

A Locality-Sensitive Hashing Framework for Reducing Bounded Distance Decoding to EDCP

Ryann Cartor^[0000–0003–0708–2363], Felice Manganiello^[0000–0002–5764–569X],
and William Youmans^{*[0000–0002–4173–4030]}

Clemson University, United States
rcartor@clemson.edu, manganm@clemson.edu, youmans.wj@gmail.com
*Corresponding author.

Abstract. Bounded Distance Decoding (BDD) is a fundamental primitive in both code- and lattice-based post-quantum cryptography. Prior work of Regev (FOCS 2002) and Brakerski, Kirshanova, Stehlé and Wen (PKC 2018) connected lattice BDD and Learning With Errors (LWE) to the Extrapolated Dihedral Coset Problem (EDCP), but these reductions rely heavily on geometric structure and do not naturally extend to coding-theoretic metrics. We present a general quantum reduction from BDD over finite Abelian groups equipped with translation-invariant metrics to EDCP, requiring only the existence of an appropriate locality-sensitive hash family. Instantiating this framework yields new reductions for decoding in the Hamming metric, rank metric, and the p -Lee metric (which coincides with the standard Lee metric when $p = 1$). For the Hamming and rank metrics, our reductions only yield a non-negligible number of EDCP states in parameter regimes where known decoding algorithms are already efficient. However, we recover the LWE-to-EDCP reduction within our framework by combining known equivalences between lattice CVP/BDD in the ℓ_p norm and p -Lee metric decoding.

Keywords: Bounded Distance Decoding · Learning With Errors · Dihedral Coset Problem · Quantum Algorithms

1 Introduction

The ongoing transition to Post-Quantum Cryptography (PQC) relies heavily on the presumed hardness of two classes of problems: decoding linear codes and decoding lattices. While recent standardization efforts have focused on lattice-based primitives (e.g., ML-KEM [31], ML-DSA [40]), code-based cryptography remains one of the most mature and trusted alternatives, dating back to the McEliece cryptosystem [30]. However, the hardness of both classes of PQC schemes is rooted in the same problem: Bounded Distance Decoding (BDD). In the context of codes, BDD asks for the closest codeword to a noisy codeword, typically in the Hamming metric. In the context of lattices, it consists of finding the closest lattice vector to a target point known to be close to the lattice, typically with respect to the Euclidean norm.

Understanding the quantum complexity of BDD is essential for determining the long-term security of these schemes. In the context of lattices, the seminal work of Regev [35] established a connection between BDD and the quantum Dihedral Coset Problem (DCP). This relationship was later strengthened by Brakerski et al. [4], who demonstrated that Learning With Errors (LWE) – an average-case variant of lattice BDD – is computationally equivalent to a generalization of DCP known as the Extrapolated Dihedral Coset Problem (EDCP). This connection is particularly significant given recent algorithmic advances, such as the quasi-polynomial time algorithms for EDCP developed by Bai et al. [1], which suggest that EDCP may be less intractable than previously assumed for certain parameter regimes.

The known BDD-to-DCP/EDCP reductions, however, use geometric features of Euclidean lattices. Regev’s reduction and the EDCP reductions of Brakerski et al. rely on partitions, grids, balls, and Gaussian structure [35,4]. These are natural objects in Euclidean space, but they do not directly explain whether similar reductions should exist for decoding problems in the discrete metrics used in coding theory. This motivates the question studied in this work: *What structural property of a metric decoding problem enables a quantum reduction from BDD to EDCP?*

1.1 Contributions

As a first step, we define a notion of convexity for Abelian groups equipped with translation-invariant metrics that we call (r, M) -convexity, which is of particular importance to the p -Lee metric. We then give a general quantum reduction from BDD over a finite subgroup $\mathcal{D}$ to EDCP over $\mathcal{D}$ in the presence of a (r, M) -convex LSH family. The reduction is controlled by a parameter

$$p_{\text{iso}} = p_1 - M(|\mathcal{D}| - 1)p_2,$$

where p_1 and p_2 are the near- and far-collision probabilities of the LSH family. In other words, the success of this reduction critically depends on the behavior of the LSH family.

By instantiating our general reduction for particular metrics, we then obtain the following results:

- A reduction from p -Lee metric decoding over $\mathbb{Z}_q$ for $p \geq 1$ (coinciding with the standard Lee metric when $p = 1$) to EDCP in Theorem 4, using a LSH family inspired by the reductions of [35,4].
- A reduction from Hamming-metric decoding (within the unique decoding radius) over $\mathbb{F}_q$ to EDCP in Theorem 7, using the standard bit-sampling LSH family of Indyk-Motwani [22].
- A reduction from rank-metric decoding to EDCP in Theorem 9, via a natural hash family based on linear transformations.

We show that the LWE-to-EDCP reduction of Brakerski et al. [4] can be recovered by our framework using the computational equivalence of CVP/BDD

in the ℓ_p norm and decoding in the p -Lee metric due to Campello et al. [23] and Blomqvist [3].

Finally, we observe that, with the specific LSH families used in our Hamming and rank instantiations, the reduction does not yield enough EDCP samples to be useful in the usual cryptographic decoding regimes. They produce non-negligible EDCP inputs only in regimes where known generic decoding algorithms are already efficient. Therefore, both the Hamming and rank metric BDD-to-EDCP reductions are best interpreted as a conceptual instantiation of the framework.

1.2 Technical Overview

We now describe the reduction at a high level. Let A be an Abelian group with a translation-invariant metric induced by a weight function wt , and let $\mathcal{D} \leq A$ be a finite subgroup with minimum distance d . A BDD instance has the form $t = s + e$, where $s \in \mathcal{D}$ and $\text{wt}(e)$ is below the unique-decoding radius. The goal is to recover s .

The reduction begins by preparing a uniform superposition over $[M] \times \mathcal{D}$:

$$\frac{1}{\sqrt{M|\mathcal{D}|}} \sum_{j \in [M]} \sum_{y \in \mathcal{D}} |j\rangle |y\rangle.$$

It then computes $y - jt$ in an auxiliary register. Since $s \in \mathcal{D}$, the change of variables $y = x + js$ gives

$$\frac{1}{\sqrt{M|\mathcal{D}|}} \sum_{j \in [M]} \sum_{x \in \mathcal{D}} |j\rangle |x + js\rangle |x - je\rangle.$$

The next step is to hash and measure the third register. Namely, we sample $h \leftarrow \mathcal{F}$, compute $h(x - je)$, and measure the resulting hash value. This abstracts the partition-and-measure step in the reductions of Regev and Brakerski–Kirshanova–Stehlé–Wen: their Euclidean partitions are replaced by LSH buckets, whose near-collision, far-collision, and convexity properties isolate the desired arithmetic progression.

If, for some x , all points

$$x, x - e, \dots, x - (M - 1)e$$

fall in the same hash bucket, and no point from any other progression falls in that bucket, then (after uncomputing the third register) the state collapses to

$$\frac{1}{\sqrt{M}} \sum_{j \in [M]} |j\rangle |x + js\rangle,$$

which is an EDCP state with secret s .

The LSH assumptions guarantee this isolation event with probability at least

$$p_{\text{iso}} = p_1 - M(|\mathcal{D}| - 1)p_2.$$

The near-collision probability p_1 ensures that the endpoints of the intended progression collide. The (r, M) -convexity condition turns this endpoint collision into a collision of the whole progression. The far-collision probability p_2 , together with a union bound over the other progressions, ensures that false collisions are unlikely.

By instantiating this reduction with various metrics and LSH families, we obtain concrete reductions from BDD in the Hamming, rank, and p -Lee metrics to EDCP, see Table 1 and Figure 1.

Table 1. Summary of the reduction for $\text{BDD}_{\mathcal{C},r}$ for various $[n, k, d]$ codes $\mathcal{C}$.

Instantiation	output EDCP states	Useful regime
p -Lee	$\Theta\left(\frac{d \log n}{M r n^{1+1/p}}\right)$	recovers LWE-to-EDCP when $r \lesssim \frac{d}{M n^{1/p}}$
Hamming	$\ell_H = \Theta\left(\frac{(d-2r) \log n}{r \log(M q^k)}\right)$	for $k, d = \Theta(n)$: $\ell_H = \Theta(\log n/r)$
Rank	$\ell_R = \Theta\left(\frac{(d-2r) \log n}{r \log(M q^{mk})}\right)$	for $m, d = \Theta(n)$: $\ell_R = \Theta(\log n/(rk))$




Fig. 1. Relevant problems and their relationships.

2 Background

Let $\mathbb{Z}_q$ denote the ring of integers modulo q with representatives in $[q] := \{0, \dots, q-1\}$. Let $\mathbb{F}_q$ denote the finite field with q elements, where $q = p^N$ is a prime power, and let $\mathbb{F}_q^n$ be the set of length n vectors over $\mathbb{F}_q$.

Given a group G we write $\exp(G)$ for the *exponent* of G , which denotes the least common multiple of the orders of all $g \in G$. If there is no least common

multiple, we say $\exp(G) = 0$. Alternatively, $\exp(G)$ is the smallest positive integer n such that $ng = 0_G$ for all $g \in G$. The exponent of a finite field is its characteristic, and the exponent of $\mathbb{Z}_q$ (for a possibly composite q) is just q .

Lowercase bold letters such as $\mathbf{x}$ represent column vectors and uppercase bold letters such as $\mathbf{A}$ represent matrices. For a vector $\mathbf{x} \in \mathbb{R}^n$ we denote its ℓ_p norm $\|\mathbf{x}\|_p = (|x_1|^p + \dots + |x_n|^p)^{1/p}$ where $p \geq 1$ is a real number. The ℓ_∞ norm is the limit of the ℓ_p norms as $p \rightarrow \infty$ and is given by $\|\mathbf{x}\|_\infty = \max\{|x_1|, \dots, |x_n|\}$. When p is omitted, we will assume we are in the ℓ_2 norm.

We write $x \leftarrow \chi$ to denote sampling from a distribution χ . For a finite set S we write $x \leftarrow S$ to denote that x is uniformly sampled from S . For any $r \in \mathbb{R}_+$ we define the spherical Gaussian function $\rho_r(\mathbf{x}) = \exp(-\pi\|\mathbf{x}\|^2/r^2)$. We denote the *discrete* Gaussian distribution over $\mathbb{Z}$ by $D_{\mathbb{Z},r}$, which has density function proportional to $\rho_r(\cdot)$. We will use the following formulation of the tail bound for a discrete Gaussian distribution.

Lemma 1 ([28, Lemma 4.4]). *For any $\kappa > 0$ and $r > 0$, we have*

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

2.1 Metric spaces and Locality Sensitive Hash Functions

A *metric* on a set X is a function $d : X \times X \rightarrow \mathbb{R}_+$ which, for any $x, y, z \in X$ it holds that $d(x, y) = 0$ if and only if $x = y$; $d(x, y) = d(y, x)$; and the triangle inequality, $d(x, z) \leq d(x, y) + d(y, z)$. Additionally, we say a metric d on X is *translation invariant* if $d(x, y) = d(x + z, y + z)$ for all $x, y, z \in X$. Similarly, a *weight function* is a function $\text{wt} : X \rightarrow \mathbb{R}_+$ such that for all $x, y \in X$ it holds that $\text{wt}(x) = 0$ and only if $x = 0$; $\text{wt}(x) = \text{wt}(-x)$; and $\text{wt}(x + y) \leq \text{wt}(x) + \text{wt}(y)$.

There is a one-to-one correspondence between weight functions and translation invariant metrics. Namely, any weight function induces a translation invariant metric by defining $d(x, y) := \text{wt}(x - y)$. For any translation invariant metric d we define the induced weight function $\text{wt}(x) := d(x, 0)$ [3].

Definition 1 (Minimum Distance). *Let (X, d) be a metric space, and $\mathcal{D} \subseteq X$ a discrete subset. Define the minimum distance of $\mathcal{D}$ by*

$$d(\mathcal{D}) = \inf\{d(x, y) \mid x, y \in \mathcal{D}, x \neq y\}.$$

Locality sensitive hashing Locality Sensitive Hashing (LSH) is a technique originally developed for the (exact or approximate) Nearest Neighbor Search (NNS) problem [22], which is defined in Problem 1.

Problem 1 (Nearest Neighbor Search (NNS)). Given a set of points $P = \{p_1, \dots, p_n\}$ in some metric space (X, d) , preprocess P so that, when given a target point $x \notin P$, one can efficiently find a point $p \in P$ close to x .

Techniques for NNS were notoriously subject to the “curse of dimensionality,” where the data structures scaled poorly with the dimension of the data [19]. By

constructing a family of hash functions with the property that “near” points collide with significantly higher probability than “far” points, one can create a hash table such that querying the table returns a set of candidate points likely to be near neighbors to a target point, enabling a potentially faster search for nearest neighbors.

