{s}\rangle}\,|j\rangle\,.$$

The output state, after rejection sampling, has the form:

$$(1/B)\cdot\sum_{j=0}^{M-1}\rho_\sigma(M)/(2^{1/4}\sqrt{\sigma})\,\omega_q^{j\cdot\langle \mathbf{a},\,\mathbf{s}\rangle}\,|j\rangle\,.$$

It is clear that the condition in Lemma 6 is satisfied. Thus the success probability is $B^2=M\cdot\rho_\sigma^2(M)/(2^{1/2}\sigma)$. For $M=c\cdot\sigma$ where c is a constant, the success probability is a constant. The lemma follows from a Chernoff bound argument. □

7 Impacts and discussions

In this section we assess the impacts of our algorithm and consider some further discussions and open questions.

7.1 LWE to $\overline{\text{U-EDCP}}$ and (non-)impacts of our algorithms

An algorithm for solving EDCP can be used to solve LWE via reductions from LWE to EDCP [55,19]. Concretely, Theorem 3 gives a reduction from LWE to G-EDCP where the amplitudes of the resulting states are Gaussian. We provide a direct reduction from LWE to U-EDCP, adapting the proof of Theorem 3. The same reduction has been given in [31, Section 5.2], but without details. For completeness, we present the reduction with concrete parameters.

Theorem 12 (LWE $\leq$ U-EDCP **reduction).** *Let κ be the security parameter. Given $m \geq n \log q$ where $m = o(2^\kappa)$ many $\text{LWE}_{n,q,\alpha}$ samples, there exists a probabilistic quantum reduction, with run-time polynomial in n, from $\text{LWE}^m_{n,q,\alpha}$ to $\text{U-EDCP}^l_{n,q,M}$ where $M \geq 1$ and $l \geq 1$ are integers, satisfying the condition:*

$$M \cdot l < 1/(64\sqrt{\kappa} m \alpha). \tag{10}$$

We refer to Appendix C for the proof of Theorem 12. Notice that, compared to the result of Theorem 3 for the case of G-EDCP, there is an asymptotic saving of $\sqrt{\kappa}$ in the denominator. Furthermore, the condition shows that it is better to keep the smallest value $m = n \log q$ in the reduction. It should also be possible to use the refined reduction in [19, Theorem 3] to further improve the above result, which we will leave as future work.

The above reduction can be specified to the constant error case, leading to a refined result that may be of independent interest. Concrete examples of LWE instances in this regime include LWE with binomial error distributions and practical homomorphic encryption schemes.

Lemma 17. *Let κ be the security parameter. Given $m \geq n \log q$ many $\text{LWE}_{n,q,B}$ samples, where $m = o(2^\kappa)$, with error bounded in ℓ_∞-norm by B, there exists a probabilistic quantum reduction, with run-time polynomial in n, from $\text{LWE}^m_{n,q,B}$ to $\text{U-EDCP}^l_{n,q,M}$ where $M \geq 1$ and $l \geq 1$ are integers, satisfying the condition:*

$$M \cdot l < q/(64Bm). \tag{11}$$

Proof. The proof is identical to the proof of Theorem 12 in Appendix C except we change the upper bound on $\|\boldsymbol{e}_0\|_\infty$. There is no need to use the tail bound in Lemma 4, as it is bounded by B. Thus we directly have:

$$M \cdot B < q/(8ck),$$

which, given that $c \geq 8$ and $k \geq ml$, leads to the statement. □

(Non-)impacts of our algorithms to standard LWE. Our algorithm of Section 4 requires $l = 2^{\Omega(\log n \log q)}$ samples to solve U-EDCP$_{n,q,M}$. Unfortunately (but *fortunately* for LWE), the two LWE to U-EDCP reductions outlined in this section put lower bounds on l for standard LWE parameters. Let $q = \mathsf{poly}(n)$. We first consider the case of fixed LWE error magnitude as in Ineq. (11) where:

$$2^{\Omega(\log n \log q)} < 1/(64MBm).$$

Re-arranging the terms of the above inequality and setting $M = 2$, we see that $B = o(1)$. This implies that the LWE instances we can tackle with our algorithm are actually $\mathsf{poly}(n)$-time solvable classically (e.g. via Gaussian elimination) since there is no error (with high probability) in them. Analogously, for LWE with Gaussian error with deviation αq, using Ineq. (10), we obtain a similar bound on αq, e.g., it must be inverse quasi-polynomially small. Note that one can also use Regev's reduction [58]. As discussed in Subsection 3.1, part of the reduction chain in [58] implies a reduction from LWE to U-EDCP$_{n,q,2}$, which can be solved by our algorithm. However this reduction, again, only produces perfect U-EDCP samples with an inverse polynomial probability (also, imperfect EDCP states are not efficiently detectable). Therefore, we conclude that our quasi-polynomial time algorithm does not affect the security of LWE with standard parameters.

