Public-Key Encryption from Continuous LWE

Andrej Bogdanov[*] Miguel Cueto Noval [†] Charlotte Hoffmann [‡] Alon Rosen [§]

January 25, 2022

Abstract

The continuous learning with errors (CLWE) problem was recently introduced by Bruna et al. (STOC 2021). They showed that its hardness implies infeasibility of learning Gaussian mixture models, while its tractability implies efficient Discrete Gaussian Sampling and thus asymptotic improvements in worst-case lattice algorithms. No reduction between CLWE and LWE is currently known, in either direction.

We propose four public-key encryption schemes based on the hardness of CLWE, with varying tradeoffs between decryption and security errors, and different discretization techniques. Some of our schemes are based on hCLWE, a homogeneous variant, which is no easier than CLWE. Our schemes yield a polynomial-time algorithm for solving hCLWE, and hence also CLWE, using a Statistical Zero-Knowledge oracle.

1 Introduction

A sample from the *continuous learning with errors* (CLWE) distribution [BRST21] is of the form $(\mathbf{a}, z)$, where $\mathbf{a} \in \mathbb{R}^n$ is a vector with individual entries sampled independently from the standard normal distribution $\mathcal{N}(0, 1)$, and

$$z := \gamma\langle \mathbf{a}, \mathbf{w}\rangle + e \mod 1.$$

Here e is the noise drawn from a Gaussian distribution with mean 0 and variance β^2 for some $\beta > 0$, $\gamma > 0$ is a fixed parameter and $\mathbf{w} \in \mathbb{R}^n$ is a secret unit vector. CLWE is the problem of distinguishing multiple CLWE samples from an equal number of samples of the form $(\mathbf{a}, u)$, where u is uniform over $[0, 1)$ and independent of $\mathbf{a}$.

The CLWE problem can be viewed as a continuous analog of Regev's LWE problem [Reg05] and is at least as (quantumly) hard as the same worst-case lattice problems underlying LWE [BRST21]. However, in spite of the similarities to LWE, and even more so to the equivalent torus LWE [Reg05] and scale-invariant LWE [CS15, AD97], there is no known reduction from one problem to the other.

Owing to its continuous nature, CLWE has been speculated to be a potentially easier target for cryptanalysis than LWE [BRST21]. It is however possible that CLWE turns out to be no easier than LWE. In this work we construct four public-key encryption schemes that are at least as hard to break as CLWE, and in fact potentially harder problems.

[*] Chinese University of Hong Kong. E-mail: andrejb@cse.cuhk.edu.hk. Part of this work done at the Simons Institute for the Theory of Computing and while visiting Bocconi University.

[†] IST Austria. E-mail: miguel.cuetonoval@ist.ac.at Part of this work done while visiting Bocconi University.

[‡] IST Austria. E-mail: charlotte.hoffmann@ist.ac.at Part of this work done while visiting Bocconi University.

[§] Bocconi University and Reichman University. E-mail: alon.rosen@unibocconi.it.

1.1 The Homogeneous CLWE Problem

Our schemes are based on the hardness of the hCLWE problem, which is a homogenenous variant of CLWE, and on a related problem, $(0, 1/2)$-hCLWE. Both problems can be shown to be no easier than CLWE. The hCLWE problem is a special case of learning high-dimensional Gaussian mixtures, a notoriously challenging problem in computational learning theory [DKS17].

Samples of hCLWE are normally distributed in every direction perpendicular to a secret direction $\mathbf{w} \in \mathbb{R}^n$. The distribution in direction $\mathbf{w}$ is a noisy discrete Gaussian, i.e. a mixture of "Gaussian pancakes" of standard deviation $\beta/\sqrt{\beta^2+\gamma^2} \approx \beta/\gamma$ and spacing $\gamma/(\beta^2+\gamma^2) \approx 1/\gamma$. The hCLWE problem is to distinguish miltiple hCLWE samples from purely normal ones.

Bruna et al. [BRST21] show that hCLWE samples are CLWE samples conditioned on $z = 0$ (Figure 1.a) and design a polynomial-time reduction from CLWE to hCLWE based on this property. Conditioning z to take some other fixed value $s \in [0, 1)$ shifts the modes of the distribution in the hidden direction $\mathbf{w}$ by a relative phase of s (Figure 1.b). To control the decryption error two of our schemes construct public keys from a labeled mixture of the two, which we call $(0, 1/2)$-hCLWE (Figure 1.c with red and blue denoting labels 0 and 1/2, respectively).

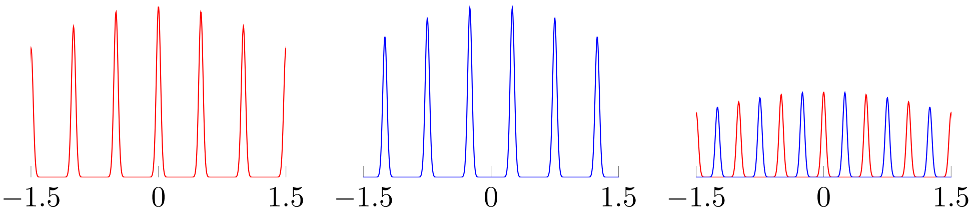

Figure 1: Probability density function of the hidden direction in the (a) hCLWE, (b) 1/2-hCLWE, and (c) $(0, 1/2)$-hCLWE distributions with parameters $\beta = 0.05$ and $\gamma = 2$

Our hCLWE-based public-key encryption schemes imply limits on the hardness of hCLWE: just as LWE, hCLWE is tractable in Statistical Zero-Knowledge. It follows that hCLWE is unlikely to be helpful for constructing encryption as secure as NP (unless NP is contained in coAM).

1.2 Four Public-Key Encryption Schemes

The schemes we propose offer varying tradeoffs between decryption and security errors, and use different techniques when disretizing continuous values.

Our first scheme ("pancake") is based on hCLWE. It has inverse polynomial decryption and constant security errors. These parameters, along with the specifics of the scheme, already suffice to prove that hCLWE can be solved in Statistical Zero-Knowledge (SZK), and therefore is in coAM.[1] The discretization step in the scheme can be performed during encryption, and so the public key is continuous. Arguing security then necessitates proving an analog of the leftover hash lemma for Gaussian matrices, which may be of independent interest.

One could in principle rely on standard techniques to reduce decryption and security errors in the first scheme [HR05] , albeit at the price of a significant loss in efficiency. Instead, we present three different ideas to reduce the errors directly.

[1] We will say that a distinguishing problem is in class $\mathcal{C}$ if there is an algorithm in $\mathcal{C}$ that accepts at least 2/3 of the yes instances and rejects at least 2/3 of the no instances.

Scheme	Assumption	Decryption error	Security error	PK size	SK size
Pancake	hCLWE	$O(1/n)$	$1/4$	$\tilde{O}(n^3)$	n
Bimodal	$(0, 1/2)$-hCLWE	0	$1/2$	$\tilde{O}(n^3)$	n
Discretized	$(0, 1/2)$-hCLWE	0	2^{-n+2}	$O(n^3)$	n
Baguette	hCLWE(ℓ)	$O(1/n^{\ell})$	$1/4$	$\tilde{O}(n^3)$	$n\ell$

Table 1: Comparison of our encryption schemes. If the assumption holds against time $t(n) + n^{O(1)}$ and advantage $\Omega(\epsilon(n))$ adversaries then the corresponding scheme is resilient against time $t(n)$ and advantage (security error + $\epsilon(n)$) adversaries.

In the second scheme ("bimodal"), we achieve perfect decryption error by publishing $(0, 1/2)$-hCLWE samples as the public key. To encrypt a 0, Bob uses samples with $z = 0$ and to encrypt a 1, he uses samples with $z = 1/2$. This eliminates the probability that a random normal ciphertext of 1 is of the form of an hCLWE sample and thus makes decryption perfect.

The third scheme ("discretized") achieves negligible security error by mapping the samples into a parallelpiped spanned by hCLWE samples; a technique due to Ajtai and Dwork [AD97]. Here the discretization step takes place already in public-key generation, allowing for the use of the standard leftover hash lemma and yielding favorable security error in comparison with the other schemes.

In the fourth scheme ("baguette") we achieve negligible decryption error assuming only hCLWE. Instead of publishing samples that have a "pancake" distribution in one direction, we sample vectors that have a pancake distribution in ℓ hidden directions. In [BRST21] the authors give a reduction from hCLWE to this hCLWE(ℓ) distribution.

The parallelepiped technique can also be applied to the fourth scheme, yielding an hCLWE-based scheme with negligible decryption and security error. We omit a formal analysis of this step as it is similar to the discretized scheme.