Definition 2 (Locality Sensitive Hash (LSH)). *Let (X, d) be a metric space, S a set, $0 \leq r_1 < r_2$. A family of hash functions $\mathcal{F} = \{h : X \rightarrow S\}$ is called (r_1, r_2, p_1, p_2) -sensitive if, for any $a, b \in X$*

- if $d(a, b) \leq r_1$ then $\Pr_{h \leftarrow U(\mathcal{F})}[h(a) = h(b)] \geq p_1$;
- if $d(a, b) \geq r_2$ then $\Pr_{h \leftarrow U(\mathcal{F})}[h(a) = h(b)] \leq p_2$.

When $p_1 \gg p_2$ there is strong separation between near and far points, significantly reducing the probability of false collisions and enabling time and memory trade-offs for NNS algorithms [26].

LSH efficiency. The efficiency of a (data-independent) (r_1, r_2, p_1, p_2) -sensitive LSH scheme is typically characterized by the exponent $\rho = \frac{\log(1/p_1)}{\log(1/p_2)}$, which determines the query time complexity $O(N^\rho)$ of LSH-based approximate NNS for a dataset of size N [22]. If an LSH family achieves the smallest possible ρ then it is called optimal.

Amplification. To improve separation between “near” and “far” pairs, LSH families are typically extended using two constructive operations: *AND amplification* and *OR amplification*. Given a (r_1, r_2, p_1, p_2) -sensitive LSH family $\mathcal{F}$, AND amplification constructs a new family by sampling t hashes from $\mathcal{F}$ independently and concatenating them, producing a (r_1, r_2, p_1^t, p_2^t) -sensitive LSH family. This can be helpful for decreasing the likelihood of false positives. OR amplification is similar, but declares two concatenated hashes to be equal if they agree on at least one entry (when viewing the concatenation as a vector). The result is a $(r_1, r_2, 1 - (1 - p_1)^t, 1 - (1 - p_2)^t)$ -sensitive LSH family.

2.2 Codes

Codes over $\mathbb{F}_q$. Let $(\mathbb{F}_q^n, d)$ be a metric space. A set $\mathcal{C} \subseteq \mathbb{F}_q^n$ is an $[n, k]$ -linear code if it is a k -dimensional linear subspace of $\mathbb{F}_q^n$ over $\mathbb{F}_q$. We call any element of $\mathcal{C}$ a *codeword*. For an $[n, k]$ code $\mathcal{C}$, a generator matrix $\mathbf{G} \in \mathbb{F}_q^{k \times n}$ is a full-rank matrix where the rows comprise a basis of $\mathcal{C}$ over $\mathbb{F}_q$. When the minimum distance d of a code is known, we write $[n, k, d]$ -linear code. The *rate* of a $[n, k, d]$ -linear code is $R = k/n$. In this work, we will primarily consider linear codes, which we will refer to simply as codes. We will always assume a code is given as a generator matrix.

Codes over $\mathbb{Z}_q$. Let $(\mathbb{Z}_q^n, d)$ be a metric space. A *ring-linear* code over $\mathbb{Z}_q$ of length n is a $\mathbb{Z}_q$ -submodule of $\mathbb{Z}_q^n$. We will call a ring-linear code over $\mathbb{Z}_q$ a $\mathbb{Z}_q$ -linear code. The *type* of a $\mathbb{Z}_q$ -linear code $\mathcal{C}$ is defined as $k = \log_q(|\mathcal{C}|)$ and the rate of $\mathcal{C}$ is defined as $R = k/n$. Like codes over $\mathbb{F}_q$, a $\mathbb{Z}_q$ -linear code $\mathcal{C}$ can be entirely represented by a generator matrix whose rows correspond to $\mathcal{C}$. Generator matrices for $\mathcal{C}$ are not necessarily unique, and if q is not prime, then even the number of rows of such matrices is not unique [42]. For background on ring-linear codes and codes over finite commutative rings, see [15].

2.3 Lattices

A lattice is an additive, finitely generated discrete subgroup of $\mathbb{Q}^m$. Let $\mathbf{B} = (\mathbf{b}_1, \dots, \mathbf{b}_n) \in \mathbb{Q}^{m \times n}$ be a full-rank matrix for $n \leq m$. The lattice $\mathcal{L}$ generated by $\mathbf{B}$ is defined as $\mathcal{L}(\mathbf{B}) = \{\mathbf{B}\mathbf{x} \mid \mathbf{x} \in \mathbb{Z}^n\}$. We say $\mathcal{L}$ has rank n , and that $\mathcal{L}$ has full rank if $n = m$. The matrix $\mathbf{B}$ is called the *basis* of $\mathcal{L}$. We define the minimum distance of a lattice by $\lambda_1^{(p)}(\mathcal{L})$, which is the length of the shortest non-zero vector of $\mathcal{L}$ in the ℓ_p norm.

A *q-ary lattice* is a lattice $\Lambda \subseteq \mathbb{Z}^n$ that satisfies $q\mathbb{Z}^n \subseteq \Lambda \subseteq \mathbb{Z}^n$. There is a one-to-one correspondence between q -ary lattices and codes $\mathcal{C} \subseteq \mathbb{Z}_q^n$ via *Construction A*, where the Construction A lattice of a code $\mathcal{C}$ is defined $\Lambda_q(\mathcal{C}) = \{x \in \mathbb{Z}^n \mid (x \bmod q) \in \mathcal{C}\}$. Note that $\Lambda_q(\mathcal{C})/q\mathbb{Z}^n \cong \mathcal{C}$.

2.4 Computational problems

We define the Bounded Distance Decoding problem for a generic metric space and discuss its translation to the metric spaces we mentioned in the previous section.

Problem 2 (Bounded Distance Decoding ($\text{BDD}_{\mathcal{D}, r}$)). Let (X, d) be a metric space, $\mathcal{D} \subseteq X$ a discrete subset. Given a target $t \in X$ and radius $r < \frac{1}{2}d(\mathcal{D})$ with the promise that $d(t, \mathcal{D}) \leq r$, $\text{BDD}_{\mathcal{D}, r}$ asks to output the unique $x^* \in \mathcal{D}$ such that $d(t, x^*) \leq r$.

This problem widely applies to multiple settings. We focus on the context of codes over finite fields, codes over rings, and lattices with various metrics.

BDD in Codes. For codes over finite fields and codes over $\mathbb{Z}_q$, the problem of decoding in the metric d within the unique decoding radius $\lfloor \frac{d-1}{2} \rfloor$ is exactly BDD.

Problem 3 (Decoding problem ($\text{DP}_{\mathcal{C}, r}$)). Given an $[n, k, d]$ -code $\mathcal{C}$, an integer $r < \frac{1}{2}d$, and $\hat{\mathbf{c}} = \mathbf{c} + \mathbf{e}$ such that $\mathbf{c} \in \mathcal{C}$ and $\text{wt}(\mathbf{e}) \leq r$, return $\mathbf{c}$.

Decoding was shown to be NP-complete in the Hamming metric by Berlekamp et al. in 1978 [2]. Consequently, for a random code, these problems are considered to be computationally hard to solve (even in other metrics). Generally, in

coding theory, the objective is to design structured codes that enable the correction of many errors while still allowing for efficient decoding. In contrast, code-based cryptography aims to disguise the underlying structure of a code so that decoding remains infeasible for unauthorized parties.

BDD in Lattices. For lattices, BDD is typically stated as the problem of finding $\mathbf{x}^* \in \mathcal{L}$ (if it exists) such that $\|\mathbf{x}^* - \mathbf{t}\|_p \leq \alpha \lambda_1^{(p)}(\mathcal{L})$ for a given target $\mathbf{t} \in \mathbb{R}^n$, where $\alpha \in (0, 1/2)$. If we add the promise that a solution exists, then this is also just a special case of Problem 2.

BDD is a central primitive in lattice-based cryptography and is tightly connected to other standard lattice problems. In particular, BDD (at appropriate approximation factors) is equivalent up to polynomial factors to the unique Shortest Vector Problem (uSVP) and the decision problem GapSVP [29].

We now introduce the Learning With Errors (LWE) problem, an average-case variant of BDD, which serves as a central computational hardness assumption in lattice-based cryptography.

Problem 4 (Search Learning With Errors [37]). With input parameters $m, n \geq 1$, a modulus $q \geq 2$ and a distribution χ , the search version of $\text{LWE}_{n,q,\chi}^m$ consists of m samples of the form $(\mathbf{a}, b) \in \mathbb{Z}_q^n \times \mathbb{Z}_q$, with $\mathbf{a}$ uniform in $\mathbb{Z}_q^n$, $b = \langle \mathbf{a}, \mathbf{s} \rangle + e \bmod q$ and $e \leftarrow \chi$ for a fixed $\mathbf{s} \in \mathbb{Z}_q^n$, and asks to recover the secret $\mathbf{s}$.

A search LWE instance is typically represented in matrix form $(\mathbf{A}, \mathbf{b}) \in \mathbb{Z}_q^{m \times n} \times \mathbb{Z}_q^m$ where $\mathbf{b} = \mathbf{A}\mathbf{s} + \mathbf{e} \bmod q$. The most common choice of distribution χ is the discrete Gaussian distribution $D_{\mathbb{Z}, \alpha q}$, which we denote $\text{LWE}_{n,q,\alpha}^m$.

In the context of LWE, we can view the image of a matrix $\mathbf{A} \in \mathbb{Z}_q^{m \times n}$ as a code $\mathcal{C}(\mathbf{A}) := \{\mathbf{A}\mathbf{x} \mid \mathbf{x} \in \mathbb{Z}_q^n\}$. The Construction A lattice $\Lambda_q(\mathbf{A}) := \Lambda_q(\mathcal{C}(\mathbf{A})) = \{\mathbf{y} \in \mathbb{Z}^m \mid \mathbf{y} = \mathbf{A}\mathbf{s} \bmod q, \mathbf{s} \in \mathbb{Z}^n\}$ is referred to as the *primal lattice* defined by $\mathbf{A}$. We have the following lower bound on the minimum distance of $\Lambda_q(\mathbf{A})$ in the ℓ_∞ norm.

Lemma 2 ([4, Lemma 5]). *Given a uniformly chosen matrix $\mathbf{A} \in \mathbb{Z}_q^{m \times n}$ for some positive integers q, m , and n such that $m \geq n$, we have $\lambda_1^\infty(\Lambda_q(\mathbf{A})) \geq q^{(m-n)/m}/2$ with probability $1 - 2^{-m}$.*

Search LWE is often referred to as an average-case variant of BDD. Concretely, let $\mathbf{A} \in \mathbb{Z}_q^{m \times n}$ be an instance of (search) $\text{LWE}_{n,q,\alpha}^m$, with $\mathbf{b} = \mathbf{A}\mathbf{s} + \mathbf{e} \bmod q$ for some secret $\mathbf{s} \in \mathbb{Z}_q^n$ and error $\mathbf{e} \leftarrow D_{\mathbb{Z}^m, \alpha q}$. Let $\kappa = \Omega(m)$ be a security parameter and $q = \text{poly}(n)$. By Lemma 1, we have $\|\mathbf{e}\|_\infty \leq \sqrt{\kappa} \alpha q$ with overwhelming probability, so we can interpret $\text{LWE}_{n,q,\alpha}$ as $\text{BDD}_{\Lambda_q(\mathbf{A}), r}$ in the ℓ_∞ norm, with $r = \sqrt{\kappa} \alpha q$.

A related problem in lattices is that of the Closest Vector Problem (CVP), which, given a lattice and a target point, asks to find the lattice point that is closest to that target point under a chosen distance metric.

Campello et al. [23] showed that CVP on the q -ary lattice $\Lambda_q(\mathcal{C})$ in the ℓ_p norm is computationally equivalent to the decoding within the unique decoding radius on $\mathcal{C}$ in the p -Lee metric. This was refined by Blomqvist using the notion

of a quotient metric [3]. We reformulate this equivalence in the language of BDD here.

Theorem 1 (Adapted from [3, Theorems 4 and 6]). *Let n, q be integers and $p \geq 1$. Let $\mathcal{C} \subseteq \mathbb{Z}_q^n$ be a code and $\Lambda_q(\mathcal{C})$ be the corresponding q -ary Construction A lattice. Let $r \in \mathbb{R}_+$ be a decoding radius. Then $\text{BDD}_{\Lambda_q(\mathcal{C}), r}$ in the ℓ_p norm is computationally equivalent to $\text{BDD}_{\mathcal{C}, r}$ in the p -Lee metric.*

3 Quantum Background

We use standard notation from quantum computation; see [32] for background. A pure state is a unit vector $|\psi\rangle$ in a finite-dimensional Hilbert space. We write states in the computational basis, so that $|\psi\rangle = \sum_x a_x |x\rangle$ where $\sum_x |a_x|^2 = 1$. Measuring in the computational basis gives outcome x with probability $|a_x|^2$. We sometimes omit normalization factors when they are clear from context.