7.2 Discussions

Directions for improvement. There are two potential avenues to make our algorithm applicable to more "standard" LWE instances. First is improved success probability of the g function from Lemma 18 in Appendix C, which restricts the number of output U-EDCP samples in the reduction. Concretely, the role of g is to map $\mathbf{A}\boldsymbol{x} + \boldsymbol{e}_1$ and $\mathbf{A}\boldsymbol{x} + \boldsymbol{e}_2$ to the same value. In Lemma 18 this function is efficient but it has a non-zero probability of failure. In contrast, one could consider a decoder for $\mathcal{L}_q(\mathbf{A})$ that can decode within the bound $\|\boldsymbol{e}_i\|_\infty$, which is smaller than $\lambda_1(\mathcal{L}_q(\mathbf{A}))$ for LWE instances. This would allow us to generate as many U-EDCP samples as we want, however, all known decoders for $\mathcal{L}_q(\mathbf{A})$ have at least exponential in $\dim(\mathcal{L}_q(\mathbf{A}))$ time, as this problem is nothing else but solving LWE classically.

The second approach is to improve the algorithm from Section 4 such that it requires less samples. We attempted to do so by considering a shorter chain of reductions (9) that directly operate on $\overline{\text{U-EDCP}}$ samples with $M > 2$, but, as explained in Subsection 4.3, this attempt failed (Lemma 12 becomes false).

Binary/small secret EDCP *and* LWE. Our algorithm from Section 4 performs slightly better when the secret is binary, namely, as soon as we know $\mathbf{s} \bmod 2$, we terminate the algorithm and save a $\mathsf{poly}(n)$ factor in the number of required samples and in runtime. On the LWE front, a binary secret can also speed up classical attacks [21], and moreover, the known LWE-to-EDCP reduction does not take advantage of the shape of the LWE secret. Hence, we do not know how to leverage small secrets in quantum attacks to achieve a super-polynomial improvement.

Extending the algorithm to general q. Since our algorithm is only applicable to a power-of-two modulus, it is natural to ask whether it can be generalized to an arbitrary modulus. For LWE, one can switch to a different modulus at the cost of a slight increase in error; see [20]. However, no such reduction tailored to EDCP is known. One possible approach is to leverage the chain of reductions EDCP $\to$ LWE $\xrightarrow{\text{mod switch}}$ LWE $\to$ EDCP. However, the first reduction produces LWE with $1/\alpha = \mathsf{poly}(n)$ when the M in EDCP is $\mathsf{poly}(n)$, and hence the final reduction [19] yields only a polynomial number of EDCP samples, rendering our algorithm inapplicable. Another possibility is to apply the modulus-switching techniques from [5] directly to EDCP samples. We leave this for future research.

Modulus vs. dimension trade-offs. Modulus-dimension switching for $\mathrm{LWE}_{n,q,\alpha}$ [20] allows reducing the LWE dimension to n/k by increasing the modulus to q^k, keeping the entropy of the LWE secret $2^{n \log q}$ fixed. If a converse reduction to LWE with dimension $n \log q$ and modulus 2 were possible, again with entropy $2^{n \log q}$, then we could potentially use the reductions EDCP $\to$ LWE $\xrightarrow{\text{mod switch}}$ LWE $\to$ EDCP to convert $\mathrm{EDCP}_{n,q}$ to $\mathrm{EDCP}_{n \log q,2}$. Our algorithm can solve this in polynomial time, so it would be interesting to exhibit such a reduction.

Application to $\mathrm{C}\,|\mathrm{LWE}\rangle$. As discussed earlier, it has been observed in [25] that, if there is a quantum algorithm that solves the $\mathrm{S}\,|\mathrm{LWE}\rangle$ problem without collapsing the input quantum states, then there is a quantum algorithm that solves the $\mathrm{C}\,|\mathrm{LWE}\rangle$ problem. Notice that our quasi-polynomial algorithm of Section 6 solves $\mathrm{S}\,|\mathrm{LWE}\rangle$, however, it does destroy the input $\mathrm{S}\,|\mathrm{LWE}\rangle$ states due to the merging steps. If this algorithm could be modified to preserve the input states then it could be relevant to the task of oblivious LWE sampling [23,30] as well as potentially impact various quantum cryptographic constructions for certified deletion and key revocation [53,8]. However, at present, the potential impacts of our algorithm on applications based on $\mathrm{C}\,|\mathrm{LWE}\rangle$ remain unclear. We leave this for future research.

Acknowledgements. We thank the reviewers for their valuable comments and suggestions, we also thank Yilei Chen, Damien Stehlé and Weiqiang Wen for insightful discussions.

References

1. Ajtai, M.: Generating hard instances of lattice problems (extended abstract). In: 28th ACM STOC. pp. 99–108. ACM Press (May 1996). https://doi.org/10.1145/237814.237838
2. Ajtai, M., Kumar, R., Sivakumar, D.: A sieve algorithm for the shortest lattice vector problem. In: 33rd ACM STOC. pp. 601–610. ACM Press (Jul 2001). https://doi.org/10.1145/380752.380857
3. Albrecht, M.R., Cid, C., Faugère, J.C., Fitzpatrick, R., Perret, L.: Algebraic algorithms for lwe problems. ACM Commun. Comput. Algebra **49**(2), 62 (Aug