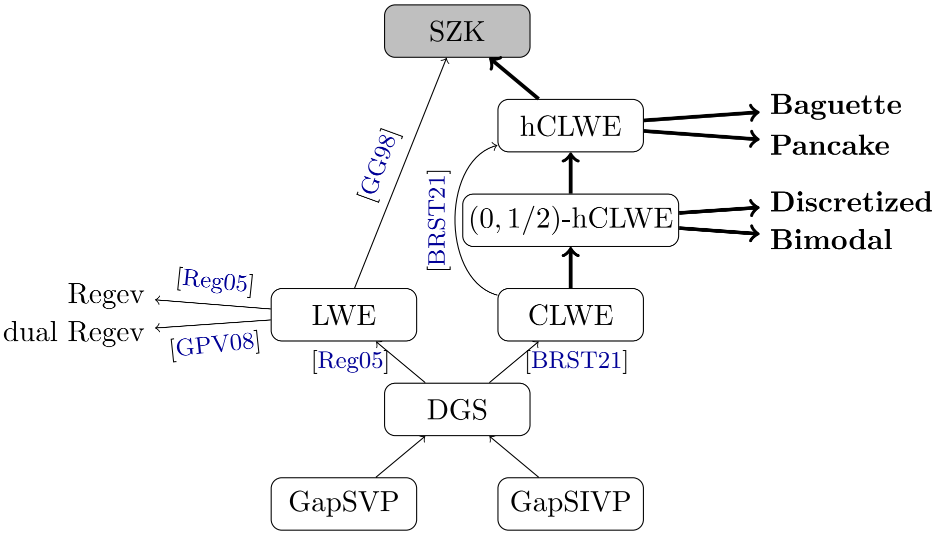

Figure 2: Reductions between problems and encryption schemes (new results are in bold).

1.3 CLWE, SZK, and Statistical-Computational Gaps

Hypothesis testing is the computational task of distinguishing whether a sequence of independent samples $X_1, \ldots, X_m$ comes from a null distribution D_0 or an alternative distribution D_1. Several works [BR13, HWX15, BB20] uncover that such problems tend to exhibit *statistical-computational gaps*: There is a range of sample complexities $m \in [m_{\text{stat}}, m_{\text{comp}}]$ for which hypothesis testing is possible, but no efficient[2] algorithm is known.

A striking feature of the hCLWE problem is that it is potentially intractable even when the sample complexity is unbounded, i.e., m_{comp} is infinite. Our Theorem 9.2 shows that when $m \geq \tilde{O}(n^2)$ samples are available hCLWE becomes solvable in SZK. Thus, in a world in which SZK = BPP, the computational threshold m_{comp} for hCLWE is at most $\tilde{O}(n^2)$.

In contrast, the statistical threshold for CLWE is $m_{\text{stat}} = O(n)$. It is an intriguing open question whether a statistical-computational gap for hCLWE exists assuming SZK = BPP. One approach for ruling out this possibility is to design a more efficient hCLWE-based PKE scheme.

Applying the reduction from CLWE to hCLWE of Bruna et al., our result also implies that CLWE is in SZK. As their reduction does not preserve sample complexity, the resulting SZK algorithm for CLWE requires a larger number of samples.

2 Technical Overview

The messages in our encryption schemes are single bits. The distributions of encryptions of zero and one, respectively, are efficiently distinguishable with the secret key but not without it. The public keys are independent samples of the hCLWE or $(0, 1/2)$-hCLWE distributions and the secret key is the hidden direction $\mathbf{w}$ of the corresponding yes instances.

As can be seen in Figure 1, the hCLWE samples used to generate the public-key have a periodic discrete structure along the secret direction $\mathbf{w}$. Encryption is designed to retain this discrete structure in the ciphertext even though the sender is oblivious to it. Decryption calculates the correlation between the secret key $\mathbf{w}$ and the ciphertext. This correlation is close to an integer multiple of the period for encryptions of zero and (typically) far from it for encryptions of one.

2.1 ”Pancake” Encryption

The first scheme (Section 4) is based on the hCLWE problem. The secret key is a random unit vector $\mathbf{w}$ and the public key is an $n \times m$ matrix $\mathbf{A}$ that consists of m hCLWE samples conditioned on the secret direction $\mathbf{w}$.

To encrypt a 0, sample a uniform vector $\mathbf{t} \leftarrow \{1/\sqrt{m}, -1/\sqrt{m}\}^m$ and compute $\mathbf{At}$. To encrypt a 1, sample a standard normal vector. The ciphertext $\mathbf{c}$ is a discretization of the resulting vector using a rounding function that divides the real line into intervals (”buckets”) of equal Gaussian measure.[3] To decrypt a ciphertext $\mathbf{c}$, compute $\gamma\sqrt{m}\langle \mathbf{w}, \mathbf{c} \rangle$ and output 0 if the result is close to an integer. Otherwise output 1.

The scheme has inverse polynomial decryption error since the probability of $\gamma\sqrt{m}\langle \mathbf{w}, \mathbf{c} \rangle$ being close to an integer is inverse polynomial for a random choice of $\mathbf{c}$. The main technical contribution in this scheme is the security proof, in particular Proposition 4.5. This result is an analog of the

[2] Efficiency is measured in terms of the length of a single sample $|X_1|$, not the number of samples m.

[3] In the body of the paper we use the notation $1/\gamma' = \gamma/(\beta^2 + \gamma^2)$ for the period of the hCLWE hidden direction. As the difference between $1/\gamma'$ and $1/\gamma$ is small we make no distinction between the two in this overview.

leftover hash lemma for the multiplication of Gaussian matrices with vectors with uniform vectors $\mathbf{t} \leftarrow \{1/\sqrt{m}, -1/\sqrt{m}\}^m$ which shows that the security error is 1/2 for our choice of parameters.

2.2 "Bimodal" Encryption

In the second scheme (Section 6) we introduce the following changes: We base the scheme on the $(0, 1/2)$-hCLWE problem and publish two matrices $(\mathbf{A}_0, \mathbf{A}_1)$ as the public key. The matrix $\mathbf{A}_0$ consists of hCLWE samples conditioned on $\mathbf{w}$ and $\mathbf{A}_1$ consists of 1/2-hCLWE samples conditioned on $\mathbf{w}$. To encrypt a 0, do the same as in the pancake scheme with the matrix $\mathbf{A}_0$. To encrypt a 1, do exactly the same with $\mathbf{A}_1$. To decrypt, check if $\gamma\sqrt{m}\langle \mathbf{w}, \mathbf{c}\rangle \mod 1$ is closer to 0 or to 1/2. Replacing one hCLWE matrix by two $(0, 1/2)$-hCLWE matrices yields perfect decryption error for all but negligibly many choices of the public key. The security error however remains constant.

2.3 "Discretized" Encryption

The third scheme (Section 7) has perfect decryption for all but an inverse polynomial fraction of public keys and negligible security error. To achieve this we make use of the parallelepiped technique due to Ataj and Dwork [AD97] to obtain uniform matrices from $(0, 1/2)$-hCLWE samples.

The public key $(\mathbf{A}_0, \mathbf{A}_1, \mathbf{B})$ now consists of 3 matrices: A square matrix $\mathbf{B}$ that consists of hCLWE samples, a matrix $\mathbf{A}_0$ that is essentially obtained by mapping hCLWE samples into the parallelpiped spanned by the columns of $\mathbf{B}$ (denoted by $\mathcal{P}(\mathbf{B})$) and a matrix $\mathbf{A}_1$ that is obtained in the same way but with 1/2-hCLWE samples mapped to $\mathcal{P}(\mathbf{B})$. This mapping into the parallelepiped transforms Gaussian vectors in $\mathbb{R}$ into uniform vectors in $\mathcal{P}(\mathbf{B})$, while preserving the pancakes in the secret direction. An additional rounding step discretizes the matrices $\mathbf{A}_0, \mathbf{A}_1$.

To encrypt a bit b, sample a vector $\mathbf{t}$ with uniform entries in $\{-1, 1\}$ and set $\mathbf{c} := \mathbf{A}_b\mathbf{t} \mod \mathbf{B}$. To decrypt, check if $\gamma\langle \mathbf{w}, \mathbf{c}\rangle \mod 1$ is closer to 0 or to 1/2. For all but an inverse polynomial fraction of choices of the matrix $\mathbf{B}$ this scheme has perfect correctness. Security follows from the classical leftover hash lemma [IZ89] since the matrices $\mathbf{A}_0$ and $\mathbf{A}_1$ are uniform and discrete.

2.4 "Baguette" Encryption

The fourth scheme (Section 8) is based on the hCLWE(ℓ) problem, which is potentially harder than the $(0, 1/2)$-hCLWE problem. We achieve negligible decryption error by modifying our first scheme as follows: Instead of publishing samples that have a pancake distribution in only one hidden direction, we publish a matrix $\mathbf{A}$ of samples that have a pancake distribution in $\log n$ many hidden directions, i.e. we replace the Gaussian pancakes with "Gaussian Baguettes".

As in the first scheme, to encrypt a 0, sample a uniform vector $\mathbf{t} \leftarrow \{1/\sqrt{m}, -1/\sqrt{m}\}^m$ and compute $\mathbf{At}$, and to encrypt a 1, sample a standard normal vector. Discretization of the ciphertext is also identical to the first scheme.

To decrypt, multiply the ciphertext with a matrix that consists of all hidden directions. If all of the entries in the resulting vectors are close to an integer, output 0, otherwise output 1. While the probability that the inner product of the ciphertext of 1 with one secret direction is close to an integer is polynomial, the probability that this happens for all of the $\log n$ directions is negligible. The security error of this scheme remains constant but could be amplified either by a standard approach or by the above parallelepiped method.

3 Preliminaries

We introduce key concepts that are used throughout the paper.

3.1 Public Key Encryption

We focus on encryption schemes with binary message space. Some of our encryptions schemes will decrypt incorrectly with bounded probability δ, and will sometimes also have noticeable (but still bounded) statistical distance ε between the distribution of encryptions of zero and those of one. Once such schemes are attained it is possible to invoke standard polarization methods to amplify security and correctness errors to be negligible [HR05].