We will also use the density-operator formalism. A mixed state is represented by a positive semidefinite operator ρ of trace 1. A pure state $|\psi\rangle$ corresponds to the rank-one density operator $|\psi\rangle\langle\psi|$. More generally, an ensemble $\{(p_i, |\psi_i\rangle)\}_i$ is represented by

$$\rho = \sum_i p_i |\psi_i\rangle\langle\psi_i|.$$

We write $X \succeq Y$ to mean that $X - Y$ is positive semidefinite. Finally, any quantum algorithm with a classical output can be represented, as far as its input state is concerned, by a POVM $\{\Pi_z\}_z$. On input density operator ρ , the probability that the algorithm outputs z is $\text{Tr}(\Pi_z \rho)$.

3.1 DCP and EDCP

Regev [35] showed that lattice problems such as uSVP and BDD reduce quantumly to the Dihedral Coset Problem. In its standard cyclic form, DCP_N^ℓ is the following problem.

Problem 5 (Dihedral Coset Problem (DCP)). Let N, ℓ be positive integers, let $s \in \mathbb{Z}_N$ be fixed, and let $x_i \leftarrow \mathbb{Z}_N$ for $i \in [\ell]$. Given ℓ quantum states of the form

$$\frac{1}{\sqrt{2}}(|0\rangle |x_i\rangle + |1\rangle |x_i + s\rangle),$$

recover the hidden shift s .

No polynomial-time quantum algorithm for DCP is known. Kuperberg gave a subexponential-time quantum algorithm [24], later made space-efficient by Regev [36] and refined through Kuperberg's collimation sieve [25]. Subsequent work extended and analyzed these methods for more general cyclic and Abelian groups [34, 6].

Brakerski et al. defined a relaxed variant of DCP known as the Extrapolated Dihedral Coset Problem (EDCP) [4], which arises naturally as an intermediate step in Regev's reduction from uSVP to DCP.

Problem 6 (Extrapolated Dihedral Coset Problem (EDCP)). Let A be a finite additive Abelian group, and let q be a positive integer such that $\exp(A) \mid q$. Let $\chi : \mathbb{Z}_q \rightarrow \mathbb{C}$ satisfy $\sum_{j \in \mathbb{Z}_q} |\chi(j)|^2 = 1$. Let $s \in A$ be fixed, and let $x_i \leftarrow A$ for $i \in [\ell]$. Given ℓ quantum states of the form

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

recover s .

Most previous works consider $A = \mathbb{Z}_q^n$ for a possibly composite modulus q [4,10,9,1]. The more general definition given here allows us to consider EDCP states over a linear code $\mathcal{C} \subseteq \mathbb{F}_q^n$. Note that if $\mathcal{C} \neq \{0\}$ is an $\mathbb{F}_q$ -linear code with $q = p^N$, then $\exp(\mathcal{C}) = p$.

This work only uses the uniform-amplitude case. For $2 \leq M \leq q$, define

$$|\psi_{x,s}\rangle := \frac{1}{\sqrt{M}} \sum_{j \in [M]} |j\rangle |x + j \cdot s\rangle.$$

We denote the corresponding problem by $\text{EDCP}_{A,q,M}^\ell$. When $q = \exp(A)$, we omit q and write $\text{EDCP}_{A,M}^\ell$; when $A = \mathbb{Z}_q^n$, we also use the standard notation $\text{EDCP}_{n,q,M}^\ell$.

Although an EDCP sample is a pure state once the offset x is fixed, the offset is sampled randomly and is not given to the solver as classical side information. Thus, for a finite EDCP group G and secret $s \in G$, the one-sample uniform EDCP input distribution is represented by the mixed state

$$\sigma_s := \frac{1}{|G|} \sum_{x \in G} |\psi_{x,s}\rangle \langle \psi_{x,s}|.$$

Likewise, the usual ℓ -sample EDCP input distribution is represented by $\sigma_s^{\otimes \ell}$. Brakerski et al. showed that (under quantum polynomial-time reductions) LWE is computationally equivalent to EDCP [4]. Later work studied variants of EDCP and related quantum LWE -type problems, including $S|\text{LWE}\rangle$ and $C|\text{LWE}\rangle$ [10,9,1]. There is also a rich and closely related line of work on quantum decoding and quantum variants of LPN/LWE -style decoding problems [8,13,7]. Our focus is different: rather than designing algorithms for these quantum decoding problems, we ask when we can efficiently reduce a BDD instance to EDCP.

4 (r, M) -Convexity

In Euclidean space, a set is convex if it contains the continuous line segment connecting any two of its points. However, for discrete Abelian groups equipped with a metric, such as those found in coding theory, this definition of convexity does not apply. This motivates the following definition.

Definition 3 ((r, M)-Convex). Let G be an additive Abelian group equipped with a weight wt . Let $r \in \mathbb{R}_+$ and $M \geq 2$ be an integer. We say a subset $S \subseteq G$ is (r, M)-convex if for every $x, e \in G$ such that $\text{wt}(e) \leq r$, if $x \in S$ and $x + (M - 1)e \in S$, we have $x + je \in S$ for all $j \in [M]$.

This means if S is an (r, M)-convex set, then for any arithmetic progression $\{x, x + e, \dots, x + (M - 1)e\}$ where the endpoints x and $x + (M - 1)e$ are less than distance $(M - 1)r$ apart and are both elements of the set S , then all of the intermediate points are also in S . By abuse of notation, we will say that a function h is (r, M)-convex if the fiber (preimage set) $h^{-1}(b)$ is a (r, M)-convex set for every $b \in \text{im } h$.

The importance of the parameter M becomes apparent when working over a finite Abelian group. The following lemma gives a simple criterion for showing that a family of functions $\mathcal{F}$ is (∞, M) -convex when each function is also a group homomorphism.

Lemma 3. Let G and H be additive Abelian groups, $h : G \rightarrow H$ a group homomorphism. Suppose that for some integer $M \geq 2$ the group H has no elements of order k , for any $1 \leq k < M$. Then h is (∞, M) -convex.

We give the proof of Lemma 3 in Appendix A. Finally, we extend the notion of (r, M)-convexity to a LSH family by the following definition.

Definition 4 ((r, M)-Convex LSH). Let G be an additive Abelian group equipped with a metric d . Let $\mathcal{F}$ be a LSH family on (G, d) . We say $\mathcal{F}$ is a (r, M)-convex LSH family if every $h \in \mathcal{F}$ is (r, M)-convex.

5 From BDD to EDCP

We now present our main result, a general reduction from BDD in an arbitrary finite Abelian group equipped with a translation-invariant metric to EDCP. For convenience we define

$$\Delta_M(e) := \max_{j, j' \in [M]} \text{wt}((j - j')e)$$

for $e \in A$.

Theorem 2 (BDD $_{\mathcal{D}, r}$ to EDCP $_{\mathcal{D}, q, M}$). Let $(A, +)$ be a (non-trivial) efficiently represented Abelian group equipped with a weight $\text{wt} : A \rightarrow \mathbb{R}_+$. Let $\mathcal{D} \leq A$ be a finite subgroup with minimum distance d , and assume there is an efficient quantum procedure to prepare the uniform superposition

$$|\mathcal{D}\rangle = \frac{1}{\sqrt{|\mathcal{D}|}} \sum_{x \in \mathcal{D}} |x\rangle.$$

Let q and M be integers such that $\exp(\mathcal{D}) \mid q$ and $2 \leq M \leq q$. Let $\mathcal{F}$ be an efficiently sampleable and computable (r_1, r_2, p_1, p_2) -sensitive, (r, M)-convex LSH family for A with $p_1 > M(|\mathcal{D}| - 1)p_2$. Define

$$p_{\text{iso}} := p_1 - M(|\mathcal{D}| - 1)p_2 > 0.$$

Let $t = s + e \in A$ be a $\text{BDD}_{\mathcal{D},r}$ instance with $s \in \mathcal{D}$, $\text{wt}(e) \leq r < d/2$, and

$$\text{wt}((M-1)e) \leq r_1 < \frac{1}{2}d < r_2 \leq d - \Delta_M(e).$$

Suppose there is a quantum algorithm $\mathcal{A}_{\text{EDCP}}$ which, on ℓ independent ideal $\text{EDCP}_{\mathcal{D},q,M}$ samples $\sigma_s^{\otimes \ell}$, outputs s with probability at least ε . Then there is a quantum algorithm for $\text{BDD}_{\mathcal{D},r}$ which outputs s with probability at least $p_{\text{iso}}^\ell \varepsilon$.

We defer the proof of Theorem 2 to Appendix B.

Remark 1. If $p_{\text{iso}} = 1$, then the reduction can produce an arbitrary number of EDCP samples. Indeed, Lemma 10 gives $\rho_s \succeq \sigma_s$. Since both ρ_s and σ_s are density operators of trace 1, this implies $\rho_s = \sigma_s$. Hence, for every $\ell \geq 1$, ℓ independent executions of $\text{Gen}(t)$ produce the ideal input $\sigma_s^{\otimes \ell}$ with no probability loss. Therefore, the only interesting case for the sample count is $0 < p_{\text{iso}} < 1$, and in the remainder of this section we assume $p_{\text{iso}} < 1$.

5.1 Decoding in the p -Lee metric

The Lee metric was introduced by C. Lee in [27]. In contrast to the Hamming metric, the Lee metric accounts for the values of a codeword's nonzero entries. This metric gained attention in code-based cryptography largely due to the FuLeecca signature scheme [38], which was subsequently eliminated from NIST's additional call for digital signatures following the cryptanalysis in [21].

Definition 5 (Lee Weight, Lee Distance). Let $x \in \mathbb{Z}_q$ for a (possibly composite) integer q . The Lee weight of x over $\mathbb{Z}_q$ is $\text{wt}_{\text{Lee}}(x) = \min(|x|, q - |x|)$. For $\mathbf{x}, \mathbf{y} \in \mathbb{Z}_q^n$, the Lee distance between $\mathbf{x}$ and $\mathbf{y}$ is

$$d_{\text{Lee}}(\mathbf{x}, \mathbf{y}) = \sum_{i=1}^n \text{wt}_{\text{Lee}}(x_i - y_i).$$

The Lee distance can be seen as an analogue of the ℓ_1 metric for $\mathbb{Z}_q$. Just like how the ℓ_1 norm can be extended to the ℓ_p norm for $p \geq 1$, the Lee metric can be generalized to the p -Lee metric.

Definition 6. For $\mathbf{x}, \mathbf{y} \in \mathbb{Z}_q^n$, the p -Lee distance between $\mathbf{x}$ and $\mathbf{y}$ is

$$d_{p,\text{Lee}}(\mathbf{x}, \mathbf{y}) := \left(\sum_{i=1}^n \text{wt}_{\text{Lee}}(x_i - y_i)^p \right)^{1/p}$$

for $p \in [1, \infty)$ [23]. Likewise, for $p = \infty$ we define

$$d_{\infty,\text{Lee}}(\mathbf{x}, \mathbf{y}) = \max\{\text{wt}_{\text{Lee}}(x_i - y_i) \mid i \in [n]\}.$$

Alternatively, Blomqvist showed that the quotient metric on $\mathbb{R}^n/q\mathbb{Z}^n$ induced by the ℓ_p norm on $\mathbb{R}^n$ is precisely the p -Lee metric [3]. More precisely, define $\bar{\mathbf{x}} = \mathbf{x} \bmod q \in \mathbb{R}^n/q\mathbb{Z}^n$ for $\mathbf{x} \in \mathbb{R}^n$. Then we have

$$d_{\mathbb{R}^n/q\mathbb{Z}^n}(\bar{\mathbf{x}}, \bar{\mathbf{y}}) := \min_{\mathbf{u} \in \mathbf{x} + q\mathbb{Z}^n} \|\mathbf{u} - \mathbf{y}\|_p = d_{p, \text{Lee}}(\bar{\mathbf{x}}, \bar{\mathbf{y}}).$$

For more information on the p -Lee metric, see [11, 23], and on the theory of linear codes over $\mathbb{Z}_q$ endowed with the Lee metric, see [42].

In this section we will consider the p -Lee metric over $\mathbb{Z}_q^n$ for q a prime power. Let $c \geq 2$ be an integer and $w = q/c$. For any $\mathbf{x} \in \mathbb{Z}_q^n$, index $i \in [n]$, and an offset $\Delta \in [0, w)$, we define the single-coordinate hash

$$h_{i, \Delta}(\mathbf{x}) = \left\lfloor \frac{(x_i + \Delta) \bmod q}{w} \right\rfloor, \quad (1)$$

where $(x_i + \Delta) \bmod q$ is computed in the torus $\mathbb{R}/q\mathbb{Z}$. Then $h_{i, \Delta}$ partitions $\mathbb{Z}_q$ evenly into c buckets of width w . This is somewhat similar to the p -stable hash of [12] for the ℓ_p norm, and motivated by the reductions of Regev [35] and Brakerski et al. [4, Lemma 12].