2015). https://doi.org/10.1145/2815111.2815158, https://doi.org/10.1145/2815111.2815158
4. Albrecht, M.R., Ducas, L., Herold, G., Kirshanova, E., Postlethwaite, E.W., Stevens, M.: The general sieve kernel and new records in lattice reduction. In: Ishai, Y., Rijmen, V. (eds.) EUROCRYPT 2019, Part II. LNCS, vol. 11477, pp. 717–746. Springer, Cham (May 2019). https://doi.org/10.1007/978-3-030-17656-3_25
5. Albrecht, M.R., Faugère, J.C., Fitzpatrick, R., Perret, L.: Lazy modulus switching for the BKW algorithm on LWE. In: Krawczyk, H. (ed.) PKC 2014. LNCS, vol. 8383, pp. 429–445. Springer, Berlin, Heidelberg (Mar 2014). https://doi.org/10.1007/978-3-642-54631-0_25
6. Albrecht, M.R., Gheorghiu, V., Postlethwaite, E.W., Schanck, J.M.: Estimating quantum speedups for lattice sieves. In: Moriai, S., Wang, H. (eds.) ASIACRYPT 2020, Part II. LNCS, vol. 12492, pp. 583–613. Springer, Cham (Dec 2020). https://doi.org/10.1007/978-3-030-64834-3_20
7. Albrecht, M.R., Player, R., Scott, S.: On the concrete hardness of Learning with Errors. Journal of Mathematical Cryptology **9**(3), 169–203 (2015)
8. Ananth, P., Poremba, A., Vaikuntanathan, V.: Revocable cryptography from learning with errors. In: Rothblum, G.N., Wee, H. (eds.) TCC 2023, Part IV. LNCS, vol. 14372, pp. 93–122. Springer, Cham (Nov / Dec 2023). https://doi.org/10.1007/978-3-031-48624-1_4
9. Aono, Y., Nguyen, P.Q., Shen, Y.: Quantum lattice enumeration and tweaking discrete pruning. In: Peyrin and Galbraith [52], pp. 405–434. https://doi.org/10.1007/978-3-030-03326-2_14
10. Arora, S., Ge, R.: New algorithms for learning in presence of errors. In: Aceto, L., Henzinger, M., Sgall, J. (eds.) ICALP 2011, Part I. LNCS, vol. 6755, pp. 403–415. Springer, Berlin, Heidelberg (Jul 2011). https://doi.org/10.1007/978-3-642-22006-7_34
11. Bai, S., Miller, S., Wen, W.: A refined analysis of the cost for solving LWE via uSVP. In: Buchmann, J., Nitaj, A., eddine Rachidi, T. (eds.) AFRICACRYPT 19. LNCS, vol. 11627, pp. 181–205. Springer, Cham (Jul 2019). https://doi.org/10.1007/978-3-030-23696-0_10
12. Bai, S., Van Hoof, I., Johnson, F., Lange, T., Ngo, T.: Concrete analysis of quantum lattice enumeration. In: Guo, J., Steinfeld, R. (eds.) ASIACRYPT 2023, Part III. LNCS, vol. 14440, pp. 131–166. Springer, Singapore (Dec 2023). https://doi.org/10.1007/978-981-99-8727-6_5
13. Banaszczyk, W.: New bounds in some transference theorems in the geometry of numbers. Mathematische Annalen **296**, 625–635 (1993)
14. Becker, A., Ducas, L., Gama, N., Laarhoven, T.: New directions in nearest neighbor searching with applications to lattice sieving. In: Krauthgamer, R. (ed.) 27th SODA. pp. 10–24. ACM-SIAM (Jan 2016). https://doi.org/10.1137/1.9781611974331.ch2
15. Bindel, N., Bonnetain, X., Tiepelt, M., Virdia, F.: Quantum lattice enumeration in limited depth. In: Reyzin, L., Stebila, D. (eds.) CRYPTO 2024, Part VI. LNCS, vol. 14925, pp. 72–106. Springer, Cham (Aug 2024). https://doi.org/10.1007/978-3-031-68391-6_3
16. Bonnetain, X.: Hidden Structures and Quantum Cryptanalysis. (Structures cachées et cryptanalyse quantique). Ph.D. thesis, Sorbonne University, France (2019), https://tel.archives-ouvertes.fr/tel-02400328
17. Bonnetain, X., Chailloux, A., Schrottenloher, A., Shen, Y.: Finding many collisions via reusable quantum walks: Application to lattice sieving. In: Hazay, C., Stam,

