

Pseudorandom Codes from LWE

Keewoo Lee

Ethereum Foundation
keewoo.lee@ethereum.org

Abstract. A pseudorandom code (PRC), introduced by Christ and Gunn (Crypto 2024), is an error-correcting code whose codewords look uniformly random to anyone without the secret key. PRCs not only are natural cryptographic objects to study in their own right, but also have interesting applications such as undetectable watermarking of AI-generated content. Previous PRC constructions have centered on Learning Parity with Noise (LPN), often combined with additional assumptions. In particular, the original PRC of Christ and Gunn is based either on standard LPN together with the planted XOR assumption of Agrawal et al. (Crypto 2024), or on subexponential LPN.

In this work, we construct PRCs from Learning with Errors (LWE), providing an alternative foundation for PRCs. Our construction is simple: a codeword is the one-bit rounding of an LWE sample, masked by a random string. The construction can be easily understood as an LWE analogue of the PRC of Christ and Gunn. Indeed, like theirs, it is based either on standard LWE together with the hardness of the planted t -SUM problem of Agrawal et al., or on subexponential LWE. However, as the rounding is not linear, its robustness is less straightforward, and we prove it using Fourier analysis.

1 Introduction

A pseudorandom code (PRC), introduced by Christ and Gunn [CG24], is an error-correcting code whose codewords look uniformly random to anyone without the secret key. PRCs not only are natural cryptographic objects to study in their own right, but also have interesting applications such as undetectable watermarking of AI-generated content [CG24, GZS25] and robust steganography [CG24]. Previous PRC constructions have centered on Learning Parity with Noise (LPN) [BFKL94], often combined with additional assumptions. In particular, the original PRC of [CG24] is based either on standard LPN together with the planted XOR assumption of Agrawal et al. [ASS⁺24], or on subexponential LPN.¹ Even recent PRCs robust to adaptive substitutions still rely on LPN, combined with new variants of the McEliece assumption [DJK⁺26, DFV26]. Beyond LPN, PRCs are known only from a few less standard assumptions: log-local

¹ Ghentiyala and Guruswami [GG25] later showed that a noisy variant of the planted XOR assumption suffices, at the cost of LPN with a subconstant noise rate. A similar relaxation applies to our result (Remark 3.9).

weak pseudorandom functions [GM24], permuted puzzles [CGG⁺26], and, with relaxed pseudorandomness, the planted hyperloop assumption [GG25]. We refer to the survey of Christ [Chr26] for an overview of the area.

1.1 Our Contribution

In this work, we construct PRCs from Learning with Errors (LWE) [Reg05], which underlies much of lattice-based cryptography and enjoys a worst-case to average-case reduction. At first sight, LWE seems ill-suited to PRCs. An LPN sample is already binary, and the bit flips of the channel simply add to its sparse Bernoulli noise. In contrast, an LWE sample lies in $\mathbb{Z}_q^m$, and its Gaussian noise is small in every entry rather than sparse. Nevertheless, we show that simply rounding LWE samples to bits does the job.

Our main construction (Figure 1) is a public-key zero-bit weak PRC against the binary symmetric channel BSC_p for every constant $p \in [0, 1/2)$ (Definition 2.4). Its pseudorandomness is based either on standard LWE together with the hardness of the planted t -SUM problem [ASS⁺24] (Theorem 3.8), or on subexponential LWE (Theorem 3.10), mirroring the two LPN-based instantiations of [CG24].² By known transformations [CG24, GM24, GG25, Chr26], these weak PRCs give full PRCs with various robustness guarantees, as well as watermarking and steganography schemes, under the same assumptions (Remark 3.11).

Our result also bears on which assumptions carry enough structure to support PRCs. Döttling, Müller, and Rajasree [DMR25] and Garg, Gunn, and Wang [GGW25] showed that black-box cryptography is useless for PRCs, so every construction must exploit the structure of a specific assumption. Until now, this role has been played mainly by LPN, and we show that LWE can play it as well. We hope that this also helps isolate which properties of LPN previous constructions rely on.

1.2 Technical Overview

Review: The Construction of [CG24]. We describe the public-key zero-bit weak PRC underlying the construction of [CG24]. A zero-bit PRC has no message to recover: its decoder only has to tell whether a received word is a noisy codeword or a string chosen independently of the key. A weak PRC only needs to do so with an inverse-polynomial advantage. The idea of [CG24] is to plant a short parity check in a random linear code. The public key consists of a matrix $\mathbf{G} \in \mathbb{F}_2^{m \times n}$ and a random mask $\mathbf{z} \in \mathbb{F}_2^m$, and the secret key is a set of t rows of $\mathbf{G}$, for a small t , planted to sum to zero. We write $\mathbf{v} \in \{0, 1\}^m$ for its indicator vector, so that $\mathbf{v}^\top \mathbf{G} = \mathbf{0}$. A codeword is $\mathbf{c} = \mathbf{G}\mathbf{s} \oplus \mathbf{e} \oplus \mathbf{z}$ for a random

² As observed in [CGG⁺26], constructions that plant structures of logarithmic size, such as those of [CG24, GM24, GG25], are only quasi-polynomially secure, and so are ours. PRCs with plausible subexponential security rely on non-standard assumptions [CGG⁺26, DJK⁺26, DfV26].

$\mathbf{s} \in \mathbb{F}_2^n$ and a Bernoulli noise vector $\mathbf{e}$, and the decoder accepts a received word $\mathbf{c}'$ if it passes the parity check $\langle \mathbf{v}, \mathbf{c}' \oplus \mathbf{z} \rangle = 0$. This works because the parity check cancels $\mathbf{G}\mathbf{s}$, as $\mathbf{v}^\top \mathbf{G} = \mathbf{0}$. For a noisy codeword $\mathbf{c} \oplus \boldsymbol{\xi}$ with channel flips $\boldsymbol{\xi}$, the parity is $\langle \mathbf{v}, \mathbf{e} \oplus \boldsymbol{\xi} \rangle$, which is biased toward 0 as $\mathbf{e} \oplus \boldsymbol{\xi}$ is a Bernoulli noise vector of rate below $1/2$ and t is small. For a string chosen independently of the key, the parity is a uniform bit as $\mathbf{z}$ is uniform, so the decoder accepts it with probability exactly $1/2$ over the choice of the key. Pseudorandomness holds as long as the planted rows stay hidden, either computationally by the planted XOR assumption, in which case it follows from standard LPN, or statistically by taking n sufficiently small, in which case it follows from subexponential LPN.