Lemma 4 (Single-coordinate collision probability). *Fix $i \in [n]$. For any $\mathbf{x}, \mathbf{y} \in \mathbb{Z}_q^n$ define $t_i = \text{wt}_{\text{Lee}}(x_i - y_i)$. Over the randomness of Δ , we have*

$$\Pr_{\Delta}[h_{i, \Delta}(\mathbf{x}) = h_{i, \Delta}(\mathbf{y})] = \max\left(0, 1 - \frac{t_i}{w}\right).$$

Proof. Since $t_i \leq q/2$ and $\Delta < w \leq q$ we have a collision $h_{i, \Delta}(\mathbf{x}) = h_{i, \Delta}(\mathbf{y})$ if and only if $\Delta < w - t_i$, so the probability of a collision is $(w - t_i)/w = 1 - t_i/w$ over the randomness of Δ . If $t_i \geq w$ then $\Delta \geq w - t_i$ for all Δ , so no collision is possible. $\square$

By extending the single-coordinate hash of Lemma 4 to all coordinates we obtain the following LSH family for the p -Lee metric on $\mathbb{Z}_q^n$.

Theorem 3 (p -Lee metric LSH). *For any $0 \leq r_1 < r_2$ with $r_1 < w$ there exists a (r_1, r_2, p_1, p_2) -sensitive LSH family for $(\mathbb{Z}_q^n, d_{p, \text{Lee}})$ where*

$$p_1 = \left(1 - \frac{r_1}{w}\right)^n \quad \text{and} \quad p_2 = \max\left(0, 1 - \frac{r_2}{n^{1/p}w}\right).$$

In particular, when $r_2 \geq n^{1/p}w$, $p_2 = 0$.

We defer the proof of Theorem 3 to Appendix E.

Lemma 5. *Let $r \in \mathbb{R}_+$ and $2 \leq M < q$ be an integer such that $(M - 1)r < w$. Then the LSH family of Theorem 3 is (r, M) -convex.*

Proof. Let $h \in \mathcal{F}$ as defined in Equation (1) and $\mathbf{x}, \mathbf{e} \in \mathbb{Z}_q^n$ such that $\text{wt}(\mathbf{e}) \leq r$ and $h(\mathbf{x}) = h(\mathbf{x} + (M-1)\mathbf{e})$. We will show that $h(\mathbf{x}) = h(\mathbf{x} + j\mathbf{e})$ for all $j \in [M]$.

Fix a coordinate $i \in [n]$ and write $\Delta = \Delta_i$. Let $\epsilon \in \mathbb{Z}$ be the minimum integer with $\epsilon = e_i \bmod q$ and $|\epsilon| = \text{wt}_{p,\text{Lee}}(e_i)$, so that $|\epsilon| \leq \text{wt}_{p,\text{Lee}}(\mathbf{e}) \leq r$. Then

$$|(M-1)\epsilon| \leq (M-1)r < w \leq \frac{q}{2}.$$

Write $k = h_{i,\Delta}(\mathbf{x})$ and define $I = [kw, (k+1)w) \subset [0, q)$. Then there are integers $a, b \in \mathbb{Z}$ such that $x_i + \Delta + aq \in I$ and $x_i + (M-1)\epsilon + \Delta + bq \in I$, so that

$$|(M-1)\epsilon - (a-b)q| < w.$$

This implies that $|a-b|q < 2w$. Since $w = q/c$ with $c \geq 2$, we find $|a-b| < 1$, hence $a = b$.

For any $j \in [M]$ the real number $x_i + j\epsilon + \Delta$ lies between $x_i + \Delta$ and $x_i + (M-1)\epsilon + \Delta$, hence it is also contained in $I \subset [0, q)$. Then $h_{i,\Delta}(x_i + je_i) = h_{i,\Delta}(x_i)$. □

Instantiating the general reduction of Theorem 2 for a $\mathbb{Z}_q$ -linear code $\mathcal{C} \subseteq \mathbb{Z}_q^n$ requires efficiently preparing a uniform superposition over the codewords of $\mathcal{C}$:

$$\frac{1}{\sqrt{|\mathcal{C}|}} \sum_{\mathbf{c} \in \mathcal{C}} |\mathbf{c}\rangle.$$

We assume $\mathcal{C}$ is given as a generator matrix $\mathbf{G} \in \mathbb{Z}_q^{k \times n}$ whose rows generate $\mathcal{C}$ as a $\mathbb{Z}_q$ -submodule.

If $\mathbf{G}$ is left-invertible then we can efficiently prepare the uniform superposition by the same generator matrix construction used in the proof of Corollary 1. However, $\mathbf{G}$ is not necessarily left-invertible without additional assumptions. Instead, we rely on the Smith normal form (SNF).

By e.g. [16], there is an efficient (polynomial in $n, k, \log p$, and N) algorithm for computing the SNF of $\mathbf{G}$, producing an invariant-factor decomposition

$$\mathcal{C} \cong \bigoplus_{i=1}^r p^{a_i} \mathbb{Z}_q \cong \bigoplus_{i=1}^r \mathbb{Z}_{p^{N-a_i}}.$$

More importantly, the SNF provides an explicit isomorphism of $\mathcal{C}$ with its decomposition into cyclic factors, and we can efficiently prepare a uniform superposition over each cyclic factor.

Theorem 4 (Lee $\text{DP}_{\mathcal{C},r}$ to $\text{EDCP}_{\mathcal{C},M}$). *Let $\mathcal{C} \subseteq \mathbb{Z}_q^n$ be a $[n, k, d]$ $\mathbb{Z}_q$ -linear code in the p -Lee metric with q a prime power, $r \in \mathbb{R}_+$, and $2 \leq M < q$ be an integer such that*

$$r < \frac{d}{(2n^{1/p} + 1)(M-1)}.$$

Given a Lee $\text{DP}_{\mathcal{C},r}$ instance $\hat{\mathbf{c}} = \mathbf{c} + \mathbf{e}$ with $\mathbf{c} \in \mathcal{C}$ and $\text{wt}_{p,\text{Lee}}(\mathbf{e}) \leq r$, there is an efficient procedure that outputs ℓ states of $\text{EDCP}_{\mathcal{C},M}$ with secret $\mathbf{c}$, where

$$\ell = \Theta\left(\frac{d \log n}{Mrn^{1+1/p}}\right).$$

Proof. By the preceding discussion we can efficiently prepare a uniform superposition over $\mathcal{C}$ given a generator matrix. Let $\mathcal{F}$ be the hash family of Theorem 3 with $r_1 = (M-1)r$ and $r_2 = d - (M-1)r$. By the upper bound on r we have

$$\text{wt}_{p,\text{Lee}}((M-1)\mathbf{e}) \leq r_1 < \frac{d}{2} < r_2 \leq d - \text{wt}_{p,\text{Lee}}((M-1)\mathbf{e})$$

as required in Theorem 2.

Lemma 5 requires we choose $w = q/c$ such that $(M-1)r < w$. To force $p_2 = 0$ we also want $r_2 \geq n^{1/p}w$. The condition

$$r < \frac{d}{(2n^{1/p} + 1)(M-1)}$$

implies that $(r_1, r_2/n^{1/p}]$ is non-empty, so there exists a suitable w . More precisely, we can fix $c = \lceil qn^{1/p}/r_2 \rceil$. Then $c \geq 2$ and $w = q/c \in (r_1, r_2/n^{1/p}]$. Substituting w into $p_{\text{iso}} = p_1$ we find the stated expression for ℓ . $\square$

5.2 Recovering the LWE-to-EDCP reduction

We now explain how our framework can be used to recover the LWE to EDCP reduction of Brakerski et al. [4]. The primary tool is Theorem 1, which states that, for a code $\mathcal{C} \subseteq \mathbb{Z}_q^n$, BDD in $\Lambda_q(\mathcal{C})$ in the ℓ_p norm is computationally equivalent to BDD in $\mathcal{C}$ in the p -Lee metric.

As a prerequisite, for $\mathbf{A} \in \mathbb{Z}_q^{m \times n}$, we relate the minimum distances of $\Lambda_q(\mathbf{A})$ in the ℓ_∞ norm and $\mathcal{C}(\mathbf{A})$ in the ∞ -Lee metric in Lemma 6, which we prove in Appendix C.

Lemma 6. *Given a uniformly chosen matrix $\mathbf{A} \in \mathbb{Z}_q^{m \times n}$ for some positive integers q , m , and n such that $m \geq n$, we have $\lambda_1^\infty(\Lambda_q(\mathbf{A})) = d_{\infty,\text{Lee}}(\mathcal{C}(\mathbf{A}))$ with probability $1 - q^{-(m-n)}$.*

By viewing LWE as an average-case variant of BDD in the ℓ_∞ norm, and using the equivalence with ∞ -Lee decoding, we obtain the following reduction from LWE to EDCP.

Theorem 5 (LWE $_{n,q,\alpha}$ to EDCP $_{n,q,M}$). *Let n, q, m be positive integers and κ be a security parameter such that $m = n \log q = \Omega(\kappa)$. Let $\alpha \in (0, 1)$ such that $\alpha < 1/(8\sqrt{\kappa}(M-1))$. There is a probabilistic polynomial time reduction from LWE $_{n,q,\alpha}^m$ to EDCP $_{n,q,M}^\ell$ where $M \geq 2$ and ℓ are integers satisfying*

$$\ell = \Theta\left(\frac{\log m}{m\sqrt{\kappa}\alpha M}\right).$$

Proof. Let $(\mathbf{A}, \mathbf{b}) \in \mathbb{Z}_q^{m \times n} \times \mathbb{Z}_q^m$ be a $\text{LWE}_{n,q,\alpha}^m$ instance with $\mathbf{b} = \mathbf{A}\mathbf{s} + \mathbf{e}$, where $\mathbf{e} \leftarrow D_{\mathbb{Z}^m, \alpha q}$. By Lemma 1 we have $\|\mathbf{e}\|_\infty \leq \sqrt{\kappa} \alpha q$ with probability $1 - 2^{-m} = 1 - 2^{-\Omega(\kappa)}$. By Theorem 1 we can recover $\mathbf{A}\mathbf{s}$ (and hence $\mathbf{s}$) given an oracle for $\text{BDD}_{\mathcal{C},r}$ where $\mathcal{C} = \{\mathbf{A}\mathbf{x} \mid \mathbf{x} \in \mathbb{Z}_q^n\}$ and $r = \sqrt{\kappa} \alpha q$.

To invoke Theorem 4 we require $r < d/(2(M-1))$ where $d = d_{\infty, \text{Lee}}(\mathcal{C})$. By Lemma 2 and Lemma 6, with probability $(1 - q^{-(m-n)})(1 - 2^{-m}) = 1 - 2^{-\Omega(\kappa)}$, it suffices to assume that

$$r < \frac{q^{(m-n)/m}}{4(M-1)} = \frac{q}{8(M-1)},$$

i.e. $\alpha < 1/(8\sqrt{\kappa}(M-1))$.

Note that $d = \Theta(q)$ with overwhelming probability (again by Lemma 2 and Lemma 6) and substitute into Theorem 4 to obtain the stated number of EDCP states. $\square$

Remark 2. Our LWE-to-EDCP reduction above is stated for $q = p^N$, while the original reduction applies to arbitrary moduli [4]. This assumption is not inherent to the LSH-based reduction mechanism: the p -Lee metric and our LSH family make sense over $\mathbb{Z}_q$ for arbitrary $q \geq 2$. We impose the prime-power condition primarily to keep the uniform superposition preparation via SNF as simple as possible. We expect the same reduction to extend to arbitrary composite moduli after replacing the normal-form preparation step by the corresponding CRT/invariant-factor decomposition. We leave this bookkeeping to future work.

Brakerski et al. state their LWE-to-EDCP reduction for *Gaussian* EDCP [4, Theorem 2], and their reduction produces

$$\ell = \mathcal{O}\left(\frac{1}{m\kappa\alpha M}\right)$$

output states with constant probability. This matches our result up to a factor of $\sqrt{\kappa}$ (and logarithmic factors), which can be explained in two ways.

First, this discrepancy is identified by Bai et al. immediately following their simplified reduction from LWE directly to uniform EDCP [1, Theorem 12], which bypasses Gaussian EDCP and obtains an asymptotically equivalent reduction to ours.

Second, Brakerski et al. provide a second reduction that, instead of separating the ambient space $\mathbb{Z}^m$ by cubes, considers the intersection of balls around points and their shifts [4, Theorem 3] based on work by Regev [36]. This improves their stated ℓ by a factor of $\sqrt{\kappa}$, bringing it in line with ours.

6 Decoding in the Hamming and Rank Metrics

By instantiating Theorem 2 with a linear code $\mathcal{C} \subseteq \mathbb{F}_q^n$ we obtain the following useful corollary.