M. (eds.) EUROCRYPT 2023, Part V. LNCS, vol. 14008, pp. 221–251. Springer, Cham (Apr 2023). https://doi.org/10.1007/978-3-031-30589-4_8
18. Bonnetain, X., Naya-Plasencia, M.: Hidden shift quantum cryptanalysis and implications. In: Peyrin and Galbraith [52], pp. 560–592. https://doi.org/10.1007/978-3-030-03326-2_19
19. Brakerski, Z., Kirshanova, E., Stehlé, D., Wen, W.: Learning with errors and extrapolated dihedral cosets. In: Abdalla, M., Dahab, R. (eds.) PKC 2018, Part II. LNCS, vol. 10770, pp. 702–727. Springer, Cham (Mar 2018). https://doi.org/10.1007/978-3-319-76581-5_24
20. Brakerski, Z., Langlois, A., Peikert, C., Regev, O., Stehlé, D.: Classical hardness of learning with errors. In: Boneh, D., Roughgarden, T., Feigenbaum, J. (eds.) 45th ACM STOC. pp. 575–584. ACM Press (Jun 2013). https://doi.org/10.1145/2488608.2488680
21. Buchmann, J.A., Göpfert, F., Player, R., Wunderer, T.: On the hardness of LWE with binary error: Revisiting the hybrid lattice-reduction and meet-in-the-middle attack. In: Pointcheval, D., Nitaj, A., Rachidi, T. (eds.) AFRICACRYPT 16. LNCS, vol. 9646, pp. 24–43. Springer, Cham (Apr 2016). https://doi.org/10.1007/978-3-319-31517-1_2
22. Chailloux, A., Loyer, J.: Lattice sieving via quantum random walks. In: Tibouchi, M., Wang, H. (eds.) ASIACRYPT 2021, Part IV. LNCS, vol. 13093, pp. 63–91. Springer, Cham (Dec 2021). https://doi.org/10.1007/978-3-030-92068-5_3
23. Chen, Y., Hu, Z., Liu, Q., Luo, H., Tu, Y.: LWE with quantum amplitudes: Algorithm, hardness, and oblivious sampling. Cryptology ePrint Archive, Paper 2023/1498 (2023), https://eprint.iacr.org/2023/1498
24. Chen, Y., Hu, Z., Liu, Q., Luo, H., Tu, Y.: On the hardness of $\mathsf{S}|\mathsf{LWE}\rangle$ with gaussian and other amplitudes. Cryptology ePrint Archive, Report 2023/1498 (2023), https://eprint.iacr.org/2023/1498
25. Chen, Y., Liu, Q., Zhandry, M.: Quantum algorithms for variants of average-case lattice problems via filtering. In: Dunkelman, O., Dziembowski, S. (eds.) EUROCRYPT 2022, Part III. LNCS, vol. 13277, pp. 372–401. Springer, Cham (May / Jun 2022). https://doi.org/10.1007/978-3-031-07082-2_14
26. Chen, Y., Nguyen, P.Q.: BKZ 2.0: Better lattice security estimates. In: Lee, D.H., Wang, X. (eds.) ASIACRYPT 2011. LNCS, vol. 7073, pp. 1–20. Springer, Berlin, Heidelberg (Dec 2011). https://doi.org/10.1007/978-3-642-25385-0_1
27. Childs, A.M., van Dam, W.: Quantum algorithm for a generalized hidden shift problem. In: Bansal, N., Pruhs, K., Stein, C. (eds.) 18th SODA. pp. 1225–1232. ACM-SIAM (Jan 2007)
28. Coppersmith, D.: An approximate fourier transform useful in quantum factoring (2002)
29. Dawson, C.M., Nielsen, M.A.: The solovay-kitaev algorithm (2005), https://arxiv.org/abs/quant-ph/0505030
30. Debris-Alazard, T., Fallahpour, P., Stehlé, D.: Quantum oblivious LWE sampling and insecurity of standard model lattice-based SNARKs. In: Mohar, B., Shinkar, I., O'Donnell, R. (eds.) 56th ACM STOC. pp. 423–434. ACM Press (Jun 2024). https://doi.org/10.1145/3618260.3649766
31. Doliskani, J.: Efficient quantum public-key encryption from learning with errors. Cryptology ePrint Archive, Report 2020/1557 (2020), https://eprint.iacr.org/2020/1557
32. Fincke, U., Pohst, M.: Improved methods for calculating vectors of short length in a lattice, including a complexity analysis. Mathematics of Computation **44**(170),