Our Construction: Rounding LWE Samples. We follow the same blueprint over $\mathbb{Z}_q$ for an even modulus q , with LWE in place of LPN. The public key consists of a matrix $\mathbf{A} \in \mathbb{Z}_q^{m \times n}$ and a random mask $\mathbf{z} \in \{0, 1\}^m$, and the secret key is again a set of t rows of $\mathbf{A}$ planted to sum to zero, with indicator vector $\mathbf{v} \in \{0, 1\}^m$, so that $\mathbf{v}^\top \mathbf{A} = \mathbf{0}$. Our idea is simply to round LWE samples to bits: a codeword is $\mathbf{c} = \lfloor \mathbf{A}\mathbf{s} + \mathbf{e} \rfloor_2 \oplus \mathbf{z}$ for a random $\mathbf{s} \in \mathbb{Z}_q^n$ and a discrete Gaussian noise vector $\mathbf{e}$, where $\lfloor u \rfloor_2$ is 0 if u is closer to 0 than to $q/2$, and 1 otherwise. The decoder accepts a received word $\mathbf{c}'$ if it passes the same parity check $\langle \mathbf{v}, \mathbf{c}' \oplus \mathbf{z} \rangle = 0$ as in [CG24] (Figure 1). Pseudorandomness follows in the same way, with the hardness of the planted t -SUM problem in place of the planted XOR assumption (Section 3.1), and soundness holds exactly as in [CG24] (Section 3.2).

The main technical step is to show that the decoder still accepts a noisy codeword with probability noticeably above $1/2$. In [CG24], the parity check cancels $\mathbf{G}\mathbf{s}$, so that only the parity of the noise remains. In our case, as the rounding is not linear, the parity check does not cancel $\mathbf{A}\mathbf{s}$: ignoring the channel flips, the decoder accepts exactly when an even number of the t planted entries of $\mathbf{A}\mathbf{s} + \mathbf{e}$ round to 1, which depends on $\mathbf{A}\mathbf{s}$ and not only on the noise. However, the planted relation still ensures that these entries are uniformly random subject to summing to $\langle \mathbf{v}, \mathbf{e} \rangle$, which is biased toward 0. Using Fourier analysis, we show that this bias carries over to the parity of the rounded planted entries (Section 3.3).

2 Preliminaries

2.1 Notation

We write $\mathbf{v}[i]$ for the i -th entry of a vector $\mathbf{v}$, and $\mathbf{A}[i]$ for the i -th row of a matrix $\mathbf{A}$. For $\mathbf{v} \in \{0, 1\}^m$, we write $\text{supp}(\mathbf{v}) = \{i : \mathbf{v}[i] = 1\}$ and $\text{wt}(\mathbf{v}) = |\text{supp}(\mathbf{v})|$. For $s > 0$, $D_{\mathbb{Z}, s}$ denotes the discrete Gaussian distribution over $\mathbb{Z}$ with parameter s , which assigns each $x \in \mathbb{Z}$ probability proportional to $\exp(-\pi x^2/s^2)$. For $p \in [0, 1/2)$, BSC_p denotes the binary symmetric channel, which flips each bit of its input independently with probability p . We write $\Delta(X, Y)$ for the statistical distance between distributions X and Y . For an integer $q \geq 2$, let $\omega_q^a := \exp(2\pi i a/q)$ for $a \in \mathbb{R}$, which depends only on a modulo q , so that ω_q^u is

well defined for $u \in \mathbb{Z}_q$. For even q not divisible by 4, the *rounding* of $u \in \mathbb{Z}_q$ is

$$\lfloor u \rfloor_2 := \begin{cases} 0 & \text{if } |\bar{u}| < q/4, \\ 1 & \text{if } |\bar{u}| > q/4, \end{cases}$$

where $\bar{u} \in (-q/2, q/2]$ is the integer lift of u . That is, $\lfloor u \rfloor_2$ is 0 if u is closer to 0 than to $q/2$ in $\mathbb{Z}_q$, and 1 if it is closer to $q/2$, where ties cannot occur as $q/4$ is not an integer. We apply $\lfloor \cdot \rfloor_2$ to vectors coordinatewise.

2.2 Learning with Errors

Definition 2.1 (LWE [Reg05]). *Given dimension n , number of samples $m = m(n)$, modulus $q = q(n)$, and error rate $\alpha = \alpha(n)$, the $\text{LWE}_{n,m,q,\alpha}$ problem is to distinguish between the following two distributions.*

- $(\mathbf{A}, \mathbf{A}\mathbf{s} + \mathbf{e})$, where $\mathbf{A} \leftarrow \mathbb{Z}_q^{m \times n}$, $\mathbf{s} \leftarrow \mathbb{Z}_q^n$, and $\mathbf{e} \leftarrow D_{\mathbb{Z}, \alpha q}^m$
- $(\mathbf{A}, \mathbf{r})$, where $\mathbf{A} \leftarrow \mathbb{Z}_q^{m \times n}$ and $\mathbf{r} \leftarrow \mathbb{Z}_q^m$

The LWE assumption with constant modulus-to-noise ratio states that, for every constant $\alpha \in (0, 1)$, every $q = q(n) \leq \text{poly}(n)$ with $\alpha q \geq 2\sqrt{n}$, and every $m = m(n) \leq \text{poly}(n)$, any non-uniform probabilistic polynomial-time algorithm solves $\text{LWE}_{n,m,q,\alpha}$ with advantage at most $\text{negl}(n)$.

Definition 2.2 (Subexponential LWE). *The subexponential LWE assumption with constant modulus-to-noise ratio states that, for every $T(n) = 2^{\tilde{O}(\sqrt{n})}$, every constant $\alpha \in (0, 1)$, every $q = q(n) \leq \text{poly}(n)$ with $\alpha q \geq 2\sqrt{n}$, and every $m = m(n) \leq T(n)$, any non-uniform probabilistic $T(n)$ -time algorithm solves $\text{LWE}_{n,m,q,\alpha}$ (Definition 2.1) with advantage at most $1/T(n)$ for large enough n .*