Definition 3.1 (Syntax)**.** *A* public key encryption scheme *is a tuple of algorithms* $(\mathrm{Gen}, \mathrm{Enc}, \mathrm{Dec})$ *such that for* $\lambda \in \mathbb{N}$*:*

- $\mathrm{Gen}(1^\lambda)$ *outputs a pair of keys* (sk, pk)*;*
- $\mathrm{Enc}(pk, m)$ *encrypts a message* m *with the public key* pk *and outputs a ciphertext* c*;*
- $\mathrm{Dec}(sk, c)$ *decrypts a ciphertext* c *using the secret key* sk *and outputs a message* m*.*

Both key-generation, Gen, and encryption, Enc, are randomized. We will allow for the decryption algorithm, Dec, to make errors.

Definition 3.2 (δ-correctness)**.** *A public key encryption scheme* $(\mathrm{Gen}, \mathrm{Enc}, \mathrm{Dec})$ *is* correct with probability δ *if*

$$\Pr\left[\mathrm{Dec}(sk, \mathrm{Enc}(pk, m)) = m\right] \geq \delta,$$

where probability is taken over the randomness of Gen *and* Enc*. We call* $1-\delta$ *the* decryption error*.*

Security is defined through indistingushability of encryptions [GM84]. To this end, we rely on the notion computational indistinguishability (defined next), which is also used more generally in our proofs of security.

Definition 3.3 (ε-indistinguishability)**.** *We say that two distributions* X, Y *are* ε-indistinguishable *if for any probabilistic polynomial time algorithm ("distinguisher")* A*:*

$$\left| \Pr_{x \leftarrow X}[A(x) = 1] - \Pr_{y \leftarrow Y}[A(y) = 1] \right| \leq \varepsilon.$$

Sometimes we quantify over size s distinguishers, in which case we say that the distributions X, Y are (s, ε)-indistinguishable. By (∞, ε)-indistinguishable we mean that the distributions X, Y have statistical distance ε.

Definition 3.4 (ε-security)**.** *A public key encryption scheme* $(\mathrm{Gen}, \mathrm{Enc}, \mathrm{Dec})$ *is said to have* security error $\varepsilon \in [0, 1]$ *if the distributions* $(pk, \mathrm{Enc}(pk, 0))$ *and* $(pk, \mathrm{Enc}(pk, 1))$ *are* ε*-indistinguishable, where probabilities are taken over the randomness of* Gen *and* Enc*.*

3.2 Singular values and matrix norms

We use the notation $\langle \mathbf{x}, \mathbf{y} \rangle = \sum_{i=1}^n x_i y_i$ for the inner product in $\mathbb{R}^n$ and $\|\mathbf{x}\| = \sqrt{\langle \mathbf{x}, \mathbf{x} \rangle}$ for the Euclidean norm in $\mathbb{R}^n$.

Given a norm $\|\cdot\|_p$ on $\mathbb{R}^n$ and a norm $\|\cdot\|_q$ on $\mathbb{R}^m$, the operator norm $\|\cdot\|_{p,q}$ on the space of matrices $\mathbb{R}^{m\times n}$ is defined as

$$\|\mathbf{A}\|_{p,q} \coloneqq \sup_{0 \neq x \in \mathbb{R}^n} \frac{\|\mathbf{A}\mathbf{x}\|_q}{\|\mathbf{x}\|_p}$$

where $\mathbf{A} \in \mathbb{R}^{m\times n}$.

We are mainly interested in the case when the vector norms on $\mathbb{R}^n$ and $\mathbb{R}^m$ are just the Euclidean norm. In this case, we use the notation $\|\mathbf{A}\|$.

Fact 3.5. *For all vectors $\mathbf{x}, \mathbf{y} \in \mathbb{R}^n$ and matrices $\mathbf{A} \in \mathbb{R}^{m\times n}$, $\mathbf{B} \in \mathbb{R}^{n\times m}$ we have*

1. $\langle \mathbf{x}, \mathbf{y} \rangle \leq \|\mathbf{x}\| \cdot \|\mathbf{y}\|$,
2. $\|\mathbf{A}\mathbf{x}\| \leq \|\mathbf{A}\| \cdot \|\mathbf{x}\|$,
3. $\|\mathbf{A}\mathbf{B}\| \leq \|\mathbf{A}\| \cdot \|\mathbf{B}\|$.

Given a matrix $\mathbf{A} \in \mathbb{R}^{m\times n}$ the singular values of $\mathbf{A}$ are the square roots of the eigenvalues of $\mathbf{A}^T\mathbf{A}$. We use the notation $s_i(\mathbf{A})$ for the i-th singular value and we order them in descending order, that is, $s_1(\mathbf{A})$ denotes the largest singular value of $\mathbf{A}$.

Fact 3.6. *Let $\mathbf{A} \in \mathbb{R}^{n\times m}$ and $s_1(\mathbf{A})$ its largest singular value. We have that*

$$s_1(\mathbf{A}) = \|\mathbf{A}\| \leq \sqrt{\sum_{i\in[m]} \sum_{j\in[n]} |a_{ij}|^2}.$$

Fact 3.7 ([Ede88])**.** *Let $\mathbf{B} \in \mathbb{R}^{n\times n}$ be a matrix with entries independently sampled from $\mathcal{N}(0, \sigma^2)$ and $s_n(B)$ be its smallest singular value. We have that for every $\varepsilon > 0$*

$$\Pr[s_n(\mathbf{B}) \leq \varepsilon] \leq \sigma^{-1}\sqrt{n}\varepsilon.$$

3.3 Normal Distribution

We consider both the continuous normal distribution and the discrete one. We refer to them as normal or Gaussian and we use these two words interchangeably. The continuous normal distribution in $\mathbb{R}^n$, denoted by $\mathcal{N}_n(\boldsymbol{\mu}, \boldsymbol{\Sigma})$, has probability density function at $\mathbf{x} \in \mathbb{R}^n$ given by

$$\frac{1}{\sqrt{(2\pi)^n \det \boldsymbol{\Sigma}}} \exp\left(-\frac{1}{2}(\mathbf{x} - \boldsymbol{\mu})^T \boldsymbol{\Sigma}^{-1} (\mathbf{x} - \boldsymbol{\mu})\right)$$

where $\boldsymbol{\mu} \in \mathbb{R}^n$ is the mean vector and $\boldsymbol{\Sigma} \in \mathbb{R}^{n\times n}$ is a positive definite matrix called the covariance matrix. We usually work with vectors with i.i.d. entries from $\mathcal{N}(0,1)$, which we denote by $\mathcal{N}_n(0,1)$ instead of $\mathcal{N}_n(0, \mathbf{I}_n)$, where $\mathbf{I}_n$ is the n-dimensional identity matrix.

Let $\boldsymbol{\mu} \in \mathbb{R}^{n \times s}$, $\mathbf{U} \in \mathbb{R}^{n \times n}$ and $\mathbf{V} \in \mathbb{R}^{s \times s}$ such that both $\mathbf{U}$ and $\mathbf{V}$ are positive definite. We say that a $n \times s$ random matrix $\mathbf{A}$ has a matrix normal distribution, denoted by $\mathcal{MN}_{n \times s}(\boldsymbol{\mu}, \mathbf{U}, \mathbf{V})$, if and only if $\mathrm{vec}(\mathbf{A})$ follows a $\mathcal{N}_{ns}(\mathrm{vec}(\boldsymbol{\mu}), \mathbf{V} \otimes \mathbf{U})$ distribution .

In the case of the discrete Gaussian we only consider covariance matrices of the form $\boldsymbol{\Sigma} = \sigma^2 \mathbf{I}_n$ for some $\sigma > 0$ and $\boldsymbol{\mu} = \mathbf{0}$. This allows us to simplify the notation for the discrete Gaussian distribution, $\mathcal{D}_{L,\sigma^2}$, where L denotes its support. If $\mathbf{x} \in L$, the probability mass function at $\mathbf{x}$ is proportional to the value of the probability density function at $\mathbf{x} \in \mathbb{R}^n$ of $\mathcal{N}(\mathbf{0}, \sigma^2 \mathbf{I}_n)$.

Fact 3.8. $\Pr[\mathcal{N}(0,1) > t] \leq \frac{1}{\sqrt{2\pi}t} e^{-t^2/2}$ *for* $t > 0$.

Fact 3.9. $\Pr[\mathcal{N}(0,1) > t] \geq e^{-t^2}$ *for* $t \geq 1.91$.

Proof. Fact 3.8 and the following lower bound are well-known [Gor41]:

$$\Pr[\mathcal{N}(0,1) > t] \geq \frac{t}{\sqrt{2\pi}(t^2+1)} e^{-t^2/2}$$

The inequality $t/\sqrt{2\pi}(t^2+1) \geq e^{-t^2/2}$ for $t \geq 1.91$ gives Fact 3.9. □

Fact 3.10. *For a random variable* X *with distribution* $\mathcal{N}(0,\sigma^2)$ *it holds that:*

$$\Pr\left[|X| > s\right] \leq \sqrt{2\sigma^2/\pi} \frac{e^{-s^2/(2\sigma^2)}}{s}.$$

Corollary 3.11. *For a vector* $x \in \mathbb{R}^n$ *with entries independently sampled from* $\mathcal{N}(0,\sigma^2)$ *we have*

$$\|x\| \leq n\sigma$$

with probability at at least $1 - \sqrt{n}e^{-n}$.