463–463 (May 1985). https://doi.org/10.1090/s0025-5718-1985-0777278-8, http://dx.doi.org/10.1090/S0025-5718-1985-0777278-8
33. Gama, N., Nguyen, P.Q., Regev, O.: Lattice enumeration using extreme pruning. In: Gilbert, H. (ed.) EUROCRYPT 2010. LNCS, vol. 6110, pp. 257–278. Springer, Berlin, Heidelberg (May / Jun 2010). https://doi.org/10.1007/978-3-642-13190-5_13
34. Gentry, C., Peikert, C., Vaikuntanathan, V.: Trapdoors for hard lattices and new cryptographic constructions. In: Ladner, R.E., Dwork, C. (eds.) 40th ACM STOC. pp. 197–206. ACM Press (May 2008). https://doi.org/10.1145/1374376.1374407
35. Hanrot, G., Stehlé, D.: Improved analysis of kannan's shortest lattice vector algorithm. In: Menezes, A. (ed.) CRYPTO 2007. LNCS, vol. 4622, pp. 170–186. Springer, Berlin, Heidelberg (Aug 2007). https://doi.org/10.1007/978-3-540-74143-5_10
36. Herold, G., Kirshanova, E., May, A.: On the asymptotic complexity of solving LWE. DCC **86**(1), 55–83 (2018). https://doi.org/10.1007/s10623-016-0326-0
37. Ivanyos, G., Prakash, A., Santha, M.: On learning linear functions from subset and its applications in quantum computing. In: Azar, Y., Bast, H., Herman, G. (eds.) 26th Annual European Symposium on Algorithms, ESA 2018, August 20-22, 2018, Helsinki, Finland. LIPIcs, vol. 112, pp. 66:1–66:14. Schloss Dagstuhl - Leibniz-Zentrum für Informatik (2018). https://doi.org/10.4230/LIPICS.ESA.2018.66, https://doi.org/10.4230/LIPIcs.ESA.2018.66
38. Kannan, R.: Improved algorithms for integer programming and related lattice problems. In: 15th ACM STOC. pp. 193–206. ACM Press (Apr 1983). https://doi.org/10.1145/800061.808749
39. Kirshanova, E.: A k-list algorithm for lwe (2020), https://crypto-kantiana.com/elena.kirshanova/talks/Simons_Hardness.pdf
40. Kitaev, A.Y.: Quantum computations: algorithms and error correction. Russ. Math. Surv. **52**(6), 1191–1249 (Dec 1997)
41. Kitaev, A.Y.: Quantum measurements and the abelian stabilizer problem. Electron. Colloquium Comput. Complex. **TR96** (1995)
42. Kuperberg, G.: A subexponential-time quantum algorithm for the dihedral hidden subgroup problem. SIAM Journal on Computing **35**(1), 170–188 (2005). https://doi.org/10.1137/S0097539703436345, https://doi.org/10.1137/S0097539703436345
43. Kuperberg, G.: Another subexponential-time quantum algorithm for the dihedral hidden subgroup problem (2011), https://arxiv.org/abs/1112.3333
44. Lyubashevsky, V.: Lattice signatures without trapdoors. In: Pointcheval, D., Johansson, T. (eds.) EUROCRYPT 2012. LNCS, vol. 7237, pp. 738–755. Springer, Berlin, Heidelberg (Apr 2012). https://doi.org/10.1007/978-3-642-29011-4_43
45. Lyubashevsky, V., Ducas, L., Kiltz, E., Lepoint, T., Schwabe, P., Seiler, G., Stehlé, D., Bai, S.: CRYSTALS-DILITHIUM. Tech. rep., National Institute of Standards and Technology (2022), available at https://csrc.nist.gov/Projects/post-quantum-cryptography/selected-algorithms-2022
46. Micciancio, D., Regev, O.: Worst-case to average-case reductions based on Gaussian measures. In: 45th FOCS. pp. 372–381. IEEE Computer Society Press (Oct 2004). https://doi.org/10.1109/FOCS.2004.72
47. Micciancio, D., Voulgaris, P.: Faster exponential time algorithms for the shortest vector problem. In: Charika, M. (ed.) 21st SODA. pp. 1468–1480. ACM-SIAM (Jan 2010). https://doi.org/10.1137/1.9781611973075.119

48. Micciancio, D., Walter, M.: Practical, predictable lattice basis reduction. In: Fischlin, M., Coron, J.S. (eds.) EUROCRYPT 2016, Part I. LNCS, vol. 9665, pp. 820–849. Springer, Berlin, Heidelberg (May 2016). https://doi.org/10.1007/978-3-662-49890-3_31
49. Nguyen, P.Q., Vidick, T.: Sieve algorithms for the shortest vector problem are practical. Journal of Mathematical Cryptology **2**(2) (2008)
50. Nielsen, M.A., Chuang, I.L.: Quantum Computation and Quantum Information: 10th Anniversary Edition. Cambridge University Press (2010)
51. Ozols, M., Roetteler, M., Roland, J.: Quantum rejection sampling. ACM Trans. Comput. Theory **5**(3) (Aug 2013)
52. Peyrin, T., Galbraith, S. (eds.): ASIACRYPT 2018, Part I, LNCS, vol. 11272. Springer, Cham (Dec 2018)
53. Poremba, A.: Quantum proofs of deletion for learning with errors. In: Kalai, Y.T. (ed.) ITCS 2023. vol. 251, pp. 90:1–90:14. LIPIcs (Jan 2023). https://doi.org/10.4230/LIPIcs.ITCS.2023.90
54. Prest, T., Fouque, P.A., Hoffstein, J., Kirchner, P., Lyubashevsky, V., Pornin, T., Ricosset, T., Seiler, G., Whyte, W., Zhang, Z.: FALCON. Tech. rep., National Institute of Standards and Technology (2022), available at https://csrc.nist.gov/Projects/post-quantum-cryptography/selected-algorithms-2022
55. Regev, O.: Quantum computation and lattice problems. In: 43rd FOCS. pp. 520–529. IEEE Computer Society Press (Nov 2002). https://doi.org/10.1109/SFCS.2002.1181976
56. Regev, O.: A subexponential time algorithm for the dihedral hidden subgroup problem with polynomial space (2004), https://arxiv.org/abs/quant-ph/0406151
57. Regev, O.: On lattices, learning with errors, random linear codes, and cryptography. In: Gabow, H.N., Fagin, R. (eds.) 37th ACM STOC. pp. 84–93. ACM Press (May 2005). https://doi.org/10.1145/1060590.1060603
58. Regev, O.: On the complexity of lattice problems with polynomial approximation factors. pp. 475–496. ISC, Springer (2010). https://doi.org/10.1007/978-3-642-02295-1
59. Regev, O., Stephens-Davidowitz, N.: A reverse Minkowski theorem. Annals of Mathematics **199**(1), 1 – 49 (2024). https://doi.org/10.4007/annals.2024.199.1.1, https://doi.org/10.4007/annals.2024.199.1.1
60. Schnorr, C., Euchner, M.: Lattice basis reduction: improved practical algorithms and solving subset sum problems. Mathematics of Programming **66**, 181–199 (1994)
61. Schwabe, P., Avanzi, R., Bos, J., Ducas, L., Kiltz, E., Lepoint, T., Lyubashevsky, V., Schanck, J.M., Seiler, G., Stehlé, D., Ding, J.: CRYSTALS-KYBER. Tech. rep., National Institute of Standards and Technology (2022), available at https://csrc.nist.gov/Projects/post-quantum-cryptography/selected-algorithms-2022
62. Stehlé, D., Steinfeld, R., Tanaka, K., Xagawa, K.: Efficient public key encryption based on ideal lattices. In: Matsui, M. (ed.) ASIACRYPT 2009. LNCS, vol. 5912, pp. 617–635. Springer, Berlin, Heidelberg (Dec 2009). https://doi.org/10.1007/978-3-642-10366-7_36
63. Wagner, D.: A generalized birthday problem. In: Yung, M. (ed.) CRYPTO 2002. LNCS, vol. 2442, pp. 288–303. Springer, Berlin, Heidelberg (Aug 2002). https://doi.org/10.1007/3-540-45708-9_19
64. Watrous, J.: The Theory of Quantum Information. Cambridge University Press (2018)