2.3 Planted t -SUM

Definition 2.3 (Planted t -SUM [ASS⁺24]). *Let n, m, t , and q be positive integers with $t \leq m$ and $q \geq 2$. The planted t -SUM distribution $\text{Plant}(n, m, q, t)$ samples a uniformly random $\mathbf{v} \in \{0, 1\}^m$ with $\text{wt}(\mathbf{v}) = t$ and $\mathbf{A} \leftarrow \mathbb{Z}_q^{m \times n}$ uniformly subject to $\sum_{i \in \text{supp}(\mathbf{v})} \mathbf{A}[i] = \mathbf{0}$, and outputs $(\mathbf{A}, \mathbf{v})$. The planted t -SUM problem is to distinguish $\mathbf{A}$ from a uniformly random matrix in $\mathbb{Z}_q^{m \times n}$, where $(\mathbf{A}, \mathbf{v}) \leftarrow \text{Plant}(n, m, q, t)$. For functions $n = n(m)$, $q = q(m)$, and $t = t(m)$, the planted t -SUM problem is hard if any non-uniform probabilistic polynomial-time algorithm solves it with advantage at most $\text{negl}(m)$.*

2.4 Pseudorandom Codes

Definition 2.4 (Weak PRC [CG24, GG25, Chr26]). *Let $p \in [0, 1/2)$. A public-key zero-bit weak pseudorandom code (PRC) of length $N = N(\lambda)$ against BSC_p consists of polynomial-time algorithms (Setup, KeyGen, Encode, Decode).*

- $\text{pp} \leftarrow \text{Setup}(1^\lambda)$: The deterministic setup algorithm takes as input the security parameter λ and outputs public parameters pp . All subsequent algorithms take pp as an implicit input.
- $(\text{pk}, \text{sk}) \leftarrow \text{KeyGen}(\text{pp})$: The key generation algorithm takes as input the public parameters pp and outputs a public key pk and a secret key sk .
- $\text{c} \leftarrow \text{Encode}(\text{pk})$: The encoding algorithm takes as input a public key pk and outputs a codeword $\text{c} \in \{0, 1\}^N$.
- $b \leftarrow \text{Decode}(\text{sk}, \text{c}')$: The decoding algorithm takes as input a secret key sk and a string $\text{c}' \in \{0, 1\}^N$, and outputs $b \in \{1, \perp\}$.

The algorithms must satisfy the following requirements, where $\text{pp} \leftarrow \text{Setup}(1^\lambda)$.

Pseudorandomness: Any non-uniform probabilistic polynomial-time adversary distinguishes (pk, c) from (pk, u) with advantage at most $\text{negl}(\lambda)$, where $(\text{pk}, \text{sk}) \leftarrow \text{KeyGen}(\text{pp})$, $\text{c} \leftarrow \text{Encode}(\text{pk})$, and $\text{u} \leftarrow \{0, 1\}^N$.

Weak Soundness: For every $\text{c}' \in \{0, 1\}^N$,

$$\Pr \left[\text{Decode}(\text{sk}, \text{c}') = 1 \mid (\text{pk}, \text{sk}) \leftarrow \text{KeyGen}(\text{pp}) \right] \leq \frac{1}{2} + \text{negl}(\lambda).$$

Weak Robustness:

$$\Pr \left[\text{Decode}(\text{sk}, \text{c}') = 1 \mid \begin{array}{l} (\text{pk}, \text{sk}) \leftarrow \text{KeyGen}(\text{pp}) \\ \text{c} \leftarrow \text{Encode}(\text{pk}) \\ \text{c}' \leftarrow \text{BSC}_p(\text{c}) \end{array} \right] \geq \frac{1}{2} + \frac{1}{\text{poly}(\lambda)}.$$

3 Pseudorandom Codes from LWE

In this section, we present our weak PRC in Figure 1 and prove its pseudorandomness (Section 3.1), soundness (Section 3.2), and robustness (Section 3.3), which give our main theorems (Section 3.4).

3.1 Pseudorandomness

In this section, we prove pseudorandomness, for Option A from LWE and the planted t -SUM problem (Lemma 3.1), and for Option B from subexponential LWE (Lemma 3.3), using that the planted matrix is then statistically close to uniform (Lemma 3.2).

Lemma 3.1 (Pseudorandomness: Option A). *Under the LWE assumption with constant modulus-to-noise ratio (Definition 2.1), if the planted t -SUM problem (Definition 2.3) is hard for $n = \lfloor \sqrt{m} \rfloor$, $q = 4\lceil 2\sqrt{n} \rceil + 2$, and $t = \lceil \log m \rceil$, then the scheme of Figure 1 with Option A satisfies pseudorandomness (Definition 2.4).*

Weak PRC from LWE

$\text{pp} \leftarrow \text{Setup}(1^\lambda)$:

1. (Option A) Set $n = \lfloor \sqrt{\lambda} \rfloor$.
 (Option B) Set $n = \lfloor (\log \lambda)^2 / (2 \log \log \lambda) \rfloor$.
2. Set $m = \lambda$, $q = 4\lceil 2\sqrt{n} \rceil + 2$,^a and $t = \lceil \log \lambda \rceil$.
3. Output $\text{pp} = (n, m, q, t)$.

$(\text{pk}, \text{sk}) \leftarrow \text{KeyGen}(\text{pp})$:

1. Sample $(\mathbf{A}, \mathbf{v}) \leftarrow \text{Plant}(n, m, q, t)$ (Definition 2.3).
2. Sample a random mask $\mathbf{z} \leftarrow \{0, 1\}^m$.
3. Output $\text{pk} = (\mathbf{A}, \mathbf{z})$ and $\text{sk} = (\mathbf{v}, \mathbf{z})$.

$\mathbf{c} \leftarrow \text{Encode}(\text{pk})$:

1. Parse pk as $(\mathbf{A}, \mathbf{z})$.
2. Sample $\mathbf{s} \leftarrow \mathbb{Z}_q^n$ and $\mathbf{e} \leftarrow D_{\mathbb{Z}, q/4}^m$.
3. Output $\mathbf{c} = \lfloor \mathbf{A}\mathbf{s} + \mathbf{e} \rfloor_2 \oplus \mathbf{z} \in \{0, 1\}^m$.

$b \leftarrow \text{Decode}(\text{sk}, \mathbf{c}')$:

1. Parse sk as $(\mathbf{v}, \mathbf{z})$.
2. Output $b = 1$ if $\langle \mathbf{v}, \mathbf{c}' \oplus \mathbf{z} \rangle = 0 \pmod{2}$, and $\perp$ otherwise.

^a We choose q to be even but not divisible by 4, so that the rounding has no ties (Section 2.1).

Fig. 1. A description of our public-key zero-bit weak PRC with two options, A and B, which differ only in the choice of n in **Setup**. We assume that λ is sufficiently large.