Proof. By Fact 3.10 we have that the absolute value of a fixed entry of x is larger than $\sqrt{n}\sigma$ with probability at most $e^{-n}/(\sqrt{n})$. Applying the union bound yields that all entries are bounded by $\sqrt{n}\sigma$ with probability $ne^{-n}/(\sqrt{n}) = \sqrt{n}e^{-n}$. It follows that $\|x\| \leq \sqrt{n \cdot (\sqrt{n}\sigma)^2} = n\sigma$ with the same probability. □

Fact 3.12. *Let* X *be a random variable with* $\Pr[|X| > t] \leq 2e^{-t^2/(2\sigma^2)}$ *then we have*

$$\mathbb{E}[|X|^k] \leq (2\sigma^2)^{k/2} k\Gamma(k/2).$$

Proof. We have

$$\mathbb{E}[|X|^k] = \int_0^\infty \Pr[|X| > t^{1/k}] dt \leq 2 \int_0^\infty e^{-t^{2/k}/2\sigma^2} dt = (2\sigma^2)^{k/2} k\Gamma(k/2),$$

where the last equality follows from replacing t with $u = t^{2/k}/(2\sigma^2)$. □

3.4 Gaussian hypercontractivity

In the proof of our analog of the leftover hash lemma we will use a Gaussian hypercontractivity result. For the sake of completeness we introduce some concepts that are only needed to understand the general hypercontractivity theorem and the proof of Corollary 3.16. Later on we will only use the result in Corollary 3.16.

Definition 3.13. *Let $L^k(\mathbb{R}^n, \gamma)$ denote the space of Borel functions $f : \mathbb{R}^n \to \mathbb{R}$ that have finite k-th moment $\|f\|_k^k$ under the Gaussian measure, i.e. $\|f\|_k^k = \mathbb{E}_{z \sim \mathcal{N}(0,1)^n}[|f(z)|^k]$ is finite.*

Definition 3.14. *Let $X = (X_1, \ldots, X_n), X' = (X'_1, \ldots, X'_n)$ be two n-dimensional standard Gaussian variables. We call X and X' ρ-correlated if each pair (X_i, X'_i) is a correlated Gaussian pair with covariance matrix $\mathbb{E}[X_i^2] = \mathbb{E}[X'^2_i] = 1$, $\mathbb{E}[X_i X'_i] = \rho$ and the n pairs are mutually independent.*

Theorem 3.15. *[O'D14, Gaussian hypercontractivity Theorem, p. 333] Let $f, g \in L^1(\mathbb{R}^n, \gamma)$, $r, s \geq 0$, $0 \leq \rho \leq \sqrt{rs} \leq 1$ and Z, Z' be ρ-correlated n-dimensional Gaussian variables. Then we have that*

$$\mathbb{E}_{(Z,Z')}[f(Z)g(Z')] \leq \|f\|_{1+r}\|g\|_{1+s}.$$

Corollary 3.16. *Let S be any event in $\mathbb{R}^n$ and $X = (X_1, \ldots, X_n), X' = (X'_1, \ldots, X'_n)$ be ρ-correlated n-dimensional standard variables. We have that*

$$\Pr[X \in S \text{ and } X' \in S] \leq \Pr[X \in S]^{1/(1+|\rho|)} \Pr[X' \in S]^{1/(1+|\rho|)}.$$

Proof. First assume that $\rho \geq 0$. Theorem 3.15 with $r = s = \rho$ and f, g being indicators of the set S gives the statement:

$$\begin{aligned}
\Pr[X \in S \text{ and } X' \in S] &= \mathbb{E}[\mathbb{1}_S(X)\mathbb{1}_S(X')] \\
&\leq \|\mathbb{1}_S\|_{1+\rho}^2 \\
&= \mathbb{E}_{Z \sim \mathcal{N}(0,1)}[\mathbb{1}_S(Z)^{1+\rho}]^{2/(1+\rho)} \\
&= \Pr[X \in S]^{1/(1+\rho)} \Pr[X' \in S]^{1/(1+\rho)}.
\end{aligned}$$

In the case where $\rho < 0$ we apply the statement to X and $-X'$ since then X and $-X$ are $-\rho$-correlated and hence

$$\Pr[X \in S \text{ and } X' \in S] = \Pr[X \in S \text{ and } -X' \in S] \leq \Pr[X \in S]^{1/(1-\rho)} \Pr[X' \in S]^{1/(1-\rho)}.$$

The two cases together give the claim of the corollary. □

3.5 Lattices

Given a basis $\mathbf{B} = \{\mathbf{b}_1, \ldots, \mathbf{b}_n\}$ of $\mathbb{R}^n$ we define the lattice $L(\mathbf{B})$ as the set of all integer linear combinations of $\mathbf{B}$, i.e.,

$$L(\mathbf{B}) := \left\{ \sum_{i=1} z_i \mathbf{b}_i \;\middle|\; z_1, \ldots, z_n \in \mathbb{Z} \right\}.$$

The minimum distance of a lattice L is $\lambda_1(L) := \min_{0 \neq \mathbf{x} \in L} \|\mathbf{x}\|$. We say that $L \subset \mathbb{R}^n$ is a (full-rank) lattice if there exists a basis $\mathbf{B}$ of $\mathbb{R}^n$ such that $L = L(\mathbf{B})$. The dual lattice of L is the set

$$L^* \coloneqq \{\mathbf{y} \in \mathbb{R}^n \mid \forall \mathbf{x} \in L \colon \langle \mathbf{y}, \mathbf{x} \rangle \in \mathbb{Z}\}\,.$$

The minimum distance of a lattice L is $\lambda_1(L) \coloneqq \min_{0 \neq \mathbf{x} \in L} \|\mathbf{x}\|$.

One parameter that is particularly useful to study the relation between normal distributions and lattices is the smoothing parameter. It is defined as

$$\eta_\epsilon(L) \coloneqq \inf \{r \mid f_{X_r}(L^* \setminus \{\mathbf{0}\}) \leq \epsilon\}$$

where X_r is a random variable with distribution $\mathcal{N}_n(\mathbf{0}_n, \frac{r^2}{2\pi}\mathbf{I}_n)$ and f_{X_r} its probability density function.

The following is a result that guarantees that under certain conditions the sum of two independent random variables with a discrete normal distribution are statistically close to a random variable with discrete normal distribution. This result appears as Lemma 4.12 in the complete version ([BF10]) of [BF11]. Here we state it in a simplified way and make the bound on the statistical distance explicit.

Lemma 3.17 (Special case of [BF10, Lemma 4.12])**.** *Let $L \subset \mathbb{Z}^n$ be a full rank lattice, $\epsilon \in \mathbb{R}$, $\sigma \in \mathbb{R}$, and X_1, X_2 two independent random variables with distribution $\mathcal{D}_{L+\mathbf{t},\sigma^2}$. If $\sigma > \eta_\epsilon(L)$, then the statistical distance between a random variable with distribution $\mathcal{D}_{L+2\mathbf{t},2\sigma^2}$ and $X_1 + X_2$ is at most $\frac{2\epsilon}{1-\epsilon}$.*

We will also need the following results:

Lemma 3.18 ([PRSD17, Lemma 2.5])**.** *For any n-dimensional lattice L, real $\epsilon > 0$, and $r \geq \eta_\epsilon(L)$, the statistical distance between $\mathcal{N}(0, \frac{r^2}{2\pi}\mathbf{I}_n) \mod L$ and the uniform distribution over $\mathbb{R}^n/L$ is at most $\epsilon/2$.*

Lemma 3.19 ([PRSD17, Lemma 2.6])**.** *Let $L \subset \mathbb{R}^n$ be an n-dimensional lattice, $c \geq 1$ and $\epsilon = \exp(-c^2 n)$. It holds that $\eta_\epsilon(L) \leq \frac{c\sqrt{n}}{\lambda_1(L^*)}$.*

And as a special case;

Lemma 3.20. *For any n-dimensional lattice $L \subset \mathbb{R}^n$ with basis $\mathbf{B} = \{\mathbf{b}_1, \ldots, \mathbf{b}_n\}$ we have $\eta_{2^{-n}}(\mathbf{B}) \leq \sqrt{n}\max_i \|\mathbf{b}_i\|$.*

3.6 The (homogeneous) CLWE distribution

Definition 3.21 (CLWE Distribution)**.** *Given a dimension n and parameters $\beta, \gamma > 0$, and a unit vector $\mathbf{w} \in \mathbb{R}^n$, samples $(\mathbf{y}, z) \in \mathbb{R}^n \times [0, 1)$ from the* CLWE *distribution $\mathcal{A}_{\mathbf{w},\beta,\gamma,n}$ are generated as follows:*

1. *Sample $\mathbf{y} \leftarrow \mathcal{N}_n(0, 1)$.*

2. *Sample $e \leftarrow \mathcal{N}(0, \beta^2)$.*

3. *Output $(\mathbf{y}, \gamma\langle \mathbf{w}, \mathbf{y} \rangle + e \mod 1)$.*

Definition 3.22 (CLWE Distinguishing Problem)**.** *For real numbers $\beta, \gamma > 0$ and $n \in \mathbb{N}$, the (average-case) distinguishing problem $\mathrm{CLWE}_{\beta,\gamma,n}$ asks to distinguish between $\mathcal{A}_{\mathbf{w},\beta,\gamma,n}$ for a uniform vector $\mathbf{w} \in \mathbb{R}^n$ and $\mathcal{N}_n(0, 1) \times \mathcal{U}$, where $\mathcal{U}$ is the uniform distribution on $[0, 1)$.*

Definition 3.23 (hCLWE Distribution)**.** *Given a dimension n, parameters $\beta, \gamma > 0$, and a unit vector $\mathbf{w} \in \mathbb{R}^n$, samples $\mathbf{y} \in \mathbb{R}^n$ from the* hCLWE *distribution $\mathcal{H}_{\mathbf{w},\beta,\gamma,n}$ are generated as follows:*