A Algorithms for LWE and SIS

We provide a brief overview of algorithms for the LWE problem relevant to the scope of this paper.

The LWE problem can be solved using algorithms that solve the approximate Shortest Vector Problem (SVP), such as lattice reduction techniques. Notable algorithms in this category include the Block-Korkine-Zolotarev (BKZ) algorithm [60] and its modern variants [35,33,26,48]. Various practical strategies have been proposed for implementing the core operations required in these algorithms, such as enumeration [38,32,60] and sieving [2,49,47,14]. Quantum variants of these core procedures have been investigated in the literature [6,22,12]. In this work, we focus primarily on the asymptotic behavior and do not delve into details. For discussions on practical aspects, we refer the readers to [7,4].

Asymptotically, given $\mathrm{LWE}^{m}_{n,q,\alpha}$ instances where the number of samples m is sufficiently large, lattice reduction algorithms solve LWE [36,11] with a running-time of $\exp(\mathcal{O}(n \log q \log n / \log^2 \alpha))$. We consider a few relevant parameter settings for the LWE problem. Let n denote the security parameter: when $q = \mathrm{poly}(n)$ and $1/\alpha = \mathrm{poly}(n)$, lattice reduction requires exponential time; when q and $1/\alpha$ are subexponential, e.g., $q = \mathcal{O}(2^{n^\epsilon})$ for $\epsilon \in (0,1)$ and $1/\alpha = \Theta(2^{n^\delta})$ for $\delta > \epsilon/2$, lattice reduction runs in subexponential time – such parameters satisfy the reduction described in [57, Theorem 1.1], which reduces the LWE problem to a worst-case lattice problem with a subexponential approximation factor $\tilde{\mathcal{O}}(n/\alpha)$.

A similar approach applies to the Short Integer Solution (SIS) problem. Given $\mathrm{SIS}^{\infty}_{m,n,q,\beta}$ instances where the number of columns m is sufficiently large, the lattice reduction algorithms solve SIS with an asymptotic running time of $\exp(\mathcal{O}(n \log q \log n / \log^2(\beta/q)))$.

Lattice reduction is one of the strategies for solving LWE; however, it is not always the most efficient approach for certain parameter settings. For example, when $q = \mathrm{poly}(n)$ and $\alpha q = \mathcal{O}(1)$, algebraic algorithms such as the Arora-Ge algorithm [10] run in polynomial time, provided a polynomial number of samples is available. In contrast, lattice reduction requires full exponential time under the same conditions. More generally, the algebraic-type algorithms have a runtime of $2^{\tilde{\mathcal{O}}((\alpha q)^2)}$ [3], thus becomes subexponential when $\alpha q = o(\sqrt{n})$.

B Kuperberg's sieve

We give some more details on Kuperberg's sieve.

The input consists of DCP samples of the form $|0\rangle |x\rangle + |1\rangle |x + s \pmod N\rangle$, where $x \leftarrow_{\$} \mathbb{Z}_N$. Implicitly, N should be at least exponential in the underlying security parameter. A quantum Fourier transform is applied to the second register over $\mathbb{Z}_N$, followed by measurement, resulting in a state:

$$|\psi_y\rangle \propto |0\rangle + \omega_N^{ys} |1\rangle , \tag{12}$$

where $y \leftarrow_{\$} \mathbb{Z}_N$ is classically known. We tensor two such states $|\psi_{y_1}\rangle, |\psi_{y_2}\rangle$ obtaining:

$$|0\rangle|0\rangle + \omega_N^{y_2 s}|0\rangle|1\rangle + \omega_N^{y_1 s}|1\rangle|0\rangle + \omega_N^{(y_1+y_2)s}|1\rangle|1\rangle .$$

Applying a CNOT gate that sends $|a, b\rangle \to |a, a \oplus b\rangle$, we get a state proportional to:

$$(|0\rangle + \omega_N^{(y_1+y_2)s}|1\rangle) \otimes |0\rangle + \omega_N^{y_2 s}(|0\rangle + \omega_N^{(y_1-y_2)s}|1\rangle) \otimes |1\rangle .$$

Measuring the second qubit leaves the first qubit in the state:

$$|\psi_{y_1 \pm y_2}\rangle \propto |0\rangle + \omega_N^{(y_1 \pm y_2)s}|1\rangle ,$$

up to a global phase, and each outcome occurs with probability 1/2.