Proof. We show that the real view $(\text{pk}, \mathbf{c})$ of the adversary is computationally indistinguishable from the ideal view $(\text{pk}, \mathbf{u})$, where $\mathbf{u} \leftarrow \{0, 1\}^m$ is uniform. By our hardness assumption on the planted t -SUM problem with $m = \lambda$ and a hybrid argument, we may assume that $\mathbf{A}$ is uniform. Then $(\text{pk}, \mathbf{c})$ and $(\text{pk}, \mathbf{u})$ result from applying the map $(\mathbf{A}, \mathbf{b}) \mapsto ((\mathbf{A}, \mathbf{z}), [\mathbf{b}]_2 \oplus \mathbf{z})$ with $\mathbf{z} \leftarrow \{0, 1\}^m$ to an LWE sample $(\mathbf{A}, \mathbf{A}\mathbf{s} + \mathbf{e})$ and to a uniform sample $(\mathbf{A}, \mathbf{r})$, respectively, where $[\mathbf{r}]_2$ is uniform as exactly $q/2$ elements of $\mathbb{Z}_q$ round to each bit. Hence they are indistinguishable by the LWE assumption. $\square$

Lemma 3.2 (Statistical Hiding). *Let n, m, t , and q be positive integers with $t \leq m$ and $q \geq 2$. Let $\mathcal{D}_0$ be the uniform distribution over $\mathbb{Z}_q^{m \times n}$, and let $\mathcal{D}_1$ be the distribution of $\mathbf{A}$ for $(\mathbf{A}, \mathbf{v}) \leftarrow \text{Plant}(n, m, q, t)$. Then*

$$\Delta(\mathcal{D}_0, \mathcal{D}_1) \leq \frac{1}{2} \sqrt{\frac{q^n - 1}{\binom{m}{t}}}.$$

Proof. Let $\mathbf{A} \leftarrow \mathcal{D}_0$, and let Z be the number of $\mathbf{v} \in \{0, 1\}^m$ with $\text{wt}(\mathbf{v}) = t$ and $\sum_{i \in \text{supp}(\mathbf{v})} \mathbf{A}[i] = \mathbf{0}$. To compute the mean and variance of Z , let $E_{\mathbf{v}}$ be the event that $\mathbf{v}$ is counted by Z .

For every $\mathbf{v}$ and every $i \in \text{supp}(\mathbf{v})$, conditioned on the rows other than $\mathbf{A}[i]$, the event $E_{\mathbf{v}}$ holds for exactly one value of $\mathbf{A}[i]$. Hence $\Pr[E_{\mathbf{v}}] = q^{-n}$, and $\mathbb{E}[Z] = \binom{m}{t} q^{-n}$. For distinct $\mathbf{v}, \mathbf{w} \in \{0, 1\}^m$ with $\text{wt}(\mathbf{v}) = \text{wt}(\mathbf{w}) = t$, take $i \in \text{supp}(\mathbf{v}) \setminus \text{supp}(\mathbf{w})$ and $j \in \text{supp}(\mathbf{w}) \setminus \text{supp}(\mathbf{v})$, which exist as $\text{wt}(\mathbf{v}) = \text{wt}(\mathbf{w})$. Conditioned on all other rows, $E_{\mathbf{v}}$ holds for exactly one value of $\mathbf{A}[i]$ and $E_{\mathbf{w}}$ for exactly one value of $\mathbf{A}[j]$, so $\Pr[E_{\mathbf{v}} \wedge E_{\mathbf{w}}] = q^{-2n}$. Hence the events $E_{\mathbf{v}}$ are pairwise independent, and $\text{Var}[Z] = \binom{m}{t} q^{-n} \cdot (1 - q^{-n})$. Given $\mathbf{v}$, the distribution $\text{Plant}(n, m, q, t)$ samples $\mathbf{A}$ uniformly from the $q^{(m-1)n}$ matrices with $\sum_{i \in \text{supp}(\mathbf{v})} \mathbf{A}[i] = \mathbf{0}$. Hence

$$\mathcal{D}_1(\mathbf{A}) = \binom{m}{t}^{-1} \cdot \frac{Z}{q^{(m-1)n}} = \frac{Z}{\mathbb{E}[Z]} \cdot \mathcal{D}_0(\mathbf{A}).$$

Using $\mathbb{E}[|X|] \leq \sqrt{\mathbb{E}[X^2]}$ for $X = Z/\mathbb{E}[Z] - 1$,

$$\Delta(\mathcal{D}_0, \mathcal{D}_1) = \frac{1}{2} \cdot \mathbb{E}\left[\left|\frac{Z}{\mathbb{E}[Z]} - 1\right|\right] \leq \frac{1}{2} \sqrt{\frac{\text{Var}[Z]}{\mathbb{E}[Z]^2}} = \frac{1}{2} \sqrt{\frac{q^n - 1}{\binom{m}{t}}}.$$

$\square$

Lemma 3.3 (Pseudorandomness: Option B). *Under the subexponential LWE assumption with constant modulus-to-noise ratio (Definition 2.2), the scheme of Figure 1 with Option B satisfies pseudorandomness (Definition 2.4).*

Proof. As in the proof of Lemma 3.1, we show that the real view $(\text{pk}, \mathbf{c})$ is computationally indistinguishable from the ideal view $(\text{pk}, \mathbf{u})$. Let $\mathcal{D}_0$ and $\mathcal{D}_1$ be the distributions of Lemma 3.2 for the parameters of Option B. As $\binom{m}{t} \geq$

$(m/t)^t$ with $m = \lambda$ and $t = \lceil \log \lambda \rceil$, we have $\log \binom{m}{t} \geq (1 - o(1))(\log \lambda)^2$. As $\log q = \frac{1}{2} \log n + O(1) \leq \log \log \lambda + O(1)$ and $n \leq (\log \lambda)^2 / (2 \log \log \lambda)$, we have $n \log q \leq (1/2 + o(1))(\log \lambda)^2$. Hence $\Delta(\mathcal{D}_0, \mathcal{D}_1) \leq 2^{-\Omega((\log \lambda)^2)}$ by Lemma 3.2, so by the data-processing inequality and a hybrid argument, we may assume that $\mathbf{A}$ is uniform. Then, as in the proof of Lemma 3.1, the two views result from applying the same map to an LWE sample and to a uniform sample. As $\sqrt{n} \log n = \Theta(\log \lambda \cdot \sqrt{\log \log \lambda})$, the function $T(n) := 2^{\sqrt{n} \log n}$, which is $2^{\tilde{O}(\sqrt{n})}$, satisfies $T(n) = \lambda^{\omega(1)}$. Hence the map followed by the adversary, which runs in time $\text{poly}(\lambda)$, is an LWE distinguisher that runs in time at most $T(n)$ on $m = \lambda \leq T(n)$ samples, so its advantage is at most $1/T(n) = \lambda^{-\omega(1)}$ by the subexponential LWE assumption. $\square$