1. *The pancake: Sample $k \in \mathbb{Z}$ with probability proportional to $\exp(-k^2/(2\gamma^2 + 2\beta^2))$.*
2. *The noise: Sample e from $\mathcal{N}(0, \beta'^2)$, where $\beta'^2 = \beta^2/(\gamma^2 + \beta^2)$.*
3. *The rest: Sample $\mathbf{w}^\perp$ as $\mathcal{N}_{n-1}(0, 1)$ on the subspace orthogonal to $\mathbf{w}$.*
4. *Output $\mathbf{w}^\perp + (k/\gamma' + e)\mathbf{w}$, where $1/\gamma' = \gamma/(\gamma^2 + \beta^2)$.*

Definition 3.24 (hCLWE Distinguishing Problem)**.** *For real numbers $\beta, \gamma > 0$ and $n \in \mathbb{N}$, the (average-case) distinguishing problem* $\text{hCLWE}_{\beta,\gamma,n}$ *asks to distinguish between $\mathcal{H}_{\mathbf{w},\beta,\gamma,n}$ for a uniform vector $\mathbf{w} \in \mathbb{R}^n$ and $\mathcal{N}_n(0, 1)$.*

The (s, ε) homogeneous CLWE (hCLWE(s, ε)) assumption [BRST21] postulates that for a random $\mathbf{w}$, a hCLWE oracle cannot be distinguished in size s from an oracle that outputs $\mathcal{N}(0, 1)$ samples on $\mathbb{R}^n$ with advantage ε. As evidence Bruna, Regev, Song, and Tang show a polynomial-time quantum reduction from the problem of sampling a discrete gaussian of width $O(\sqrt{n}/\beta)$ times the smoothing parameter assuming $\gamma \geq 2\sqrt{n}$. Specifically, if γ and β are polynomial in n then it is plausible that hCLWE holds with s and $1/\varepsilon$ exponential in n. Note that they define the standard normal distribution as $\mathcal{N}(0, 1/(2\pi))$ instead of $\mathcal{N}(0, 1)$.

It can be shown that all hCLWE versions with different variances are equivalent by rescaling the samples and the problem parameters γ and β. In particular hCLWE with normal distribution $\mathcal{N}(0, 1/(2\pi))$ and problem parameters γ and β is equivalent to hCLWE with normal distribution $\mathcal{N}(0, 1)$ and problem parameters $\gamma/\sqrt{2\pi}$ and $\beta/\sqrt{2\pi}$. We will always work with the $\mathcal{N}(0, 1)$ distribution for which $\gamma \geq \sqrt{n}$ is sufficient.

4 Scheme 1: Pancake Encryption

The first encryption scheme relies on the hCLWE assumption and has polynomial decryption- and constant security error. It is the basis for all of the following encryption schemes that achieve better error bounds but either rely on an assumption that is potentially easier to break and/or incur a blow-up in the key size. Furthermore, this scheme enables us to prove that hCLWE is in the complexity class SZK. Before presenting the scheme, we define a rounding function that we will need to discretize the ciphertexts of the scheme.

4.1 Rounding into buckets of equal measure

In the encryption we use of the following Gaussian rounding function $\text{round}_r \colon \mathbb{R} \to \{1, \ldots, r\}$ given by

$$\text{round}_r(x) = \lceil r \cdot \mu((-\infty, x)) \rceil,$$

where μ is the standard Gaussian measure on the line. In words, partition $\mathbb{R}$ into r intervals ("buckets") $J_1, J_2, \ldots, J_r$ of equal Gaussian measure, and set $\text{round}_r(x)$ to be the unique i such that $x \in J_i$. We extend the definition over $\mathbb{R}^n$ coordinate-wise, i.e. $\text{round}_r(x_1, \ldots, x_n) = (\text{round}_r(x_1), \ldots, \text{round}_r(x_n))$.

Some of the buckets are very wide (at least two of them are infinite!) so the rounding will cause encryption errors with some probability. We will argue that this is an unlikely event using the following regularity property of round_r. The *width* of an interval $J = (a, b)$ is $b - a$.

Proposition 4.1. *For every $0 < \alpha < 1$ and all r such that $r^{1-\alpha} \geq 19$, the number of i for which the width of $J_i = \mathrm{round}_r^{-1}(i)$ exceeds $r^{-\alpha}$ is at most $2r^\alpha/\sqrt{\ln r^{1-\alpha}} + 2$.*

The k widest intervals capture a k/r fraction of the probability mass μ at the tails of the normal distribution. If t is chosen so that $\mu((-\infty, t) \cup (t, \infty)) = k/r$ then the next widest interval is of the form (t', t) and t' is uniquely determined by the constraint $\mu((t', t)) = 1/r$. Using suitable analytic approximations for the normal CDF the maximum width $t - t'$ of all remaining intervals can be bounded by $r^{-\alpha}$ when $k = \lfloor 2r^\alpha/\sqrt{\ln r^{1-\alpha}} + 2\rfloor$.

Proof of Proposition 4.1. By monotonicity, the width of the intervals increases the farther the interval is from zero. Assuming r is even or not all intervals have width exceeding $r^{-\alpha}$, there are exactly two narrowest intervals of width exceeding $r^{-\alpha}$ of the form $J_{i-} = (-t, -t')$ and $J_{i+} = (t', t)$ for some $0 < t' < t$. We will later justify the assumption. The intervals of width at least $r^{-\alpha}$ are then J_{i-}, J_{i+}, and all those contained in the set $B = (-\infty, t] \cup [t, \infty)$. As $\mu(B) = \sum_{i\colon J_i \subseteq B} \mu(J_i) = \sum_{i\colon J_i \subseteq B} 1/r$, the number of intervals of width exceeding $r^{-\alpha}$ must equal $r \cdot \mu(B) + 2$. By Fact 3.8,

$$\frac{\mu(B)}{2} = \Pr[\mathcal{N}(0,1) > t] \leq \frac{e^{-t^2/2}}{\sqrt{2\pi} t},$$

from where,

$$\frac{1}{r} = \mu(J_{i+}) \geq \frac{e^{-t^2/2}}{\sqrt{2\pi}} \cdot \mathrm{width}(J_{i+}) > \frac{e^{-t^2/2}}{\sqrt{2\pi}} \cdot r^{-\alpha} \geq \frac{t\mu(B)}{2} \cdot r^{-\alpha}. \tag{1}$$

If $t \geq 2$ then by Fact 3.9, $\mu(B)/2 \geq e^{-t^2}$, so $t \geq \sqrt{\ln(2/\mu(B))}$. Plugging into (1) we get $\mu(B)\sqrt{\ln(2/\mu(B))} \leq 2r^{\alpha-1}$, and hence

$$\mu(B) \leq \frac{2r^{\alpha-1}}{\sqrt{\ln(2/\mu(B))}} \leq \frac{2r^{\alpha-1}}{\sqrt{\ln tr^{1-\alpha}}} \leq \frac{2r^{\alpha-1}}{\sqrt{\ln r^{1-\alpha}}}.$$

We conclude that $r\mu(B) + 2 \leq 2r^\alpha/\sqrt{\ln r^{1-\alpha}} + 2$ in this case. If $t < 2$ then by (1) we get $r^{1-\alpha} < \sqrt{2\pi}e^{t^2/2} < 19$.

If r is odd, at least 3, and all intervals including the middle one $J_i = (-t, t)$ have width exceeding $r^{-\alpha}$, then $t < 0.5$ and

$$\frac{1}{r} = \mu((-t, t)) \geq \frac{e^{-t^2/2}}{\sqrt{2\pi}} \cdot r^{-\alpha} > \frac{r^{-\alpha}}{3},$$

so $r^{1-\alpha} < 3$. □

4.2 The encryption scheme

The scheme is parametrized by $\gamma > 0$; $\beta > 0$; $r > 0$ and $n, m \in \mathbb{Z}$.

- The secret key is a uniformly random unit vector $\mathbf{w} \in \mathbb{R}^n$.
- The public key is a matrix $\mathbf{A} \in \mathbb{R}^{n \times m}$ whose columns are independent hCLWE samples from $\mathcal{H}_{\mathbf{w},\beta,\gamma,n}$.
- To encrypt a 0, sample a vector $\mathbf{t} \in \{-1/\sqrt{m}, +1/\sqrt{m}\}^m$ uniformly at random and output

$$\mathbf{c} := \mathrm{round}_r(\mathbf{At}).$$

- To encrypt a 1, sample a vector $\mathbf{c} \leftarrow \{1, 2, \ldots, r\}^n$ uniformly at random and output $\mathbf{c}$.
- To decrypt a ciphertext $\mathbf{c}$, take any $\mathbf{z}$ such that $\mathrm{round}_r(\mathbf{z}) = \mathbf{c}$, compute

  $$\gamma' \sqrt{m} \langle \mathbf{w}, \mathbf{z} \rangle \mod 1$$

  and check if it is in the interval $(-1/2n, 1/2n)$. If yes, output 0, else output 1.

Theorem 4.2. *Set the parameters of the scheme to* $\gamma = \sqrt{n}$,$\beta = (40000 n^{3/2} \log(n))^{-1}$,$r = (40000 n^3 \log(n))^{5/3}$ *and* $m = 10^8 \log(n)^2 n^2$*. Assuming* $\mathrm{hCLWE}(s, \varepsilon)$*, the scheme has decryption error* $O(1/n) + \varepsilon$ *and security error at most* $1/4 + 2\varepsilon$*.*

We prove correctness and security of the scheme separately in the next two subsections.