Corollary 1 ($\text{DP}_{\mathcal{C},r}$ to $\text{EDCP}_{\mathcal{C},p,M}$). *Let $\mathcal{C} \subseteq \mathbb{F}_q^n$ be a $[n, k, d]$ code with $q = p^N$ and $\text{wt} : \mathbb{F}_q^n \rightarrow \mathbb{R}_+$. Let $2 \leq M \leq p$ be an integer and $\hat{\mathbf{c}} = \mathbf{c} + \mathbf{e}$ be a $\text{DP}_{\mathcal{C},r}$ instance with $\mathbf{c} \in \mathcal{C}$ and $\text{wt}(\mathbf{e}) \leq r$.*

Given an efficiently sampleable and computable (r_1, r_2, p_1, p_2) -sensitive, (r, M) -convex LSH family $\mathcal{F}$ for $\mathbb{F}_q^n$ with $p_1 > M(q^k - 1)p_2$ and

$$\text{wt}((M-1)\mathbf{e}) \leq r_1 < \frac{1}{2}d < r_2 \leq d - \Delta_M(\mathbf{e})$$

there is an efficient procedure that outputs ℓ states of $\text{EDCP}_{\mathcal{C},p,M}$ with secret $\mathbf{c}$, where

$$\ell = \Theta\left(\frac{\log n}{\log(1/p_{\text{iso}})}\right).$$

We defer the proof of Corollary 1 to Appendix D.

6.1 Hamming metric

Originating with Hamming's foundational work [20], the Hamming metric is the most extensively studied metric in coding theory. It is also the metric underlying the well-studied McEliece cryptosystem [30] and HQC [18], the latter of which has been selected by NIST for standardization in its PQC process.

Definition 7 (Hamming distance). *Let $\mathbf{x}, \mathbf{y} \in \mathbb{F}_q^n$. The Hamming distance between $\mathbf{x}$ and $\mathbf{y}$ is*

$$d_H(\mathbf{x}, \mathbf{y}) = |\{i \in \{1, \dots, n\} \mid x_i \neq y_i\}|.$$

For the Hamming metric we consider the standard bit-sampling LSH family introduced by Indyk-Motwani [22]:

$$\mathcal{F} = \{h : \mathbf{x} \mapsto x_i \mid i \in [n]\}, \quad (2)$$

i.e. where hash functions are projections onto randomly selected coordinates. We extend this using AND amplification.

Remark 3. When $r_2 = cr_1$ for a constant $c > 1$ this LSH family yields an exponent $\rho \approx 1/c$. O'Donnell et al. later established a tight lower bound for ρ in the class of data-independent LSH families in the Hamming metric, proving that no scheme can achieve an exponent significantly smaller than $1/c$ [33]. In this sense, the bit-sampling LSH of Equation (2) for the Hamming metric is optimal.

Theorem 6 (Hamming metric LSH). *For any $0 \leq r_1 < r_2 \leq n$ and integer $t \geq 1$ there exists a $(r_1, r_2, (1 - r_1/n)^t, (1 - r_2/n)^t)$ -sensitive LSH family for $(\mathbb{F}_q^n, d_H)$.*

We defer the proof of Theorem 6 to Appendix F.

Lemma 7. *The LSH family of Theorem 6 is (∞, p) -convex.*

Proof. Since $\exp(\mathbb{F}_q) = p$ and each $h : \mathbb{F}_q^n \rightarrow \mathbb{F}_q^t$, $h \in \mathcal{F}$ is a homomorphism, $\mathcal{F}$ is (∞, p) -convex by Lemma 3. $\square$

Theorem 7 (Hamming $\text{DP}_{\mathcal{C},r}$ to $\text{EDCP}_{\mathcal{C},M}$). *Let $\mathcal{C} \subseteq \mathbb{F}_q^n$ be a $[n, k, d]$ code in the Hamming metric with $q = p^N$ and $2 \leq M \leq p$ be an integer. Given a Hamming $\text{DP}_{\mathcal{C},r}$ instance $\hat{\mathbf{c}} = \mathbf{c} + \mathbf{e}$ with $\mathbf{c} \in \mathcal{C}$ and $\text{wt}_H(\mathbf{e}) \leq r < d/2$, there is an efficient procedure that outputs ℓ states of $\text{EDCP}_{\mathcal{C},M}$ states with secret $\mathbf{c}$, where*

$$\ell = \Theta \left(\frac{(d - 2r) \log n}{r \log(Mq^k)} \right).$$

Proof. Let $\mathcal{F}$ be the LSH family from Theorem 6 with $r_1 = r$, $r_2 = d - r$, and $t \geq \log(Mq^k)/\alpha$ where

$$\alpha = \log \left(\frac{n - r}{n - d + r} \right) > 0.$$

This choice of t forces $p_1 > M(q^k - 1)p_2$, and since the Hamming metric satisfies $\text{wt}_H(j\mathbf{e}) \leq \text{wt}_H(\mathbf{e})$ with equality when $j \in \mathbb{F}_p^\times$ we have

$$\text{wt}_H(\mathbf{e}) \leq r_1 \leq \frac{d}{2} < r_2 \leq d - \text{wt}_H(\mathbf{e})$$

so we can apply Corollary 1.

We want to find a lower bound for $p_0 := p_1 - Mq^k p_2$. Factor out p_1 to rewrite p_0 as

$$p_0 = p_1 \left[1 - Mq^k \left(\frac{n - d + r}{n - r} \right)^t \right] = p_1 \left[1 - e^{\log(Mq^k) - \alpha t} \right].$$

When $t \geq \log(2Mq^k)/\alpha$ we have $p_0 \geq p_1/2$. Using the inequality $\log(1 + x) \geq x/(1 + x)$ for $x > -1$ we have $\alpha \geq (d - 2r)/(n - r)$, so it is sufficient to take

$$t = \frac{n - r}{d - 2r} \log(2Mq^k). \quad (3)$$

Since $p_{\text{iso}} \geq p_0 \geq p_1/2$ for this choice of t , we obtain ℓ states with probability $1/\text{poly}(n)$ when $(p_1/2)^\ell = 1/\text{poly}(n)$, which holds when $\ell = \Theta(\log n / \log(1/p_1))$. Using the inequality $-\log(1 - x) \geq x$ for $x \in (0, 1)$ we have $\log(1/p_1) \geq tr/n$, so we find

$$\ell = \Theta \left(\frac{n \log n}{rt} \right).$$

Substitute Equation (3) to complete the proof. $\square$

Random binary codes. To put our reduction from Theorem 7 into context, we consider a standard benchmark: random binary codes with rate $R = 1/2$. By the Gilbert-Varshamov bound, a random $[n, k, d]$ code with rate $R = k/n$ has relative distance d/n concentrating near $\delta_q = h_q^{-1}(1 - R)$ with high probability

as $n \rightarrow \infty$, where h_q is the q -ary entropy function [41, Theorem 25 and 26]. For random binary codes with rate $1/2$ we find $\delta_2 = h_2^{-1}(1/2) \approx 0.1100$.

For simplicity, assume $r \leq \gamma d$ for a fixed constant $\gamma < 1/2$. Substituting the above into Theorem 7, the number of output EDCP states in the reduction is $\ell = \Theta(\log n/r)$. It follows that we can only produce, at best, a constant number of EDCP states when the error weight is $r = \Theta(\log n)$.

Note that, for general linear codes, bounded-distance decoding via ISD has complexity exponential in the error weight r . When $r = \Theta(\log n)$, this already yields a classical polynomial time decoding algorithm. Thus this reduction should be viewed only as a conceptual instantiation of our framework.

Relation to LPN. Consider an LPN instance $(\mathbf{A}, \mathbf{b} = \mathbf{A}\mathbf{s} + \mathbf{e})$ with secret dimension n , $m = \text{poly}(n)$ samples, and independent Bernoulli errors with rate $\eta = \Theta(\log^2(n)/n)$. Brakerski, Lyubashevsky, Vaikuntanathan, and Wichs [5] use code smoothing to place LPN at this noise scale in Search-BPP^{SZK}; i.e. they show LPN with these parameters can be solved in randomized polynomial time, given a statistical zero knowledge (SZK) oracle. Independently, classical information-set decoding solves LPN in this regime in quasi-polynomial time $2^{O(\log^2 n)}$.

Compare this with our reduction: with high probability, the code $\mathcal{C} = \text{im}(\mathbf{A}) \subseteq \mathbb{F}_2^m$ has dimension n and minimum distance $d = \Theta(m)$, and the error lies within a decoding radius $r = \Theta(\eta m)$. Substituting these parameters and $q = M = 2$ into the sample bound of Theorem 7 gives

$$\frac{(d - 2r) \log m}{r \log(2^{n+1})} = \Theta\left(\frac{\log n}{\eta n}\right) = \Theta\left(\frac{1}{\log n}\right).$$

Thus, Theorem 7 at these parameters does not guarantee even one EDCP sample with inverse-polynomial loss.

6.2 Rank metric

Rank metric codes were introduced in [14], and rediscovered in [17]. Rank-metric cryptosystems became an attractive focus for code-based cryptography because they offer much smaller public keys and different than Hamming-metric schemes.

Let $\mathcal{B} = \{\beta_1, \dots, \beta_m\}$ be an ordered basis of $\mathbb{F}_{q^m}$ over $\mathbb{F}_q$. We define the $\mathcal{B}$ -extension $\Psi_{\mathcal{B}}$ as

$$\Psi_{\mathcal{B}} : \mathbb{F}_{q^m} \rightarrow \mathbb{F}_q^m, \quad \Psi_{\mathcal{B}}(x) = \begin{pmatrix} x^{(1)} \\ x^{(2)} \\ \vdots \\ x^{(m)} \end{pmatrix}, \quad \text{where } x = \sum_{j=1}^m \beta_j x^{(j)}.$$

We will also denote for any $\mathbf{x} = (x_1, \dots, x_n) \in \mathbb{F}_{q^m}^n$

$$\Psi_{\mathcal{B}}(\mathbf{x}) = (\Psi_{\mathcal{B}}(x_1) \cdots \Psi_{\mathcal{B}}(x_n)), \quad \text{where } x_i = \sum_{j=1}^m \beta_j x_i^{(j)}.$$

Definition 8 (Rank Weight, Rank Distance). Let $\mathbf{x} = (x_1, \dots, x_n) \in \mathbb{F}_{q^m}^n$ and let $\mathcal{B}$ be an ordered basis for $\mathbb{F}_{q^m}$ over $\mathbb{F}_q$. Then the $\mathbb{F}_q$ -rank weight, or simply rank weight if the field is clear from context, of $\mathbf{x}$ is

$$\text{wt}_{\text{rank}}(\mathbf{x}) = \text{rank}(\Psi_{\mathcal{B}}(\mathbf{x})).$$

The rank distance between $\mathbf{x}$ and $\mathbf{y} \in \mathbb{F}_{q^m}^n$ is defined by

$$d_{\text{rank}}(\mathbf{x}, \mathbf{y}) = \text{wt}_{\text{rank}}(\mathbf{x} - \mathbf{y}).$$

We note that the rank weight/distance is independent of the choice of a basis.

In this section, we provide a reduction from BDD in the rank metric to EDCP. We first show the existence of a simple (∞, p) -convex LSH family for the rank metric in Theorem 8 and Lemma 8. We then use this to conclude our result in Theorem 9.

Theorem 8 (Rank metric LSH). For any $0 \leq r_1 < r_2 \leq n$ and integer $t \geq 1$ there exists a $(r_1, r_2, q^{-tr_1}, q^{-tr_2})$ -sensitive LSH for $(\mathbb{F}_q^{m \times n}, d_{\text{rank}})$.

We defer the proof of Theorem 8 to Appendix G.

Remark 4. The LSH family of Theorem 8 has exponent $\rho = \log(1/p_1)/\log(1/p_2) = r_1/r_2$. When $r_2 = cr_1$ this is $1/c$, like the bit-sampling LSH discussed in Subsection 2.1. Unlike the bit-sampling LSH, it is not clear if this is optimal.

Lemma 8. The LSH family of Theorem 8 is (∞, p) -convex.

Proof. We have $\exp(\mathbb{F}_q) = p$ and each hash $h : \mathbb{F}_q^{m \times n} \rightarrow \mathbb{F}_q^{t \times n}$ is a linear transformation, hence a homomorphism. Then $\mathcal{F}$ is (∞, p) -convex by Lemma 3. $\square$

Theorem 9 (Rank $\text{DP}_{\mathcal{C}, r}$ to $\text{EDCP}_{\mathcal{C}, M}$). Let $\mathcal{C} \subseteq \mathbb{F}_{q^m}^n$ be a $[n, k, d]$ -rank metric code over $\mathbb{F}_{q^m}$ with $q = p^N$ and $2 \leq M \leq p$ be an integer. Given a rank $\text{DP}_{\mathcal{C}, r}$ instance $\hat{\mathbf{c}} = \mathbf{c} + \mathbf{e}$ with $\mathbf{c} \in \mathcal{C}$ and $\text{wt}_{\text{rank}}(\mathbf{e}) \leq r < d/2$, there is an efficient procedure that outputs ℓ states of $\text{EDCP}_{\mathcal{C}, M}$ with secret $\mathbf{c}$, where

$$\ell = \Theta \left(\frac{(d - 2r) \log n}{r \log(Mq^{mk})} \right).$$

Proof. Let $\mathcal{F}$ be the LSH family from Theorem 8 with $r_1 = r$, $r_2 = d - r$, and $t \geq (\log_q(M) + mk)/(d - 2r)$ so that $p_1 > M(q^{mk} - 1)p_2$. Let $\mathbf{E} = \Psi_{\mathcal{B}}(\mathbf{e})$ with $\text{rank}(\mathbf{E}) \leq r$. Like the Hamming metric, we have $\text{rank}(j\mathbf{E}) = \text{rank}(\mathbf{E})$ when $j \in \mathbb{F}_p^\times$, so the inequality

$$\text{rank}(\mathbf{E}) \leq r_1 < \frac{d}{2} < r_2 \leq d - \text{rank}(\mathbf{E})$$

is satisfied and we can use Corollary 1.