3.2 Soundness

Lemma 3.4 (Weak Soundness). *The scheme of Figure 1 with Option A and with Option B satisfies weak soundness (Definition 2.4).*

Proof. Fix $\mathbf{c}' \in \{0, 1\}^m$ and condition on $(\mathbf{A}, \mathbf{v})$. As some entry of $\mathbf{v}$ is 1 and $\mathbf{z} \leftarrow \{0, 1\}^m$ is independent of $(\mathbf{A}, \mathbf{v})$, the parity $\langle \mathbf{v}, \mathbf{z} \rangle$ is a uniform bit. As $\langle \mathbf{v}, \mathbf{c}' \oplus \mathbf{z} \rangle = \langle \mathbf{v}, \mathbf{c}' \rangle + \langle \mathbf{v}, \mathbf{z} \rangle \pmod{2}$ and $\langle \mathbf{v}, \mathbf{c}' \rangle$ is fixed, $\text{Decode}(\text{sk}, \mathbf{c}') = 1$ with probability exactly $1/2$, which gives weak soundness. $\square$

3.3 Robustness

In this section, we prove weak robustness (Lemma 3.7) using the Fourier coefficients of $f(u) := (-1)^{\lfloor u \rfloor_2}$ for $u \in \mathbb{Z}_q$ and of the discrete Gaussian $D_{\mathbb{Z}, q/4}$. As $f(u) = \text{sgn}(\cos(2\pi u/q))$, where sgn denotes the sign, f is a discrete square wave, and its Fourier coefficients are discrete analogues of those of the classical square wave. For completeness, we state the facts that we use in Lemmas 3.5 and 3.6, with short proofs.

Lemma 3.5 (Fourier Coefficients of Rounding). *Let q be even and not divisible by 4, let $f(u) := (-1)^{\lfloor u \rfloor_2}$ for $u \in \mathbb{Z}_q$, and let $\hat{f}(k) := \mathbb{E}_{u \leftarrow \mathbb{Z}_q}[f(u) \cdot \omega_q^{-ku}]$ for $k \in \mathbb{Z}_q$. Then the following hold.*

- (a) $\hat{f}(k) \in \mathbb{R}$ and $\hat{f}(-k) = \hat{f}(k)$ for every $k \in \mathbb{Z}_q$.
- (b) $\hat{f}(k) = 0$ for every even $k \in \mathbb{Z}_q$.
- (c) $\hat{f}(k) = \frac{2}{q} \cdot \frac{(-1)^{(\bar{k}-1)/2}}{\sin(\pi \bar{k}/q)}$ for every odd $k \in \mathbb{Z}_q$, where $\bar{k} \in (-q/2, q/2]$ is the integer lift of k .
- (d) $|\hat{f}(k)| \leq 1/|\bar{k}|$ for every odd $k \in \mathbb{Z}_q$.
- (e) $\hat{f}(1) \geq 1/2$.
- (f) $\sum_{k \in \mathbb{Z}_q \setminus \{1, -1\}} |\hat{f}(k)|^t \leq 9 \cdot (1/3)^t$ for every integer $t \geq 2$.

Proof. By definition, $f(u) = 1$ if $|\bar{u}| < q/4$, and $f(u) = -1$ otherwise.

- (a) This follows from $f(-u) = f(u)$.

- (b) This follows from $f(u + q/2) = -f(u)$ and $\omega_q^{-k(u+q/2)} = \omega_q^{-ku}$ for even k .
(c) As $f(u+q/2) \cdot \omega_q^{-k(u+q/2)} = f(u) \cdot \omega_q^{-ku}$ for odd k , we have $\hat{f}(k) = \frac{2}{q} \sum_{|\bar{u}| < q/4} \omega_q^{-ku}$.
As the integers $\bar{u}$ with $|\bar{u}| < q/4$ are $q/2$ consecutive integers centered at 0, the formula for the Dirichlet kernel gives

$$\sum_{|\bar{u}| < q/4} \omega_q^{-ku} = \frac{\sin(\pi \bar{k}/2)}{\sin(\pi \bar{k}/q)} = \frac{(-1)^{(|\bar{k}|-1)/2}}{\sin(\pi |\bar{k}|/q)},$$

where the last equality holds as $\bar{k}$ is odd.

- (d) As $\sin x \geq 2x/\pi$ for $0 \leq x \leq \pi/2$, (c) gives $|\hat{f}(k)| \leq \frac{2}{q} \cdot \frac{q}{2|\bar{k}|} = \frac{1}{|\bar{k}|}$.
(e) As $\sin x \leq x$ for $x \geq 0$, (c) gives $\hat{f}(1) = \frac{2}{q} \cdot \frac{1}{\sin(\pi/q)} \geq \frac{2}{\pi} \geq \frac{1}{2}$.
(f) Every $k \notin \{1, -1\}$ is even or has $|\bar{k}| \geq 3$, so $|\hat{f}(k)| \leq 1/3$ by (b) and (d). As $\sum_{k \in \mathbb{Z}_q} \hat{f}(k)^2 = \mathbb{E}_{u \leftarrow \mathbb{Z}_q} [f(u)^2] = 1$ by Parseval's identity,

$$\sum_{k \notin \{1, -1\}} |\hat{f}(k)|^t \leq \left(\frac{1}{3}\right)^{t-2} \cdot \sum_{k \in \mathbb{Z}_q} \hat{f}(k)^2 = 9 \cdot \left(\frac{1}{3}\right)^t.$$

□

Lemma 3.6 (Fourier Coefficients of Discrete Gaussians). *Let $q \geq 2$ be an integer, let $\alpha > 0$, let $\chi := D_{\mathbb{Z}, \alpha q}$, and let $\hat{\chi}(k) := \mathbb{E}_{e \leftarrow \chi} [\omega_q^{-ke}]$ for $k \in \mathbb{Z}_q$. Then the following hold.*

- (a) $0 \leq \hat{\chi}(k) \leq 1$ and $\hat{\chi}(-k) = \hat{\chi}(k)$ for every $k \in \mathbb{Z}_q$.
(b) If $\alpha = 1/4$ and $q \geq 8$, then $\hat{\chi}(1) \geq 3/4$.