Now, note that the classical values of y_1 and y_2 are known, allowing us to target merging those y_1 and y_2 that share certain least significant bits (LSBs). This results in the state $|\psi_{y_1 - y_2}\rangle$ with its LSBs zeroed.This combinatorial approach is central to all Kuperberg-like algorithms.

For simplicity, we assume that N is a power-of-two and denote $n = \log_2(N)$. To balance running time and the number of samples required, Kuperberg's sieving sets the number of least significant bits to be zeroed at each step to $m \approx \sqrt{n}$ and proceeds with approximately n/m steps. Once all the bits (except the first one) are zeroed, the algorithm produces states $|\psi_{2^{n-1}}\rangle$ or $|\psi_0\rangle$. Applying a Hadamard transform to the first type states reveals the least significant bit of the secret. The procedure is then repeated to recover additional bits.

We provide a rough complexity analysis. The algorithm proceeds through $\sqrt{n}$ stages, as each step zeros out $\sqrt{n}$ bits. Furthermore, each step requires at least $\mathcal{O}(2^{\sqrt{n}})$ samples to achieve a collision on the LSBs of $\sqrt{n}$ bits. At each iteration, one can start with a slightly larger subexponential number of samples, and after merging a smaller subexponential number, the result still contains a subexponential number of samples, which can then be used for the next iteration.

C LWE to $\overline{\text{U-EDCP}}$ reduction

To solve the LWE problem using our algorithm for $\overline{\text{U-EDCP}}$ described in Section 4, the following reduction chain can be used based on the previous work:

(1) reduce LWE to G-EDCP using [19, Theorems 2, 3];
(2) reduce G-EDCP to U-EDCP using [19, Lemma 8];
(3) reduce U-EDCP to $\overline{\text{U-EDCP}}$ and invoke our algorithm to solve $\overline{\text{U-EDCP}}$.

This reduction chain is depicted in dashed lines in Figure C.

In this section, we provide a direct reduction from LWE to U-EDCP, adapting the proof of Theorem 3. The same reduction has been given in [31, Section 5.2], but without details. This section gives more details on the LWE to $\overline{\text{U-EDCP}}$ reduction. The following lemma from [19] is of use.

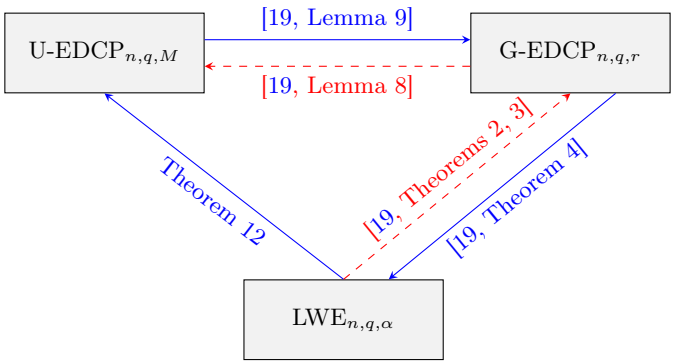

Fig. 2. Relations between EDCP variants and LWE.

Lemma 18 (Adapted from [19, Lemma 12]). *Let n, q and $m \geq n \log q$ be integers and $\mathbf{A} \in \mathbb{Z}_q^{m \times n}$. Let $c \geq 8$ be a constant. We are given m LWE samples $\mathbf{b} \equiv \mathbf{A}\mathbf{x} + \mathbf{e} \pmod{q}$. Take k to be the largest integer such that:*

$$k \leq \frac{\lambda_1^\infty(\Lambda_q(\mathbf{A}))}{2c \cdot \|\mathbf{e}_i\|_\infty}.$$

Set $z = q/c$ and $\bar{q} = q/z = c$. We define a function:

$$g : (b_1, \ldots, b_m) \mapsto (\lfloor b_1/z - w_1 \bmod \bar{q} \rceil, \ldots, \lfloor b_m/z - w_m \bmod \bar{q} \rceil),$$

where $w_1, \ldots, w_m$ are uniformly chosen from $[0, 1)$. Then for any $\mathbf{x} \in \mathbb{Z}_q^n$ the following two statements hold.

- *For any $\mathbf{u} = \mathbf{A}\mathbf{x} + \mathbf{e}_1, \mathbf{v} = \mathbf{A}\mathbf{x} + \mathbf{e}_2$ where $\|\mathbf{e}_i\|_\infty \leq \lambda_1^\infty(\Lambda_q(\mathbf{A}))/(2ck)$, with probability $(1 - 1/k)^m$, over the randomness of $w_1, \ldots, w_m$, we have $g(\mathbf{u}) = g(\mathbf{v})$.*
- *For any $\mathbf{u} = \mathbf{A}\mathbf{x} + \mathbf{e}_1, \mathbf{v} = \mathbf{A}\hat{\mathbf{x}} + \mathbf{e}_2$ where $\|\mathbf{e}_i\|_\infty \leq \lambda_1^\infty(\Lambda_q(\mathbf{A}))/(2ck)$ and $\mathbf{x} \neq \hat{\mathbf{x}}$, we have $g(\mathbf{u}) \neq g(\mathbf{v})$.*

Proof. This lemma is identical to [19, Lemma 12], except that it allows $m > n \log q$. □

Now we give the proof of Theorem 12.