Now set

$$t = \left\lceil \frac{\log_q(2M) + mk}{d - 2r} \right\rceil$$

so that $p_0 := p_1 - Mq^{mk}p_2 = p_1/2$. Since $p_{\text{iso}} \geq p_1/2$ we obtain

$$\ell = \Theta\left(\frac{\log n}{\log(1/p_1)}\right) = \Theta\left(\frac{\log n}{rt \log q}\right)$$

output EDCP states. Substitute t to complete the proof. $\square$

For a concrete parameter regime, consider rank-metric codes with $m = \Theta(n)$, $d = \Theta(n)$, and fixed q and M . When $r \leq \gamma d$ for a fixed constant $\gamma < 1/2$, Theorem 9 simplifies to

$$\ell_R = \Theta\left(\frac{\log n}{rk}\right).$$

Thus, for the natural rank-metric LSH used here, obtaining even a constant number of EDCP samples requires $rk = O(\log n)$. In the usual cryptographic rank-metric regimes, where k and r are much larger, the reduction does not produce useful EDCP inputs. As in the Hamming case, this instantiation is therefore best viewed as evidence that the LSH framework applies beyond Euclidean geometry, rather than as a reduction from standard hard rank-decoding parameters.

7 Conclusion and Future Work

We introduced a locality-sensitive hashing framework for quantum reductions from bounded distance decoding to EDCP. The framework abstracts the partition-and-measure step appearing in the reductions of Regev and Brakerski–Kirshanova–Stehlé–Wen, replacing Euclidean partitions by LSH families satisfying suitable near-collision, far-collision, and (r, M) -convexity properties. This gives a general criterion for when a BDD instance over a subgroup $\mathcal{D}$ of an Abelian group can be transformed into EDCP states with the same hidden secret.

The p -Lee instantiation in Subsection 5.1 recovers the LWE-to-EDCP connection through the known equivalence between Construction-A lattice decoding in the ℓ_p norm and decoding in the p -Lee metric. This shows that the Regev/BKSW reduction can be viewed as a special case of the LSH framework. We also instantiated the framework in the Hamming and rank metrics in Section 6. These examples demonstrate that the mechanism is not specific to Euclidean geometry, but they also show a limitation: with the specific LSH families used here, the reductions only produce useful EDCP inputs in parameter regimes where known generic decoding algorithms are already efficient. Thus, the Hamming and rank instantiations should be viewed as conceptual evidence for the scope of the framework, rather than as reductions from standard cryptographic decoding parameters.

The central limitation of our present analysis is its reliance on isolated progressions: all points of the desired progression must share a hash bucket, while no point from any competing progression may lie in that bucket. We certify this event using

$$p_{\text{iso}} = p_1 - M(|\mathcal{D}| - 1)p_2.$$

Suppressing false collisions through amplification also decreases the probability of preserving the desired progression. In the constant-rate Hamming benchmark considered in Section 6.1, our chosen amplification gives $p_{\text{iso}} = \exp(-\Theta(r))$, an exponential loss in the error weight already for a single sample. This resembles the error-avoidance cost of information-set decoding (ISD). However, an error-free information set already determines a classically verifiable decoding candidate, whereas our black-box guarantee for an ℓ -sample EDCP solver incurs the factor p_{iso}^ℓ . Consequently, a non-negligible one-sample guarantee does not by itself yield a decoding algorithm competitive with ISD.

A first direction for future work is to understand whether this limitation is inherent to the chosen hash families, the requirement of progression isolation, or the use of pairwise collision bounds in the analysis. As discussed in Remark 3, the coordinate-sampling LSH used for the Hamming metric is optimal in the usual data-independent nearest-neighbor sense, but this optimality does not immediately imply optimality for our reduction. It would be interesting to determine whether known LSH lower bounds can be translated into lower bounds for isolation-based BDD-to-EDCP reductions, or whether this setting leads to a different notion of optimality.

A second direction is to look beyond the data-independent LSH families used in this work. The Hamming and rank instantiations use hash families that depend only on the ambient metric space. Cryptographic instances often contain additional algebraic structure, such as quasi-cyclic, module, ideal, or rank-metric structure. It is natural to ask whether data-dependent or structure-aware hash families can improve the isolation probability p_{iso} , even if they do not improve standard nearest-neighbor search in the ambient space. Such hash families could potentially yield more meaningful reductions for code-based or structured lattice-based parameters.

A complementary direction is to replace progression isolation by a distributional analysis. Brakerski, Lyubashevsky, Vaikuntanathan, and Wichs [5] exploit code balancedness through Fourier analysis: for sufficiently balanced codes of dimension n , random sums of $\Theta(n/\log n)$ sample equations yield nearly uniform coefficient vectors while retaining an inverse-polynomial error bias at noise rate $\eta = \Theta(\log^2(n)/n)$. Their approach suggests investigating whether smoothing can be incorporated into a quantum construction, or whether non-isolated hash buckets retain useful information that our present analysis does not exploit. Such an extension would require guarantees on the resulting quantum states and a way to convert them into EDCP inputs, or to use them in a suitable decoding algorithm.

Finally, our reductions are one-way: they reduce BDD to EDCP. In contrast, the work of Brakerski et al. gives an equivalence between LWE and EDCP. The converse direction there relies on Fourier-analytic properties of Gaussian distributions, in particular their stability under the Fourier transform. It remains unclear whether the distributions relevant to the Hamming, rank, or Lee settings support analogous converse reductions. Developing suitable Fourier-analytic tools, potentially informed by code smoothing, would clarify whether

the resulting EDCP problems capture these decoding assumptions as tightly as in the LWE setting.

Acknowledgments. This work was supported, in part, at Clemson University by the State of South Carolina through funding for the Battelle Savannah River Alliance Workforce Development Program.

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this work.

A Proof of Lemma 3

Proof. Fix an arbitrary $b \in \text{im } h$. For any $x, e \in G$ such that $x \in h^{-1}(b)$ and $x + (M-1)e \in h^{-1}(b)$ we have $h(x) = h(x + (M-1)e)$. Since h is a homomorphism, this implies $h(x) = h(x) + (M-1)h(e)$ so $(M-1)h(e) = 0_H$.

Assume $h(e) \neq 0_H$. Then $h(e)$ must have order k dividing $M-1$, a contradiction. It follows that $h(x + je) = h(x) + jh(e) = h(x)$, so $h(x) = h(x + je)$ holds for all $j \in [M]$. Since there is no distance constraint, h is (∞, M) -convex. $\square$

B Proof of Theorem 2

Lemma 9. *Assume the hypotheses of Theorem 2. In particular, let $e \in A$ with $\text{wt}(e) \leq r < d/2$ and*

$$\text{wt}((M-1)e) \leq r_1 < \frac{1}{2}d < r_2 \leq d - \Delta_M(e).$$

Fix $x \in \mathcal{D}$ and $j \in [M]$. Over $h \leftarrow \mathcal{F}$, with probability at least p_{iso} , the following two conditions simultaneously hold:

- $h(x + je) = h(x + j'e)$ for all $j' \in [M]$, and
- $h(x + je) \neq h(x' + j'e)$ for all $x' \in \mathcal{D} \setminus \{x\}$, $j' \in [M]$.

Proof. Since

$$d(x, x + (M-1)e) = \text{wt}((M-1)e) \leq r_1,$$

by LSH sensitivity we have a good collision with probability

$$\Pr_h[h(x) = h(x + (M-1)e)] \geq p_1.$$

By (r, M) -convexity, this endpoint collision implies

$$h(x) = h(x + j'e) \quad \text{for all } j' \in [M].$$

In particular,

$$h(x + je) = h(x + j'e) \quad \text{for all } j' \in [M].$$

Now fix $x' \in \mathcal{D} \setminus \{x\}$ and $j' \in [M]$. Since $\mathcal{D}$ has minimum distance d ,

$$\text{wt}(x - x') \geq d.$$

Moreover,

$$d(x + je, x' + j'e) = \text{wt}(x - x' + (j - j')e) \geq d - \text{wt}((j - j')e) \geq d - \Delta_M(e) \geq r_2.$$

Therefore, by LSH sensitivity, the probability of a false collision is

$$\Pr_h[h(x + je) = h(x' + j'e)] \leq p_2.$$

Taking a union bound over the $M(|\mathcal{D}| - 1)$ choices of $(x', j') \in (\mathcal{D} \setminus \{x\}) \times [M]$ and using the assumption $1 \geq p_1 > M(|\mathcal{D}| - 1)p_2$, the probability of any false collision is at most $M(|\mathcal{D}| - 1)p_2$. By another union bound, combining this with the good-collision event gives the claimed lower bound

$$p_{\text{iso}} = p_1 - M(|\mathcal{D}| - 1)p_2 \geq 0.$$

□

Lemma 10. *Assume the hypotheses of Theorem 2. There is an efficient randomized quantum procedure $\text{Gen}(t)$ whose output has density operator ρ_s satisfying*

$$\rho_s \succeq p_{\text{iso}} \sigma_s$$

where $\succeq$ denotes the positive semidefinite order.

Proof. The procedure $\text{Gen}(t)$ is described as follows. First prepare a uniform superposition over $[M]$ (by e.g. [39]). Tensoring this with a uniform superposition over $\mathcal{D}$ yields the initial state

$$\frac{1}{\sqrt{M|\mathcal{D}|}} \sum_{j \in [M]} \sum_{x \in \mathcal{D}} |j\rangle |x\rangle.$$

Compute $x - jt$ in a new register. Since $\mathcal{D}$ is an Abelian group, for every fixed j the map $x \mapsto x + js$ is a bijection. Hence the resulting state can be written as

$$\frac{1}{\sqrt{M|\mathcal{D}|}} \sum_{j \in [M]} \sum_{x \in \mathcal{D}} |j\rangle |x + js\rangle |x - je\rangle.$$

Sample a hash function $h \leftarrow U(\mathcal{F})$ and compute $h(x - je)$ in a new register to obtain the state

$$\frac{1}{\sqrt{M|\mathcal{D}|}} \sum_{j \in [M]} \sum_{x \in \mathcal{D}} |j\rangle |x + js\rangle |x - je\rangle |h(x - je)\rangle.$$

Measure the fourth register to obtain some hash value c , and uncompute the third register by applying the map

$$|j\rangle |a\rangle |b\rangle \mapsto |j\rangle |a\rangle |b - (a - jt)\rangle$$

to the first three registers. The resulting unnormalized output state associated to the hash value c is

$$|\Phi_{h,c}\rangle := \frac{1}{\sqrt{M|\mathcal{D}|}} \sum_{\substack{j \in [M], x \in \mathcal{D} \\ h(x-je)=c}} |j\rangle |x+js\rangle.$$

Hence, for fixed h , the post-measurement output density operator is

$$\rho_s(h) = \sum_c |\Phi_{h,c}\rangle \langle \Phi_{h,c}|,$$

and the averaged output state is

$$\rho_s = \mathbb{E}_{h \leftarrow \mathcal{F}} [\rho_s(h)].$$

Fix $x \in \mathcal{D}$. Apply Lemma 9 with $-e \in A$. Since $\text{wt}(-x) = \text{wt}(x)$ for all x the hypotheses are unchanged. Therefore, with probability at least p_{iso} over $h \leftarrow \mathcal{F}$, the progression

$$x, x-e, \dots, x-(M-1)e$$

is isolated: all points $x-je$, $j \in [M]$, have the same hash value, and no point $x'-je$ with $x' \in \mathcal{D} \setminus \{x\}$ and $j \in [M]$ has that hash value. Denote this isolation event by $I_x(h)$. With this notation, Lemma 9 says: $\Pr_h[I_x(h)] \geq p_{\text{iso}}$.