Proof. Let $s := \alpha q$. By the Poisson summation formula, for every integer k ,

$$\hat{\chi}(k) = \frac{\sum_{j \in \mathbb{Z}} \exp(-\pi s^2 (k/q + j)^2)}{\sum_{j \in \mathbb{Z}} \exp(-\pi s^2 j^2)}.$$

- (a) The right-hand side is positive, is at most 1 as $\hat{\chi}(k)$ is an expectation of complex numbers of absolute value 1, and is unchanged under $k \mapsto -k$ by substituting $-j$ for j .
(b) For $k = 1$, the numerator is at least its term for $j = 0$, which is $\exp(-\pi/16)$. As $s = q/4 \geq 2$ and $j^2 \geq j$, the denominator is at most $1 + 2 \sum_{j \geq 1} \exp(-4\pi j) \leq 1 + 3 \exp(-4\pi)$. Hence $\hat{\chi}(1) \geq \exp(-\pi/16) / (1 + 3 \exp(-4\pi)) \geq 3/4$.

□

Lemma 3.7 (Weak Robustness). *For every constant $p \in [0, 1/2)$, the scheme of Figure 1 with Option A and with Option B satisfies weak robustness against BSC_p (Definition 2.4).*

Proof. Let $\text{pp} \leftarrow \text{Setup}(1^\lambda)$, $(\text{pk}, \text{sk}) \leftarrow \text{KeyGen}(\text{pp})$, $\mathbf{c} \leftarrow \text{Encode}(\text{pk})$, and $\mathbf{c}' \leftarrow \text{BSC}_p(\mathbf{c})$. We show that $\text{Decode}(\text{sk}, \mathbf{c}') = 1$ with probability at least $\frac{1}{2} + \frac{1}{\text{poly}(\lambda)}$, which gives weak robustness against BSC_p . Let $\mathbf{s}$ and $\mathbf{e}$ be the samples of Encode , and let $\mathbf{b} := \mathbf{A}\mathbf{s} + \mathbf{e}$, so that $\mathbf{c} = \lfloor \mathbf{b} \rfloor_2 \oplus \mathbf{z}$. Write $\mathbf{c}' = \mathbf{c} \oplus \boldsymbol{\xi}$, where the entries

of $\xi \in \{0, 1\}^m$ are independent of each other and of everything else, and each is 1 with probability p . We may assume that $\gcd(\mathbf{s}[1], \dots, \mathbf{s}[n], q) = 1$. Indeed, q has at most $\log q$ prime divisors, and all entries of $\mathbf{s}$ are divisible by a given prime divisor r with probability $r^{-n} \leq 2^{-n}$, so this fails with probability at most $\log q \cdot 2^{-n} = \text{negl}(\lambda)$ as $n = \omega(\log \lambda)$ by Setup. Then the homomorphism $\mathbf{a} \mapsto \langle \mathbf{a}, \mathbf{s} \rangle$ from $\mathbb{Z}_q^n$ to $\mathbb{Z}_q$ is onto, as its image contains 1 by Bézout's identity, so $\langle \mathbf{a}, \mathbf{s} \rangle$ is uniform over $\mathbb{Z}_q$ for uniform $\mathbf{a} \in \mathbb{Z}_q^n$. Condition on $\mathbf{v}$ and on such an $\mathbf{s}$, and let $S_{\mathbf{v}} := \text{supp}(\mathbf{v})$.

As $\mathbf{c}' \oplus \mathbf{z} = \lfloor \mathbf{b} \rfloor_2 \oplus \xi$, the decoder outputs 1 exactly when $\langle \mathbf{v}, \lfloor \mathbf{b} \rfloor_2 \rangle + \langle \mathbf{v}, \xi \rangle = 0 \pmod{2}$. As $(-1)^{\langle \mathbf{v}, \lfloor \mathbf{b} \rfloor_2 \rangle} = \prod_{i \in S_{\mathbf{v}}} f(\mathbf{b}[i])$ for $f(u) = (-1)^{\lfloor u \rfloor_2}$ as in Lemma 3.5, the decoder thus outputs 1 with probability $\frac{1}{2} + \frac{1}{2} \cdot \rho$, where

$$\rho := \mathbb{E} \left[(-1)^{\langle \mathbf{v}, \xi \rangle} \cdot \prod_{i \in S_{\mathbf{v}}} f(\mathbf{b}[i]) \right]$$

over $\mathbf{A}$, $\mathbf{e}$, and ξ . Since $t = \lceil \log \lambda \rceil$ by Setup, we have $2^{-O(t)} = 1/\text{poly}(\lambda)$, so it suffices to show that $\rho \geq 2^{-O(t)}$. The flips ξ are independent of everything else, and $\text{wt}(\mathbf{v}) = t$, so $\mathbb{E}[(-1)^{\langle \mathbf{v}, \xi \rangle}] = (1 - 2p)^t$, and hence

$$\rho = (1 - 2p)^t \cdot \mathbb{E} \left[\prod_{i \in S_{\mathbf{v}}} f(\mathbf{b}[i]) \right].$$

We next express ρ in terms of the Fourier coefficients of f and of the discrete Gaussian $D_{\mathbb{Z}_q/4}$. Fix $i^* \in S_{\mathbf{v}}$. By Definition 2.3, the rows $\mathbf{A}[i]$ for $i \in S_{\mathbf{v}} \setminus \{i^*\}$ are independent and uniform over $\mathbb{Z}_q^n$, so the entries $\mathbf{b}[i] = \langle \mathbf{A}[i], \mathbf{s} \rangle + \mathbf{e}[i]$ for $i \in S_{\mathbf{v}} \setminus \{i^*\}$ are independent and uniform over $\mathbb{Z}_q$, and independent of $\mathbf{e}$. By the planted relation, we have $\mathbf{b}[i^*] = \langle \mathbf{v}, \mathbf{e} \rangle - \sum_{i \in S_{\mathbf{v}} \setminus \{i^*\}} \mathbf{b}[i]$. For $\hat{\chi}$ as in Lemma 3.6 with $\alpha = 1/4$, we thus have

$$\begin{aligned} \mathbb{E} \left[\prod_{i \in S_{\mathbf{v}}} f(\mathbf{b}[i]) \right] &= \sum_{k \in \mathbb{Z}_q} \hat{f}(k) \cdot \mathbb{E} \left[\omega_q^{k \mathbf{b}[i^*]} \cdot \prod_{i \in S_{\mathbf{v}} \setminus \{i^*\}} f(\mathbf{b}[i]) \right] \\ &= \sum_{k \in \mathbb{Z}_q} \hat{f}(k) \cdot \mathbb{E} \left[\omega_q^{k \langle \mathbf{v}, \mathbf{e} \rangle} \right] \cdot \prod_{i \in S_{\mathbf{v}} \setminus \{i^*\}} \mathbb{E} [f(\mathbf{b}[i]) \cdot \omega_q^{-k \mathbf{b}[i]}] \\ &= \sum_{k \in \mathbb{Z}_q} \hat{f}(k)^t \cdot \hat{\chi}(k)^t. \end{aligned}$$