4.3 Correctness

There are two sources of error in this encryption scheme: key generation error and encryption error. While the key generation error is negligible, the encryption error may be noticeable.

We will call a public key $\mathbf{A}$ *good* if in all its column samples the noise e has magnitude at most $\sqrt{n}\beta$. By $\mathrm{hCLWE}(s, \varepsilon)$, Fact 3.10 and a union bound a public key is good except with probability $m/e^n + \varepsilon$.

Claim 4.3. *Assuming* $\mathrm{hCLWE}(s, \varepsilon)$ *where* s *is the complexity of rounding, the probability that* $\mathrm{Dec}(\mathbf{w}, \mathrm{Enc}(\mathbf{A}, 0)) \neq 0$ *is at most* $1/2n + \varepsilon$ *for all but a fraction of* $m/e^n + \varepsilon$ *choices of* $\mathbf{A}$*.*

Proof. Given a ciphertext $\mathbf{c}$, the decryption chooses a vector $\mathbf{z}$ that satisfies $\mathrm{round}_r(\mathbf{z}) = \mathbf{c}$ and outputs

$$\gamma' \sqrt{m} \langle \mathbf{w}, \mathbf{z} \rangle = \gamma' \sqrt{m} \langle \mathbf{wA}, \mathbf{t} \rangle + \gamma' \sqrt{m} \langle \mathbf{w}, \mathbf{z} - \mathbf{At} \rangle.$$

Since the public key is good, all entries of $\mathbf{wA}$ are $\sqrt{n}\beta$-close to multiples of $1/\gamma'$ (i.e. they are a multiple of $1/\gamma'$ plus an error term of magnitude at most $\sqrt{n}\beta$) , so $\langle \mathbf{wA}, \mathbf{t} \rangle$ must be $\sqrt{mn}\beta$-close to a multiple of $1/\gamma'\sqrt{m}$. By our choice of parameters we get

$$\left| \gamma' \sqrt{m} \langle \mathbf{wA}, \mathbf{t} \rangle \bmod 1 \right| \leq 1/4n.$$

It remains to bound the absolute value of the term $\gamma' \sqrt{m} \langle \mathbf{w}, \mathbf{z} - \mathbf{At} \rangle$, which arises from the rounding error. By the $\mathrm{hCLWE}(s, \varepsilon)$ assumption and the fact that $f_{\mathbf{t}}(\mathbf{X}) := \mathrm{round}_r(\mathbf{Xt})$ is an efficiently computable function, the probability that at least one entry of $\mathrm{round}_r(\mathbf{At})$ falls into an interval of width more than $r^{-3/5}$ is within ε of the probability of the same event when $\mathbf{A}$ is replaced by a standard normal $n \times m$ matrix $\mathbf{N}$. By Proposition 4.1 and a union bound, this probability is at most $2nr^{-2/5} + 2n/r$ which is at most $1/2n$ by our choice of r. Assume this does not happen. Since $\mathrm{round}_r(\mathbf{z}) = \mathrm{round}_r(\mathbf{Nt})$, the entries of $\mathbf{z} - \mathbf{Nt}$ are bounded in magnitude by $r^{-3/5}$, so

$$\left| \gamma' \sqrt{m} \langle \mathbf{w}, \mathbf{z} - \mathbf{Nt} \rangle \right| \leq \|\mathbf{w}\| \cdot \|\mathbf{z} - \mathbf{Nt}\| \, \gamma' \sqrt{m} \leq \sqrt{n} r^{-3/5} \gamma' \sqrt{m} \leq \frac{1}{4n}$$

by our choice of parameters. By the triangle inequality $|\gamma' \sqrt{m} \langle \mathbf{w}, \mathbf{z} \rangle \bmod 1| \leq 1/2n$ as desired. As this happens except with probability at most $1/2n + \varepsilon$, the claim follows. □

Claim 4.4. *The probability that* $\mathrm{Dec}(\mathbf{w}, \mathrm{Enc}(\mathbf{A}, 1)) \neq 1$ *is at most* $3/2n$*.*

Proof. The ciphertext $\mathbf{c} := \mathrm{Enc}(\mathbf{A}, 1)$ is a vector with i.i.d. uniform entries in $\{1, 2, \ldots, r\}$. The decryption chooses a vector $\mathbf{z}$ that satisfies $\mathrm{round}_r(\mathbf{z}) = \mathbf{c}$. By definition of the rounding function this is a standard Gaussian $\mathbf{g} \in \mathbb{R}^n$ plus the rounding error $\mathbf{z} - \mathbf{g}$. We have already seen that the absolute value of $\gamma' \sqrt{m} \langle \mathbf{w}, \mathbf{z} - \mathbf{g} \rangle \bmod 1$ is at most $1/4n$ except with probability $1/2n$.

Since $\mathbf{w}$ is a unit vector, $\langle \mathbf{w}, \mathbf{g} \rangle$ is a standard normal random variable. By the smoothing property of Gaussians modulo the integers (Lemmas 3.18 and 3.19) $\gamma' \sqrt{m} \langle \mathbf{w}, \mathbf{g} \rangle \bmod 1$ is $\exp(-\gamma'^2 m)$-close to a uniform random variable on the interval $(-1/2, 1/2)$. The probability that its absolute value is $1/2n$ or less is at most $1/n$. It follows that the decryption error is at most $1/2n + 1/n = 3/2n$ in this case. □

4.4 Security

We show that the above scheme has constant security error by the following argument:

1. Under the $\mathrm{hCLWE}(s, \varepsilon)$ assumption, the tuple $(\mathbf{A}, \mathrm{Enc}(\mathbf{A}, b))$ is ε-indistinguishable from $(\mathbf{N}, \mathrm{Enc}(\mathbf{N}, b))$ for both $b = 0$ and $b = 1$, where $\mathbf{N}$ is a $n \times m$ matrix with i.i.d. entries sampled from $\mathcal{N}(0, 1)$.
2. The distributions $(\mathbf{N}, \mathrm{Enc}(\mathbf{N}, 0))$ and $(\mathbf{N}, \mathrm{Enc}(\mathbf{N}, 1))$ are 1/4-statistically close.
3. It follows that the distributions $(\mathbf{A}, \mathrm{Enc}(\mathbf{A}, 0))$ and $(\mathbf{A}, \mathrm{Enc}(\mathbf{A}, 1))$ are at most $(1/4 + 2\varepsilon)$-indistinguishable.

The first claim follows directly from the hCLWE assumption using the fact that the encryption is an efficiently computable function of the public-key. To prove the second claim (Proposition 4.8) we will argue that for each possible set (bucket) S that is the of the form $\mathrm{round}_r^{-1}(\mathbf{c})$, the random variable $\Pr[\mathbf{Nt} \in S | \mathbf{N}]$ is unlikely to deviate from its mean $\mathbb{E}\left[\Pr[\mathbf{Nt} \in S | \mathbf{N}]\right] = \Pr[\mathbf{g} \in S]$ by much, where $\mathbf{g}$ is a standard normal vector. Then by a union bound over all the buckets we can say that with high probability over the choice of $\mathbf{N}$ the statistical distance between the two distributions is small (given $\mathbf{N}$). Recall that $\mu(S) = \Pr[\mathbf{g} \in S]$ is the standard Gaussian measure over $\mathbb{R}^n$.

Proposition 4.5. *Let $\mathbf{N}$ be a $m \times n$ matrix of independent $\mathcal{N}(0, 1)$ random variables, $\mathbf{t}$ a random m-dimensional $\{-1/\sqrt{m}, +1/\sqrt{m}\}$ vector, and S be any event in $\mathbb{R}^n$. Assuming $\mu(S) \geq \exp(-\sqrt{m}/4e)$, we have*

$$\mathrm{Var}\left[\Pr[\mathbf{Nt} \in S | \mathbf{N}]\right] \leq \frac{4e\mu(S)^2 \ln(1/\mu(S))}{\sqrt{m}}.$$

Proof. Using the definition $\mathrm{Var}[Z] = \mathbb{E}[Z^2] - \mathbb{E}[Z]^2$ for any random variable Z we get:

$$\mathrm{Var}\big[\Pr[\mathbf{Nt} \in S | \mathbf{N}]\big] = \Pr[\mathbf{Nt} \in S \text{ and } \mathbf{Nt}' \in S] - \Pr[\mathbf{Nt} \in S] \Pr[\mathbf{Nt}' \in S], \tag{2}$$

where $\mathbf{t}, \mathbf{t}'$ are two independent copies of a random $\pm 1/\sqrt{m}$-valued m-dimensional vector. Let $X = (X_1, \ldots, X_n) = \mathbf{Nt}$ and $X = (X'_1, \ldots, X'_n) = \mathbf{Nt}'$. Conditioned on $\mathbf{t}$ and $\mathbf{t}'$, each pair (X_i, X'_i) is a correlated Gaussian pair (independent of the others) with covariance matrix $\mathbb{E}[X_i^2] = \mathbb{E}[X'^2_i] = 1$, $\mathbb{E}[X_i X'_i] = \rho$, where $\rho = \langle \mathbf{t}, \mathbf{t}' \rangle$ is the inner product of the vectors $\mathbf{t}$ and $\mathbf{t}'$. By Corollary 3.16 we get

$$\Pr[\mathbf{Nt} \in S \text{ and } \mathbf{Nt}' \in S] \leq \Pr[\mathbf{Nt} \in S]^{1/(1+|\rho|)}] \Pr[\mathbf{Nt}' \in S]^{1/(1+|\rho|)}$$

for fixed choices of $\mathbf{t}$ and $\mathbf{t}'$. The quantities $\Pr[\mathbf{N}\mathbf{t} \in S]$ and $\Pr[\mathbf{N}\mathbf{t}' \in S]$ are simply the Gaussian measure $\mu(S)$ of the bucket S, so (2) gives

$$\operatorname{Var}\bigl[\Pr[\mathbf{N}\mathbf{t} \in S|\mathbf{N}]\bigr] \le \mathbb{E}[\mu(S)^{2/(1+|\rho|)} - \mu(S)^2] = \mathbb{E}\bigl[\mu(S)^{-2|\rho|/(1+|\rho|)} - 1\bigr]\mu(S)^2. \tag{3}$$