When $I_x(h)$ holds, let c_x denote the common hash value of the isolated progression. Then the corresponding hash bucket contributes exactly

$$|\Phi_{h,c_x}\rangle = \frac{1}{\sqrt{M|\mathcal{D}|}} \sum_{j \in [M]} |j\rangle |x+js\rangle = \frac{1}{\sqrt{|\mathcal{D}|}} |\psi_{x,s}\rangle$$

to the measured output. Thus $\rho_s(h)$ contains the positive semidefinite term

$$|\Phi_{h,c_x}\rangle \langle \Phi_{h,c_x}| = \frac{1}{|\mathcal{D}|} |\psi_{x,s}\rangle \langle \psi_{x,s}|$$

whenever $I_x(h)$ occurs.

Moreover, if $I_x(h)$ and $I_{x'}(h)$ both hold for $x \neq x'$, then the corresponding hash buckets are distinct; otherwise the isolation condition would be violated. All remaining hash buckets contribute positive semidefinite operators. Hence, for every fixed h ,

$$\rho_s(h) \succeq \frac{1}{|\mathcal{D}|} \sum_{x \in \mathcal{D}} \mathbf{1}_{I_x(h)} |\psi_{x,s}\rangle \langle \psi_{x,s}|,$$

where $\mathbf{1}_{I_x(h)}$ is the indicator of the event $I_x(h)$.

Averaging over h , we obtain

$$\rho_s = \mathbb{E}_h[\rho_s(h)] \succeq \frac{1}{|\mathcal{D}|} \sum_{x \in \mathcal{D}} \Pr_h[I_x(h)] |\psi_{x,s}\rangle \langle \psi_{x,s}| \succeq \frac{p_{\text{iso}}}{|\mathcal{D}|} \sum_{x \in \mathcal{D}} |\psi_{x,s}\rangle \langle \psi_{x,s}| = p_{\text{iso}} \sigma_s.$$

□

Lemma 11. *Let ρ, σ be positive semidefinite operators and let $p \geq 0$. If*

$$\rho \succeq p\sigma,$$

then, for every $\ell \geq 1$,

$$\rho^{\otimes \ell} \succeq p^\ell \sigma^{\otimes \ell}.$$

Proof. Since $\rho \succeq p\sigma$, there exists a positive semidefinite operator $R \succeq 0$ such that

$$\rho = p\sigma + R.$$

Taking tensor powers and expanding,

$$\rho^{\otimes \ell} = (p\sigma + R)^{\otimes \ell} = p^\ell \sigma^{\otimes \ell} + R',$$

where $R' \succeq 0$, since every remaining term in the tensor expansion is a tensor product of positive semidefinite operators and is therefore positive semidefinite. Hence

$$\rho^{\otimes \ell} \succeq p^\ell \sigma^{\otimes \ell}.$$

□

Proof of Theorem 2. On input t , run the procedure $\text{Gen}(t)$ from Lemma 10 independently ℓ times. By Lemma 11, the resulting state is

$$\rho_s^{\otimes \ell} \succeq p_{\text{iso}}^\ell \sigma_s^{\otimes \ell}.$$

Feed the generated state $\rho_s^{\otimes \ell}$ to $\mathcal{A}_{\text{EDCP}}$. Let Π_s be the POVM element corresponding to the event that $\mathcal{A}_{\text{EDCP}}$ outputs s . Since $\mathcal{A}_{\text{EDCP}}$ succeeds with probability at least ε on ideal EDCP samples, we have

$$\text{Tr}(\Pi_s \sigma_s^{\otimes \ell}) \geq \varepsilon.$$

Because $\Pi_s \succeq 0$, the above implies

$$\text{Tr}(\Pi_s \rho_s^{\otimes \ell}) \geq p_{\text{iso}}^\ell \text{Tr}(\Pi_s \sigma_s^{\otimes \ell}) \geq p_{\text{iso}}^\ell \varepsilon.$$

Thus the full procedure outputs the correct secret s with probability at least $p_{\text{iso}}^\ell \varepsilon$.

Note that we can easily verify a candidate $\hat{s}$ returned by this procedure by checking that $\hat{s} \in \mathcal{D}$ and $\text{wt}(t - \hat{s}) \leq r$, and such an $\hat{s}$ is necessarily unique. Therefore, if $p_{\text{iso}}^\ell \varepsilon$ is inverse-polynomial in λ , then polynomially many independent repetitions give constant success probability, and the overall reduction runs in quantum polynomial time. □

C Proof of Lemma 6

Proof. A uniformly random matrix $\mathbf{A} \in \mathbb{Z}_q^{m \times n}$ has full rank with probability at least $1 - q^{-(m-n)}$. If $\mathbf{A}$ has full rank then $\det(\Lambda_q(\mathbf{A})) = q^{m-n}$.

By Minkowski's theorem, there exists a non-zero vector $\mathbf{x} \in \Lambda_q(\mathbf{A})$ such that $\|\mathbf{x}\|_\infty \leq \det(\Lambda_q(\mathbf{A}))^{1/m}$. Then with probability $1 - q^{-(m-n)}$ we have

$$\lambda_1^\infty(\Lambda_q(\mathbf{A})) \leq (q^{m-n})^{1/m} = q^{1-n/m} < q.$$

Note that we can write $\Lambda_q(\mathbf{A})$ as a union of cosets of $q\mathbb{Z}^n$:

$$\Lambda_q(\mathbf{A}) = \bigcup_{\mathbf{c} \in \mathcal{C}(\mathbf{A})} (\mathbf{c} + q\mathbb{Z}^n).$$

By the above, this means

$$\lambda_1^\infty(\Lambda_q(\mathbf{A})) = \min_{\mathbf{c} \in \mathcal{C}(\mathbf{A}) \setminus \{\mathbf{0}\}} \left(\min_{\mathbf{x} \in \mathbf{c} + q\mathbb{Z}^n} \|\mathbf{x}\|_\infty \right)$$

with probability $1 - q^{-(m-n)}$. As noted in [3] (discussed in Subsection 5.1) and the definition of minimum distance, the right-hand side is precisely $d_{\infty, \text{Lee}}(\mathcal{C}(\mathbf{A}))$. $\square$

D Proof of Corollary 1

Proof. Prepare the uniform superposition

$$\frac{1}{\sqrt{q^k}} \sum_{\mathbf{x} \in \mathbb{F}_q^k} |\mathbf{x}\rangle.$$

Let $\mathbf{G} \in \mathbb{F}_q^{n \times k}$ be a generator matrix for $\mathcal{C}$. Because $\mathbf{G}$ has full rank, it has a left inverse $\mathbf{H}$ such that $\mathbf{H}\mathbf{G} = \mathbf{I}_k$. Then we can implement a unitary that performs the transformation $|\mathbf{x}\rangle |\mathbf{0}\rangle \mapsto |\mathbf{x}\rangle |\mathbf{x}\mathbf{G}^T\rangle \mapsto |\mathbf{x} - \mathbf{x}(\mathbf{H}\mathbf{G})^T\rangle |\mathbf{x}\mathbf{G}^T\rangle = |\mathbf{0}\rangle |\mathbf{x}\mathbf{G}^T\rangle$, resulting in the state

$$\frac{1}{\sqrt{|\mathcal{C}|}} \sum_{\mathbf{c} \in \mathcal{C}} |\mathbf{c}\rangle.$$

Since matrix multiplication requires $\mathcal{O}(nk)$ operations over $\mathbb{F}_q$, this can be implemented with $\mathcal{O}(nk \cdot \text{poly}(\log q))$ gates. Finally, since $\exp(\mathbb{F}_q^n) = p$ we can apply Theorem 2, noting that $|\mathcal{C}| = q^k$. $\square$

E Proof of Theorem 3

Proof. Define $\mathcal{F} = \{h : \mathbb{Z}_q^n \rightarrow \mathbb{R}_+^n\}$ where h is the concatenation of h_{i, Δ_i} for $i \in [n]$ and Δ_i are i.i.d. uniform from $[0, w)$. Assume $\mathbf{x}, \mathbf{y} \in \mathbb{Z}_q^n$ with $\text{wt}_{p, \text{Lee}}(\mathbf{x} - \mathbf{y}) \leq r_1$, and define $t_i = \text{wt}_{\text{Lee}}(x_i - y_i)$. For all $i \in [n]$

$$t_i \leq \|(t_1, \dots, t_n)\|_p = \text{wt}_{p, \text{Lee}}(\mathbf{x} - \mathbf{y}) \leq r_1.$$

Since $r_1 < w$, each single-coordinate probability from Lemma 4 satisfies

$$\max\left(0, 1 - \frac{t_i}{w}\right) = 1 - \frac{t_i}{w} \geq 1 - \frac{r_1}{w}$$

so $\Pr[h(\mathbf{x}) = h(\mathbf{y})] \geq (1 - r_1/w)^n$ over the randomness of the Δ_i .

Similarly, assume $\text{wt}_{p,\text{Lee}}(\mathbf{x} - \mathbf{y}) \geq r_2$ and define $t_{\max} = \max_i t_i$. Then

$$r_2 \leq \text{wt}_{p,\text{Lee}}(\mathbf{x} - \mathbf{y}) = \|(t_1, \dots, t_n)\|_p \leq n^{1/p} t_{\max}$$

so $t_{\max} \geq r_2/n^{1/p}$.

Let $j \in [n]$ be an index such that $t_j = t_{\max}$. Then

$$\max\left(0, 1 - \frac{t_j}{w}\right) \leq \max\left(0, 1 - \frac{r_2}{n^{1/p}w}\right)$$

and $\max(0, 1 - t_i/w) \leq 1$ for all other i . Therefore

$$\Pr[h(\mathbf{x}) = h(\mathbf{y})] \leq \max\left(0, 1 - \frac{r_2}{n^{1/p}w}\right)$$

over the randomness of the Δ_i . $\square$

F Proof of Theorem 6

Proof. For $i \in [n]$ define $g_i : \mathbb{F}_q^n \rightarrow \mathbb{F}_q$ by $g_i(\mathbf{x}) = x_i$. Set $\mathcal{G} = \{g_0, \dots, g_{n-1}\}$. If $\text{wt}_H(\mathbf{x} - \mathbf{y}) \leq r_1$ then

$$\Pr_{h_i \leftarrow U(\mathcal{G})}[g_i(\mathbf{x}) = g_i(\mathbf{y})] = 1 - \frac{\text{wt}_H(x_i - y_i)}{n} \geq 1 - \frac{r_1}{n}.$$

Similarly, if $\text{wt}_H(\mathbf{x} - \mathbf{y}) \geq r_2$ then $\Pr_{g_i \leftarrow U(\mathcal{G})}[g_i(\mathbf{x}) = g_i(\mathbf{y})] \leq 1 - r_2/n$, so $\mathcal{G}$ is $(r_1, r_2, 1 - r_1/n, 1 - r_2/n)$ -sensitive.

Now let $\mathcal{F}$ be the set of all $h(\mathbf{x}) = (g_{i_1}(\mathbf{x}), \dots, g_{i_t}(\mathbf{x}))$ where the functions g_i are i.i.d. uniform from $\mathcal{G}$, with replacement. Then $\mathcal{F}$ is $(r_1, r_2, (1 - r_1/n)^t, (1 - r_2/n)^t)$ -sensitive. $\square$

G Proof of Theorem 8

Proof. Let $\Psi_{\mathcal{B}}(\mathbf{x}) := \mathbf{X}$ and $\Psi_{\mathcal{B}}(\mathbf{y}) := \mathbf{Y}$ for $\mathbf{x}, \mathbf{y} \in \mathbb{F}_{q^m}^n$. Sample $\mathbf{A} \leftarrow \mathbb{F}_q^{t \times m}$ with i.i.d. uniform rows and define the hash $h_{\mathbf{A}} : \mathbb{F}_{q^m}^n \rightarrow \mathbb{F}_q^{t \times n}$ by

$$h_{\mathbf{A}}(\mathbf{x}) = \mathbf{A}\mathbf{X} \in \mathbb{F}_q^{t \times n}.$$

We now show that if $d(\mathbf{x}, \mathbf{y}) \leq r$, then $h_{\mathbf{A}}(\mathbf{x}) = h_{\mathbf{A}}(\mathbf{y})$ with probability q^{-tr} . Denote $\mathbf{M} = \mathbf{X} - \mathbf{Y}$. If $\text{rank}(\mathbf{M}) = r$, we have $h_{\mathbf{A}}(\mathbf{x}) = h_{\mathbf{A}}(\mathbf{y})$ iff $\mathbf{A}\mathbf{M} = \mathbf{0}$. The left kernel of $\mathbf{M}$ has dimension $m - r$, so for a uniform row $\mathbf{a}$ of $\mathbf{A}$ we have $\Pr[\mathbf{a}\mathbf{M} = \mathbf{0}] = q^{m-r}/q^m = q^{-r}$. Since the rows are independent, we have $\Pr_{\mathbf{A}}[h_{\mathbf{A}}(\mathbf{x}) = h_{\mathbf{A}}(\mathbf{y})] = \Pr[\mathbf{A}\mathbf{M} = \mathbf{0}] = q^{-tr}$. $\square$