The first equality follows by expanding $f(\mathbf{b}[i^*]) = \sum_{k \in \mathbb{Z}_q} \hat{f}(k) \cdot \omega_q^{k \mathbf{b}[i^*]}$, and the second by the planted relation and independence. The last holds as $\mathbb{E}[\omega_q^{k \langle \mathbf{v}, \mathbf{e} \rangle}] = \hat{\chi}(-k)^t = \hat{\chi}(k)^t$ by Lemma 3.6(a), since the $\mathbf{e}[i]$ for $i \in S_{\mathbf{v}}$ are independent samples from $D_{\mathbb{Z}_q/4}$, and as $\mathbb{E}[f(\mathbf{b}[i]) \cdot \omega_q^{-k \mathbf{b}[i]}] = \hat{f}(k)$ for $i \in S_{\mathbf{v}} \setminus \{i^*\}$, since $\mathbf{b}[i]$ is uniform. Hence

$$\rho = (1 - 2p)^t \cdot \sum_{k \in \mathbb{Z}_q} \hat{f}(k)^t \cdot \hat{\chi}(k)^t.$$

We now bound this sum. As $q \geq 8$, the frequencies $k = \pm 1$ are distinct, and $\hat{\chi}(-1) = \hat{\chi}(1) \geq 3/4$ by Lemma 3.6(a) and (b). As $\hat{f}(-1) = \hat{f}(1) \geq 1/2$ by Lemma 3.5(a) and (e),

$$\sum_{k \in \{1, -1\}} \hat{f}(k)^t \cdot \hat{\chi}(k)^t = 2\hat{f}(1)^t \cdot \hat{\chi}(1)^t \geq 2 \cdot \left(\frac{1}{2}\right)^t \cdot \left(\frac{3}{4}\right)^t = 2 \cdot \left(\frac{3}{8}\right)^t.$$

For $t \geq 2$, as $|\hat{\chi}(k)| \leq 1$ by Lemma 3.6(a), Lemma 3.5(f) gives

$$\left| \sum_{k \notin \{1, -1\}} \hat{f}(k)^t \cdot \hat{\chi}(k)^t \right| \leq \sum_{k \notin \{1, -1\}} |\hat{f}(k)|^t \leq 9 \cdot \left(\frac{1}{3}\right)^t.$$

Hence $\rho \geq (1 - 2p)^t \cdot (2 \cdot (3/8)^t - 9 \cdot (1/3)^t) = 2^{-O(t)}$. $\square$

3.4 Main Theorems and Consequences

Theorem 3.8 (Weak PRC from LWE and Planted t -SUM). *Under the LWE assumption with constant modulus-to-noise ratio (Definition 2.1), if the planted t -SUM problem (Definition 2.3) is hard for $n = \lfloor \sqrt{m} \rfloor$, $q = 4\lceil 2\sqrt{n} \rceil + 2$, and $t = \lceil \log m \rceil$, then for every constant $p \in [0, 1/2)$, there is a public-key zero-bit weak PRC (Definition 2.4) against BSC_p .*

Proof. The scheme of Figure 1 with Option A is such a PRC for every constant $p \in [0, 1/2)$. It satisfies pseudorandomness by Lemma 3.1, weak soundness by Lemma 3.4, and weak robustness against BSC_p by Lemma 3.7. Note that $D_{\mathbb{Z}, q/4}$ can be sampled in $\text{poly}(\lambda)$ time up to statistical distance $\text{negl}(\lambda)$ [GPV08]. $\square$

Remark 3.9 (Noisy Planted Relations). Analogously to [GG25], the planted relation $\mathbf{v}^\top \mathbf{A} = \mathbf{0}$ can be relaxed to $\mathbf{v}^\top \mathbf{A} = \mathbf{r}^\top$ for a short residual $\mathbf{r} \leftarrow D_{\mathbb{Z}, \sigma}^n$, which yields a weaker assumption. The planted entries of $\mathbf{A}\mathbf{s}$ then sum to $\langle \mathbf{r}, \mathbf{s} \rangle$ instead of 0. To keep this residual small, Encode samples a short secret $\mathbf{s} \leftarrow D_{\mathbb{Z}, \tau}^n$, as in the normal form of LWE [ACPS09], and $\mathbf{e} \leftarrow D_{\mathbb{Z}, \tau}^m$. The proof of Lemma 3.7 goes through with an extra factor $\mathbb{E}[\omega_q^{k \langle \mathbf{r}, \mathbf{s} \rangle}] \in [0, 1]$ in the term for each k , which is bounded below by a constant for $k = \pm 1$ if $\sigma\tau\sqrt{n} = O(q)$. As $\tau = \alpha q \geq 2\sqrt{n}$ as in Definition 2.1, this requires LWE with a polynomial, rather than constant, modulus-to-noise ratio.

Theorem 3.10 (Weak PRC from Subexponential LWE). *Under the subexponential LWE assumption with constant modulus-to-noise ratio (Definition 2.2), for every constant $p \in [0, 1/2)$, there is a public-key zero-bit weak PRC (Definition 2.4) against BSC_p .*

Proof. The proof is the same as that of Theorem 3.8, with Option B and Lemma 3.3 in place of Option A and Lemma 3.1. $\square$

Remark 3.11 (Consequences). By the transformation sketched in [Chr26, Sec. 5], a variant of [GG25, Lem. 19] with independent keys, the weak PRCs of Theorems 3.8 and 3.10 can be amplified to public-key zero-bit PRCs against BSC_p , in which Decode outputs 1 on any fixed string with probability at most $\text{negl}(\lambda)$ and on a noisy codeword with probability at least $1 - \text{negl}(\lambda)$. The transformation of [Chr26, Thm. 2], which builds on [CG24], further gives public-key zero-bit PRCs robust to oblivious substitutions at rate p . By [CG24], these give public-key multi-bit PRCs of constant rate robust to oblivious substitutions, and watermarking and steganography schemes. The former in turn give public-key multi-bit PRCs robust to random deletions and substitutions [CG24] and public-key multi-bit PRCs over polynomial-size alphabets robust to a constant rate of edits [GM24]. All these PRCs except those from [GM24] are binary.