Proof of Theorem 12. Input an LWE$_{n,q,\alpha}$ instance $(\mathbf{A}, \mathbf{b}_0)$ with $\mathbf{b}_0 \equiv \mathbf{A}\mathbf{s}_0 + \mathbf{e}_0 \pmod{q}$. We prepare the uniform superposition:

$$\sum_{\mathbf{s} \in \mathbb{Z}_q^n} \sum_{j=0}^{M-1} |j\rangle |\mathbf{s}\rangle .$$

Evaluate $f(j, \mathbf{s}) = \mathbf{A}\mathbf{s} - j \cdot \mathbf{b} \pmod q$ and store the result in the third register.

$$\sum_{\mathbf{s}\in\mathbb{Z}_q^n} \sum_{j=0}^{M-1} |j\rangle\,|\mathbf{s}\rangle\,|\mathbf{A}\mathbf{s} - j\cdot\mathbf{b} \pmod q\rangle\,.$$

By a change of variables on $\mathbf{s}$ we have

$$\sum_{\mathbf{s}\in\mathbb{Z}_q^n} \sum_{j=0}^{M-1} |j\rangle\,|\mathbf{s}\rangle\,|\mathbf{s} + j\cdot\mathbf{s}_0 \pmod q\rangle\,|\mathbf{A}\mathbf{s} - j\cdot\mathbf{e}_0 \pmod q\rangle\,.$$

Sample $w_1, \ldots, w_m$ uniformly from $[0, 1)$ and set $z = q/c$ for some constant $c \geq 8$. By Lemma 1, we have $z = q/c \leq q/8 \leq \lambda_1^\infty(\Lambda_q(\mathbf{A}))/2$ with probability $1 - 2^{-\Omega(m)}$ over the randomness of $\mathbf{A}$. Further, since $\lambda_1^\infty(\Lambda_q(\mathbf{A})) \leq q$ due to q-ary vectors, we have $z \in [1/c, 1/2] \cdot \lambda_1^\infty(\Lambda_q(\mathbf{A}))$ with probability $1 - 2^{-\Omega(m)}$. Set $\bar{q} := q/z = c$ and define

$$g : (x_1, \ldots, x_m) \mapsto (\lfloor x_1/z - w_1 \bmod \bar{q}\rfloor, \ldots, \lfloor x_m/z - w_m \bmod \bar{q}\rfloor).$$

Evaluate g on the third register, store the result in a new register and measure.

First, we upper bound $\|\mathbf{e}_0\|_\infty$. By the one-dimensional tail bound in Lemma 4, one coordinate of the vector $\mathbf{e}_0$ smaller than $\sqrt{\kappa}\alpha q$ with probability $\geq 1 - e^{-\Omega(\kappa)}$. Since $\mathbf{e}_0$ has m coordinates, $\|\mathbf{e}_0\|_\infty \leq \sqrt{\kappa}\alpha q$ with probability $\geq (1 - e^{-\Omega(\kappa)})^m$. Thus $\|M \cdot \mathbf{e}_0\|_\infty \leq M\sqrt{\kappa}\alpha q$ with overwhelming probability when $m = \mathrm{o}(2^\kappa)$. Second, in order to use Lemma 18, we need

$$M\|\mathbf{e}_0\|_\infty \leq \frac{\lambda_1^\infty(\Lambda_q(\mathbf{A}))}{2c \cdot k}.$$

for a success probability of $(1-1/k)^m$ for *one* sample. Aiming at l-many U-EDCP samples with constant success probability, we require $k \geq ml$.

Finally, we combine the two upper bounds on $\|M\mathbf{e}_0\|_\infty$ to obtain an upper bound on M. Concretely, we want to achieve the "goal" inequality:

$$M\|\mathbf{e}_0\|_\infty \overset{\text{Lemma 4}}{\leq} M\sqrt{\kappa}\alpha q \overset{goal}{\leq} q/(8ck) \overset{\text{Lemma 1}}{\leq} \lambda_1^\infty(\Lambda_q(\mathbf{A}))/(2ck). \tag{13}$$

Therefore, recalling that $c \geq 8$ due to Lemma 12 it is sufficient to set the parameters:

$$M < \frac{1}{(64\sqrt{\kappa}k\alpha)} \leq \frac{1}{(64\sqrt{\kappa}lm\alpha)}.$$

When the above inequality is satisfied, we have the following state after the measurement:

$$\sum_{j=0}^{M-1} |j\rangle\,|\mathbf{s} + j\cdot\mathbf{s}_0\rangle\,|\mathbf{A}\mathbf{s} - j\cdot\mathbf{e}_0\rangle$$

for some known $\mathbf{s} \in \mathbb{Z}_q^n$, with constant probability $(1 - 1/k)^m$. Uncompute and discard the third register using the function $(j, \mathbf{s}, \mathbf{b}) \mapsto \mathbf{b} - \mathbf{A}\mathbf{s} + j \cdot \mathbf{b}_0 \pmod q$

to obtain the desired states of the form

$$\sum_{j=0}^{M-1} |j\rangle \, |\mathbf{s} + j \cdot \mathbf{s}_0 \pmod{q}\rangle \, .$$

The success probability over l many states is a constant when $k \gtrapprox ml$. $\square$