The expectation here is taken over the choice of $\rho = \langle \mathbf{t}, \mathbf{t}' \rangle = (Z_1 + \cdots + Z_m)/m$, where Z_i are i.i.d. ± 1. If we further use $\mu(S) \le 1$ and $|\rho| \ge 0$, we get that

$$\mathbb{E}\bigl[\mu(S)^{-2|\rho|/(1+|\rho|)} - 1\bigr] \le \mathbb{E}[\mu(S)^{-2|\rho|}] - 1.$$

We further bound this expression by using the following claim:

Claim 4.6. $\mathbb{E}[\mu^{-2|\rho|}] \le \sum_{k=0}^{\infty}(es)^k$, *where* $s = (2\ln 1/\mu)/\sqrt{m}$.

By our assumption $\mu(S) \ge \exp(-\sqrt{m}/4e)$, we have $0 \le es \le 1/2$ so we get $\sum_k (es)^k = 1/(1-es) \le 1 + 2es$. Plugging into (3) we get the proposition. □

Proof of Claim 4.6. The random variable $|\rho|\sqrt{m}$ is subgaussian: $\Pr[|\rho|\sqrt{m} \ge t] \le 2\exp(-t^2/2)$, but doesn't have mean zero. Then

$$\begin{aligned}
\mathbb{E}[\mu^{-2|\rho|}] &= \mathbb{E}[\exp(s\,|\rho|\,\sqrt{m})] \\
&= \sum_{k=0}^{\infty} \frac{s^k \mathbb{E}[(|\rho|\,\sqrt{m})^k]}{k!} \\
&\le 1 + \sum_{k=1}^{\infty} \frac{s^k \cdot 2^{k/2} k \Gamma(k/2)}{k!} && \text{(by Fact 3.12)} \\
&\le 1 + \sum_{k=1}^{\infty} \frac{(es)^k}{k^{k/2-1}} && (\Gamma(k/2) \le (k/2)^{k/2} \text{ and } k! \ge (k/e)^k) \\
&\le \sum_{k=0}^{\infty} (es)^k && (k^{k/2-1} \ge 1 \text{ for } k \ge 1.)
\end{aligned}$$
□

Using Proposition 4.5 we can now bound the statistical distance between $(\mathbf{N}, \mathrm{round}_r(\mathbf{N}\mathbf{t}))$ and $(\mathbf{N}, \mathrm{round}_r(\mathbf{g}))$ which are basically encryptions of 0 and 1 with a standard normal matrix instead of a public key. Security of the scheme then follows from the fact that under the hCLWE assumption $\mathbf{N}$ is indistinguishable from a public key.

Corollary 4.7. *Let* round *be any discrete-valued function on* $\mathbb{R}^n$ *such that* $\mu(\mathrm{round}^{-1}(\mathbf{c})) \ge \alpha$ *for all* $\mathbf{c}$ *in the range of* round. *Then the statistical distance between* $(\mathbf{N}, \mathrm{round}(\mathbf{N}\mathbf{t}))$ *and* $(\mathbf{N}, \mathrm{round}(\mathbf{g}))$ *is at most* $\sqrt{4e\ln(1/\alpha)/\sqrt{m}}$.

Proof. We will assume $\alpha \ge \exp(-\sqrt{m}/4e)$ for otherwise $\sqrt{4e\ln(1/\alpha)/\sqrt{m}} \ge 1$ and the claim is true. Fix $\mathbf{c}$ and let $S = \mathrm{round}^{-1}(\mathbf{c})$. Applying the Cauchy-Schwarz inequality to Proposition 4.5 we have

$$\mathbb{E}\bigl|\Pr[\mathbf{N}\mathbf{t} \in S|\mathbf{N}] - \mu(S)\bigr| \le \sqrt{\frac{4e\ln(1/\mu(S))}{\sqrt{m}}} \cdot \mu(S).$$

In particular, if $\mu(\mathrm{round}^{-1}(\mathbf{c})) \geq \alpha \geq \exp(-\sqrt{m}/4e)$ for every $\mathbf{c}$, then

$$\begin{aligned}\Delta((\mathbf{N}, \mathrm{round}(\mathbf{Nt})); (\mathbf{N}, \mathrm{round}(\mathbf{g}))) &= \frac{1}{2}\mathbb{E}\left[\sum_{\mathbf{c}} \left|\Pr[\mathrm{round}(\mathbf{Nt}) = \mathbf{c}|\mathbf{N}] - \Pr[\mathrm{round}(\mathbf{g}) = \mathbf{c}|\mathbf{N}]\right|\right] \\ &\leq \frac{1}{2}\sum_{\mathbf{c}} \sqrt{\frac{4e\ln(1/\mu(\mathrm{round}^{-1}(\mathbf{c})))}{\sqrt{m}}} \cdot \mu(\mathrm{round}^{-1}(\mathbf{c})) \\ &\leq \sqrt{\frac{e\ln(1/\alpha)}{\sqrt{m}}} \sum_{\mathbf{c}} \mu(\mathrm{round}^{-1}(\mathbf{c})),\end{aligned}$$

which is at most the desired expression as the summation equals $\mu(\mathbb{R}^n) = 1$. □

Proposition 4.8. *The distributions* $(\mathbf{N}, \mathrm{Enc}(\mathbf{N}, 0))$ *and* $(\mathbf{N}, \mathrm{Enc}(\mathbf{N}, 1))$ *are* $1/4$*-statistically close for a matrix* $\mathbf{N}$ *of independent standard Gaussians.*

Proof. By construction, $\mu(\mathrm{round}_r^{-1}(b)) = r^{-n}$ for all b. By Corollary 4.7 the statistical distance between encryptions is then at most $\sqrt{4e\ln r^n/\sqrt{m}}$ which is at most $1/4$ by our choice of parameters. □

Corollary 4.9. *Assuming* $\mathrm{hCLWE}(s, \varepsilon)$*,* $(\mathbf{A}, \mathrm{Enc}(\mathbf{A}, 0))$ *and* $(\mathbf{A}, \mathrm{Enc}(\mathbf{A}, 1))$ *are* $(s - \mathrm{poly}(n), 1/4 + 2\varepsilon)$*-indistinguishable where* $\mathbf{A}$ *is the public key matrix.*

Proof. Let $\mathbf{N}$ be a random normal matrix. By $\mathrm{hCLWE}(s, \varepsilon)$, $(\mathbf{A}, \mathrm{Enc}(\mathbf{A}, b))$ and $(\mathbf{N}, \mathrm{Enc}(\mathbf{N}, b))$ are $(s - \mathrm{poly}(n), \varepsilon)$-indistinguishable for both $b = 0$ and $b = 1$. By Proposition 4.8, $(\mathbf{N}, \mathrm{Enc}(\mathbf{N}, 0))$ and $(\mathbf{N}, \mathrm{Enc}(\mathbf{N}, 1))$ are $(\infty, 1/4)$-indistinguishable. The corollary follows from the triangle inequality. □

5 The s-hCLWE and $(0, 1/2)$-hCLWE Distributions

In this section we introduce two distributions that are indistinguishable from $\mathcal{N}_n(0, 1)$ (i.e. n-dimensional vectors with i.i.d. entries from $\mathcal{N}(0, 1)$) by the CLWE assumption: the s-hCLWE and the $(0, 1/2)$-hCLWE distributions. Samples from the s-hCLWE distribution are CLWE samples $(\mathbf{y}_i, z_i)$ with $z_i = s$. Note that by definition the 0-hCLWE distribution is just the hCLWE distribution. Samples from the $(0, 1/2)$-hCLWE distribution are CLWE samples $(\mathbf{y}_i, z_i)$ with $z_i \in \{0, 1/2\}$. We obtain them by flipping a coin and, depending on the outcome, generating either an hCLWE sample or a 1/2-hCLWE sample. In the next two encryption schemes (”bimodal” in Section 6 and ”discretized” in Section 7) we use samples from the $(0, 1/2)$-hCLWE distribution to construct the public key.

To argue that these two distributions are indistinguishable from $\mathcal{N}_n(0, 1)$, we give a reduction from CLWE to both distributions. We also give a reduction from 1/2-hCLWE to hCLWE for completeness even though it is not needed in the rest of the paper.

5.1 The s-hCLWE Distribution

We begin by formally defining the distribution and then we show that there exists a reduction from CLWE.