References

1. Bai, S., Jangir, H., Kirshanova, E., Ngo, T., Youmans, W.: A quasi-polynomial time algorithm for the extrapolated dihedral coset problem over power-of-two moduli. In: Kalai, Y.T., Kamara, S.F. (eds.) *Advances in Cryptology - CRYPTO 2025 - 45th Annual International Cryptology Conference*, Santa Barbara, CA, USA, August 17-21, 2025, Proceedings, Part II. *Lecture Notes in Computer Science*, vol. 16001, pp. 416–448. Springer (2025). https://doi.org/10.1007/978-3-032-01878-6_14, https://doi.org/10.1007/978-3-032-01878-6_14
2. Berlekamp, E.R., McEliece, R.J., van Tilborg, H.C.A.: On the inherent intractability of certain coding problems. *IEEE Transactions on Information Theory* **24**(3), 384–386 (May 1978). <https://doi.org/10.1109/TIT.1978.1055873>, funding by NASA/JPL/Caltech
3. Blomqvist, F.: On quotient metrics and decoding problems. In: *2018 International Symposium on Information Theory and Its Applications (ISITA)*. pp. 198–202 (2018). <https://doi.org/10.23919/ISITA.2018.8664238>
4. 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
5. Brakerski, Z., Lyubashevsky, V., Vaikuntanathan, V., Wichs, D.: Worst-case hardness for LPN and cryptographic hashing via code smoothing. In: Ishai, Y., Rijmen, V. (eds.) *EUROCRYPT 2019, Part III*. LNCS, vol. 11478, pp. 619–635. Springer, Cham (May 2019). https://doi.org/10.1007/978-3-030-17659-4_21
6. Castryck, W., Dooms, A., Emerencia, C., Lemmens, A.: A fusion algorithm for solving the hidden shift problem in finite abelian groups. In: Cheon, J.H., Tillich, J.P. (eds.) *Post-Quantum Cryptography - 12th International Workshop, PQCrypto 2021*. pp. 133–153. Springer, Cham (2021). https://doi.org/10.1007/978-3-030-81293-5_8
7. Chailloux, A., Hermouet, P.: On the quantum equivalence between $s|lwe\rangle$ and ISIS. *CoRR* **abs/2510.06097** (2025). <https://doi.org/10.48550/ARXIV.2510.06097>, <https://doi.org/10.48550/arXiv.2510.06097>
8. Chailloux, A., Tillich, J.: The quantum decoding problem. In: Magniez, F., Grilo, A.B. (eds.) *19th Conference on the Theory of Quantum Computation, Communication and Cryptography, TQC 2024*, Okinawa, Japan, September 9-13, 2024. pp. 6:1–6:14. *LIPICs, Schloss Dagstuhl - Leibniz-Zentrum für Informatik* (2024). <https://doi.org/10.4230/LIPICS.TQC.2024.6>, <https://doi.org/10.4230/LIPICS.TQC.2024.6>
9. 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>
10. 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
11. Costa, S.I.R., Campello, A., Jorge, G.C., Strapasson, J., Qureshi, C.M.: Codes and lattices in the lp metric. In: *2014 Information Theory and Applications Workshop, ITA 2014*, San Diego, CA, USA, February 9-14, 2014. pp. 1–4. IEEE (2014). <https://doi.org/10.1109/ITA.2014.6804217>, <https://doi.org/10.1109/ITA.2014.6804217>

12. Datar, M., Immorlica, N., Indyk, P., Mirrokni, V.S.: Locality-sensitive hashing scheme based on p-stable distributions. In: Snoeyink, J., Boissonnat, J. (eds.) *Proceedings of the 20th ACM Symposium on Computational Geometry*, Brooklyn, New York, USA, June 8–11, 2004. pp. 253–262. ACM (2004). <https://doi.org/10.1145/997817.997857>, <https://doi.org/10.1145/997817.997857>
13. Debris-Alazard, T., Remaud, M., Tillich, J.: Quantum reduction of finding short code vectors to the decoding problem. *IEEE Trans. Inf. Theory* **70**(7), 5323–5342 (2024). <https://doi.org/10.1109/TIT.2023.3327759>, <https://doi.org/10.1109/TIT.2023.3327759>
14. Delsarte, P.: Bilinear forms over a finite field, with applications to coding theory. *Journal of Combinatorial Theory, Series A* **25**(3), 226–241 (1978). [https://doi.org/10.1016/0097-3165\(78\)90015-8](https://doi.org/10.1016/0097-3165(78)90015-8), <https://www.sciencedirect.com/science/article/pii/0097316578900158>
15. Dougherty, S.T.: *Algebraic Coding Theory Over Finite Commutative Rings*. SpringerBriefs in Mathematics, Springer (2017). <https://doi.org/10.1007/978-3-319-59806-2>
16. Elsheikh, M., Giesbrecht, M., Novocin, A., Saunders, B.D.: Fast computation of smith forms of sparse matrices over local rings. In: van der Hoeven, J., van Hoeij, M. (eds.) *International Symposium on Symbolic and Algebraic Computation, ISSAC’12, Grenoble, France - July 22 - 25, 2012*. pp. 146–153. ACM (2012). <https://doi.org/10.1145/2442829.2442853>, <https://doi.org/10.1145/2442829.2442853>
17. Gabidulin, E.: Theory of codes with maximum rank distance (translation). *Problems of Information Transmission* **21**, 1–12 (01 1985)
18. Gaborit, P., Aguilar-Melchor, C., Aragon, N., Bettaleb, S., Bidoux, L., Blazy, O., Deneuville, J.C., Persichetti, E., Zémor, G., Bos, J., Dion, A., Lacan, J., Robert, J.M., Véron, P., Barreto, P.L., Ghosh, S., Gueron, S., Güneysu, T., Misoczki, R., Richter-Brokmann, J., Sendrier, N., Tillich, J.P., Vasseur, V.: *Hamming quasi-cyclic (HQC)* (2025)
19. Gionis, A., Indyk, P., Motwani, R.: Similarity search in high dimensions via hashing. In: Atkinson, M.P., Orłowska, M.E., Valduriez, P., Zdonik, S.B., Brodie, M.L. (eds.) *VLDB’99, Proceedings of 25th International Conference on Very Large Data Bases*, September 7–10, 1999, Edinburgh, Scotland, UK. pp. 518–529. Morgan Kaufmann (1999), <http://www.vldb.org/conf/1999/P49.pdf>
20. Hamming, R.W.: Error detecting and error correcting codes. *The Bell System Technical Journal* **29**(2), 147–160 (1950). <https://doi.org/10.1002/j.1538-7305.1950.tb00463.x>
21. Hörmann, F., van Woerden, W.P.J.: Fuleakage: Breaking fuleeca by learning attacks. In: Reyzin, L., Stebila, D. (eds.) *Advances in Cryptology - CRYPTO 2024 - 44th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 18–22, 2024, Proceedings, Part VI. Lecture Notes in Computer Science*, vol. 14925, pp. 253–286. Springer (2024). https://doi.org/10.1007/978-3-031-68391-6_8, https://doi.org/10.1007/978-3-031-68391-6_8
22. Indyk, P., Motwani, R.: Approximate nearest neighbors: Towards removing the curse of dimensionality. In: Vitter, J.S. (ed.) *Proceedings of the Thirtieth Annual ACM Symposium on the Theory of Computing*, Dallas, Texas, USA, May 23–26, 1998. pp. 604–613. ACM (1998). <https://doi.org/10.1145/276698.276876>, <https://doi.org/10.1145/276698.276876>
23. Jorge, G.C., Campello, A., Costa, S.I.: q -ary lattices in the l_p norm and a generalization of the lee metric. In: *Proceedings of the International Workshop on Coding and Cryptography* (2013)

24. Kuperberg, G.: A subexponential-time quantum algorithm for the dihedral hidden subgroup problem. *SIAM J. Comput.* **35**(1), 170–188 (2005). <https://doi.org/10.1137/S0097539703436345>, <https://doi.org/10.1137/S0097539703436345>
25. Kuperberg, G.: Another subexponential-time quantum algorithm for the dihedral hidden subgroup problem. In: Severini, S., Brandão, F.G.S.L. (eds.) 8th Conference on the Theory of Quantum Computation, Communication and Cryptography, TQC 2013, Guelph, Canada, May 21–23, 2013. *LIPIcs*, vol. 22, pp. 20–34. Schloss Dagstuhl - Leibniz-Zentrum für Informatik (2013). <https://doi.org/10.4230/LIPIcs.TQC.2013.20>, <https://doi.org/10.4230/LIPIcs.TQC.2013.20>
26. Laarhoven, T.: Sieving for shortest vectors in lattices using angular locality-sensitive hashing. In: Gennaro, R., Robshaw, M. (eds.) *Advances in Cryptology - CRYPTO 2015 - 35th Annual Cryptology Conference*, Santa Barbara, CA, USA, August 16–20, 2015, *Proceedings, Part I. Lecture Notes in Computer Science*, vol. 9215, pp. 3–22. Springer (2015). https://doi.org/10.1007/978-3-662-47989-6_1, https://doi.org/10.1007/978-3-662-47989-6_1
27. Lee, C.: Some properties of nonbinary error-correcting codes. *IRE Transactions on Information Theory* **4**(2), 77–82 (1958). <https://doi.org/10.1109/TIT.1958.1057446>
28. 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
29. Lyubashevsky, V., Micciancio, D.: On bounded distance decoding, unique shortest vectors, and the minimum distance problem. In: Halevi, S. (ed.) *CRYPTO 2009. LNCS*, vol. 5677, pp. 577–594. Springer, Berlin, Heidelberg (Aug 2009). https://doi.org/10.1007/978-3-642-03356-8_34
30. McEliece, R.J.: A public key cryptosystem based on algebraic coding theory (1978), <https://api.semanticscholar.org/CorpusID:56502909>
31. National Institute of Standards and Technology: Module-lattice-based key-encapsulation mechanism standard. Tech. Rep. Federal Information Processing Standards Publication (FIPS) 203, U.S. Department of Commerce, Washington, D.C. (August 2024). <https://doi.org/10.6028/NIST.FIPS.203>, <https://doi.org/10.6028/NIST.FIPS.203>
32. Nielsen, M.A., Chuang, I.L.: *Quantum Computation and Quantum Information: 10th Anniversary Edition*. Cambridge University Press (2010)
33. O’Donnell, R., Wu, Y., Zhou, Y.: Optimal lower bounds for locality-sensitive hashing (except when q is tiny). *ACM Trans. Comput. Theory* **6**(1), 5:1–5:13 (2014). <https://doi.org/10.1145/2578221>, <https://doi.org/10.1145/2578221>
34. Peikert, C.: He gives c -sieves on the CSIDH. In: Canteaut, A., Ishai, Y. (eds.) *Advances in Cryptology - EUROCRYPT 2020 - 39th Annual International Conference on the Theory and Applications of Cryptographic Techniques*, Zagreb, Croatia, May 10–14, 2020, *Proceedings, Part II. Lecture Notes in Computer Science*, vol. 12106, pp. 463–492. Springer (2020). https://doi.org/10.1007/978-3-030-45724-2_16, https://doi.org/10.1007/978-3-030-45724-2_16
35. 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>
36. Regev, O.: A subexponential time algorithm for the dihedral hidden subgroup problem with polynomial space (2004), <https://arxiv.org/abs/quant-ph/0406151>

37. 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>
38. Ritterhoff, S., Bitzer, S., Karl, P., Maringer, G., Schamberger, T., Schupp, J., Sigl, G., Wachter-Zeh, A., Weger, V.: Fuleeca: Algorithm specifications and supporting documentation. Tech. rep., National Institute of Standards and Technology (May 2023), <https://csrc.nist.gov/csrc/media/Projects/pqc-dig-sig/documents/round-1/spec-files/FuLeeca-spec-web.pdf>, NIST Post-Quantum Cryptography Additional Signature Schemes Submission
39. Shukla, A., Vedula, P.: An efficient quantum algorithm for preparation of uniform quantum superposition states. *Quantum Information Processing* **23**(2), 38 (2024). <https://doi.org/10.1007/s11128-024-04258-4>
40. National Institute of Standards and Technology: Module-lattice-based digital signature standard. Tech. Rep. Federal Information Processing Standards Publication (FIPS) 204, National Institute of Standards and Technology (August 2024). <https://doi.org/10.6028/NIST.FIPS.204>, <https://doi.org/10.6028/NIST.FIPS.204>
41. Weger, V., Gassner, N., Rosenthal, J.: A survey on code-based cryptography. *CoRR abs/2201.07119* (2022), <https://arxiv.org/abs/2201.07119>
42. Weger, V., Khathuria, K., Horlemann, A., Battaglioni, M., Santini, P., Persichetti, E.: On the hardness of the lee syndrome decoding problem. *Adv. Math. Commun.* **18**(1), 233–266 (2024). <https://doi.org/10.3934/AMC.2022029>, <https://doi.org/10.3934/amc.2022029>