References

- ACPS09. Benny Applebaum, David Cash, Chris Peikert, and Amit Sahai. Fast cryptographic primitives and circular-secure encryption based on hard learning problems. In Shai Halevi, editor, *CRYPTO 2009*, volume 5677 of *LNCS*, pages 595–618. Springer, Berlin, Heidelberg, August 2009. doi:[10.1007/978-3-642-03356-8_35](https://doi.org/10.1007/978-3-642-03356-8_35). 11
- ASS⁺24. Shweta Agrawal, Sagnik Saha, Nikolaj I. Schwartzbach, Akhil Vanukuri, and Prashant Nalini Vasudevan. k-SUM in the sparse regime: Complexity and applications. In Leonid Reyzin and Douglas Stebila, editors, *CRYPTO 2024, Part II*, volume 14921 of *LNCS*, pages 315–351. Springer, Cham, August 2024. doi:[10.1007/978-3-031-68379-4_10](https://doi.org/10.1007/978-3-031-68379-4_10). 1, 2, 4
- BFKL94. Avrim Blum, Merrick L. Furst, Michael J. Kearns, and Richard J. Lipton. Cryptographic primitives based on hard learning problems. In Douglas R. Stinson, editor, *CRYPTO ’93*, volume 773 of *LNCS*, pages 278–291. Springer, Berlin, Heidelberg, August 1994. doi:[10.1007/3-540-48329-2_24](https://doi.org/10.1007/3-540-48329-2_24). 1
- CG24. Miranda Christ and Sam Gunn. Pseudorandom error-correcting codes. In Leonid Reyzin and Douglas Stebila, editors, *CRYPTO 2024, Part VI*, volume 14925 of *LNCS*, pages 325–347. Springer, Cham, August 2024. doi:[10.1007/978-3-031-68391-6_10](https://doi.org/10.1007/978-3-031-68391-6_10). 1, 2, 3, 4, 12
- CGG⁺26. Miranda Christ, Noah Golowich, Sam Gunn, Ankur Moitra, and Daniel Wichs. Improved pseudorandom codes from permuted puzzles. In *Proceedings of the 58th Annual ACM Symposium on Theory of Computing, STOC 2026*, pages 2130–2139. ACM, 2026. doi:[10.1145/3798129.3800916](https://doi.org/10.1145/3798129.3800916). 2
- Chr26. Miranda Christ. A survey of pseudorandom codes. Cryptology ePrint Archive, Report 2026/2162, 2026. URL: <https://eprint.iacr.org/2026/2162>. 2, 4, 12
- DFV26. Victor Dýseryn, Danilo Francati, and Daniele Venturi. Public-key pseudorandom codes from distorted McEliece assumptions. Cryptology ePrint Archive, Report 2026/1184, 2026. URL: <https://eprint.iacr.org/2026/1184>. 1, 2
- DJK⁺26. Nico Döttling, Antoine Joux, Venkata Koppula, Mahesh Sreekumar Rajasree, and Hendrik Waldner. Chosen ciphertext secure pseudorandom codes in the standard model. In Nadia Heninger and Mike Rosulek, editors, *Advances in Cryptology - CRYPTO 2026 - 46th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 17-20, 2026, Proceedings*,

- Part I*, Lecture Notes in Computer Science, pages 282–313. Springer, 2026. doi:10.1007/978-3-032-35367-2_10. 1, 2
- DMR25. Nico Döttling, Anne Müller, and Mahesh Sreekumar Rajasree. Separating pseudorandom codes from local oracles. In Benny Applebaum and Huijia (Rachel) Lin, editors, *TCC 2025, Part IV*, volume 16271 of *LNCS*, pages 225–257. Springer, Cham, December 2025. doi:10.1007/978-3-032-12290-2_8. 2
- GG25. Surendra Ghentiyala and Venkatesan Guruswami. New constructions of pseudorandom codes. In Alina Ene and Eshan Chattopadhyay, editors, *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques, APPROX/RANDOM 2025*, volume 353 of *LIPICs*, pages 54:1–54:22. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2025. doi:10.4230/LIPICs.APPROX/RANDOM.2025.54. 1, 2, 4, 11, 12
- GGW25. Sanjam Garg, Sam Gunn, and Mingyuan Wang. Black-box crypto is useless for pseudorandom codes. In Benny Applebaum and Huijia (Rachel) Lin, editors, *TCC 2025, Part IV*, volume 16271 of *LNCS*, pages 205–224. Springer, Cham, December 2025. doi:10.1007/978-3-032-12290-2_7. 2
- GM24. Noah Golowich and Ankur Moitra. Edit distance robust watermarks via indexing pseudorandom codes. In Amir Globersons, Lester Mackey, Danielle Belgrave, Angela Fan, Ulrich Paquet, Jakub M. Tomczak, and Cheng Zhang, editors, *Advances in Neural Information Processing Systems 37: Annual Conference on Neural Information Processing Systems 2024, NeurIPS 2024, Vancouver, BC, Canada, December 10 - 15, 2024*, 2024. URL: http://papers.nips.cc/paper_files/paper/2024/hash/24c53bfa5b53fc2cf05644f5a7a26bb0-Abstract-Conference.html. 2, 12
- GPV08. Craig Gentry, Chris Peikert, and Vinod Vaikuntanathan. Trapdoors for hard lattices and new cryptographic constructions. In Richard E. Ladner and Cynthia Dwork, editors, *40th ACM STOC*, pages 197–206. ACM Press, May 2008. doi:10.1145/1374376.1374407. 11
- GZS25. Sam Gunn, Xuandong Zhao, and Dawn Song. An undetectable watermark for generative image models. In *The Thirteenth International Conference on Learning Representations, ICLR 2025, Singapore, April 24-28, 2025*. OpenReview.net, 2025. 1
- Reg05. Oded Regev. On lattices, learning with errors, random linear codes, and cryptography. In Harold N. Gabow and Ronald Fagin, editors, *37th ACM STOC*, pages 84–93. ACM Press, May 2005. doi:10.1145/1060590.1060603. 2, 4