Definition 5.1 (s-hCLWE Distribution)**.** *For a unit vector* $\mathbf{w} \in \mathbb{R}^n$*, real numbers* $\beta, \gamma > 0$*,* $n \in \mathbb{N}$ *and* $s \in [0,1]$*, samples* $\mathbf{y} \in \mathbb{R}^n$ *for the* s-hCLWE *distribution* $\mathcal{H}^s_{\mathbf{w},\beta,\gamma,n}$ *are generated as follows:*

1. *Sample* $k \in \mathbb{Z} + s$ *with probability proportional to* $\exp(-k^2/(2\gamma^2 + 2\beta^2))$*.*
2. *Sample* $e \leftarrow \mathcal{N}(0, \beta'^2)$*, where* $\beta'^2 := \beta^2/(\gamma^2 + \beta^2)$*.*
3. *Sample* $\mathbf{v}$ *as* $\mathcal{N}_{n-1}(0,1)$ *from the subspace orthogonal to* $\mathbf{w}$*.*
4. *Output* $\mathbf{y} := \mathbf{v} + (k/\gamma' + e)\mathbf{w}$*, where* $\gamma' := (\gamma^2 + \beta^2)/\gamma$*.*

It follows from the definition that hCLWE corresponds to the case $s = 0$. When $s = 0$, we write $\mathcal{H}_{\mathbf{w},\beta,\gamma,n}$ instead of $\mathcal{H}^0_{\mathbf{w},\beta,\gamma,n}$. The s-hCLWE distinguishing problem asks to distinguish between s-hCLWE samples and standard normal ones.

Definition 5.2 (s-hCLWE Distinguishing Problem)**.** *For real numbers* $\beta, \gamma > 0$*,* $n \in \mathbb{N}$ *and* $s \in [0,1]$*, the (average-case) distinguishing problem* $s\text{-hCLWE}_{\beta,\gamma,n}$ *asks to distinguish between* $\mathcal{H}^s_{\mathbf{w},\beta,\gamma,n}$ *for a uniform unit vector* $\mathbf{w} \in \mathbb{R}^n$ *and* $\mathcal{N}_n(0,1)$*.*

We do not consider the worst-case formulation of this problem as it is equivalent to the average-case one. The proof is analogous to [BRST21, Claim 2.22] for hCLWE and CLWE.

We now proceed to compare s-hCLWE to hCLWE and CLWE. First of all, using rejection sampling it is possible to obtain s-hCLWE samples from CLWE samples. This result follows from [BRST21, Lemma 4.1], which shows this for the case $s = 0$. Let $\mathcal{A}_{\mathbf{w},\beta,\gamma,n}$ denote the distribution of CLWE samples.

Lemma 5.3. *For a unit vector* $\mathbf{w} \in \mathbb{R}^n$*, real numbers* $\beta, \gamma > 0$*,* $n \in \mathbb{N}$ *and* $s \in [0,1]$*, there exists a probabilistic algorithm that runs in time poly*$(n, 1/\delta)$ *and that on input* $\delta \in (0,1)$ *and samples from* $\mathcal{A}_{\mathbf{w},\beta,\gamma,n}$*, outputs samples from* $\mathcal{H}^s_{\mathbf{w},\sqrt{\beta^2+\delta^2},\gamma,n}$*.*

Proof. The same proof as the one of Lemma 4.1 in [BRST21] with $g_0(z) := \sum_{k\in\mathbb{Z}} \rho_\delta(z + s + k)$. □

If we take $\delta = \beta/\sqrt{2}$, we obtain as a corollary the following reduction:

Proposition 5.4. *For* $s \in [0,1]$*,* $n \in \mathbb{N}$ *and real numbers* $\beta = \beta(n), \gamma = \gamma(n) > 0$ *such that* β *is the inverse of a polynomial in* n*, there exists a polynomial-time reduction from* $\text{CLWE}_{\beta/\sqrt{2},\gamma,n}$ *to* $s\text{-hCLWE}_{\beta,\gamma,n}$*.*

Now that we have given a reduction from CLWE to s-hCLWE it is a natural question to ask whether there is a reduction from s-hCLWE to CLWE. However, we do not know if this is possible for any value of s.

5.2 The $(0, 1/2)$-hCLWE Distribution

We now define the $(0,1/2)$-hCLWE distribution, which is the distribution on which the following two encryptions schemes are based. Afterwards we show that there is a reduction from CLWE to $(0,1/2)$-hCLWE.

Definition 5.5 ($(0,1/2)$-hCLWE Distribution)**.** *For a unit vector* $\mathbf{w} \in \mathbb{R}^n$ *and real numbers* $\beta, \gamma > 0$*,* $n \in \mathbb{N}$ *, samples* $(\mathbf{y}, z) \in \mathbb{R}^n \times \{0, 1/2\}$ *for the* $(0,1/2)$-hCLWE *distribution* $\mathcal{H}^{(0,\frac{1}{2})}_{\mathbf{w},\beta,\gamma,n}$ *are generated as follows:*

1. *Sample* $z \leftarrow \{0, 1/2\}$.

2. *Sample* $\mathbf{y} \leftarrow \mathcal{H}^z_{\mathbf{w},\beta,\gamma,n}$.

3. *Output* $(\mathbf{y}, z)$.

Definition 5.6 ($(0, 1/2)$-hCLWE Distinguishing Problem)**.** *For real numbers* $\beta, \gamma > 0$ *and* $n \in \mathbb{N}$ *, the (average-case) distinguishing problem* $(0, 1/2)$-hCLWE$_{\beta,\gamma,n}$ *asks to distinguish between* $\mathcal{H}^{(0,\frac{1}{2})}_{\mathbf{w},\beta,\gamma,n}$ *for a uniform unit vector* $\mathbf{w} \in \mathbb{R}^n$ *and* $\mathcal{N}_n(0, 1) \times \mathcal{U}(\{0, 1/2\})$.

Lemma 5.7. *For a unit vector* $\mathbf{w} \in \mathbb{R}^n$*,* $n \in \mathbb{N}$ *and real numbers* $\beta, \gamma > 0$ *, there exists a probabilistic algorithm that runs in time* $\mathrm{poly}(n, 1/\delta)$ *and that on input* $\delta \in (0, 1)$ *and samples from* $\mathcal{A}_{\mathbf{w},\beta,\gamma,n}$*, outputs samples from* $\mathcal{H}^{(0,\frac{1}{2})}_{\mathbf{w},\sqrt{\beta^2+\delta^2},\gamma,n}$.

Proof. We first sample $z \leftarrow \{0, 1/2\}$ uniformly at random. By Lemma 5.3 we can obtain a sample $\mathbf{y}$ from $\mathcal{H}^z_{\mathbf{w},\sqrt{\beta^2+\delta^2},\gamma,n}$ using samples from $\mathcal{A}_{\mathbf{w},\beta,\gamma,n}$ in time $\mathrm{poly}(n, 1/\delta)$ and $(\mathbf{y}, z)$ is a sample from $\mathcal{H}^{(0,\frac{1}{2})}_{\mathbf{w},\sqrt{\beta^2+\delta^2},\gamma,n}$. □

If we take $\delta = \beta/\sqrt{2}$, we obtain as a corollary the following result:

Proposition 5.8. *For* $n \in \mathbb{N}$ *and real numbers* $\beta = \beta(n), \gamma = \gamma(n) > 0$ *such that* β *is the inverse of a polynomial in* n*, there exists a polynomial-time reduction from* CLWE$_{\beta/\sqrt{2},\gamma,n}$ *to* $(0, 1/2)$-hCLWE$_{\beta,\gamma,n}$.

5.3 A reduction from $1/2$-hCLWE to hCLWE

Finally, we show that there exists a reduction from 1/2-hCLWE to hCLWE (with slightly different parameters) to get a finer understanding of the relative hardness of these phased hCLWE problems. We obtain the reduction by constructing samples from $\mathcal{H}_{\mathbf{w},\sqrt{2}\beta,\sqrt{2}\gamma,n}$ using samples from $\mathcal{H}^{1/2}_{\mathbf{w},\beta,\gamma,n}$.

Lemma 5.9. *For a unit vector* $\mathbf{w} \in \mathbb{R}^n$*,* $n \in \mathbb{N}$*, real numbers* $\beta, \gamma > 0$ *such that* $\gamma > \sqrt{n}$*, and independent random variables* Y_1, Y_2 *with distribution* $\mathcal{H}^{1/2}_{\mathbf{w},\beta,\gamma,n}$*, the distribution of* $(Y_1 - Y_2)/\sqrt{2}$ *is* e^{1-n}*-statistically close to* $\mathcal{H}_{\mathbf{w},\sqrt{2}\beta,\sqrt{2}\gamma,n}$.

Proof. By definition, $Y_i = \mathbf{v}_i + (k_i/\gamma' + e_i)\mathbf{w}$ for $i = 1, 2$ and

$$\frac{1}{\sqrt{2}}(Y_1 - Y_2) = \frac{1}{\sqrt{2}}(\mathbf{v}_1 - \mathbf{v}_2) + \left(\frac{1}{\sqrt{2}}\frac{k_1 - k_2}{\gamma'} + \frac{1}{\sqrt{2}}(e_1 - e_2)\right)\mathbf{w}$$

By standard properties of the normal distribution, it follows that $(\mathbf{v}_1 - \mathbf{v}_2)/\sqrt{2}$ has a $\mathcal{N}_{n-1}(0, 1)$ distribution and $(e_1 - e_2)/\sqrt{2}$ has a $\mathcal{N}(0, \beta'^2)$ distribution.

It remains to show that the distribution of $k_1 - k_1$ is statistically close to the discrete normal distribution over $\mathbb{Z}$ with variance $2(\gamma^2 + \beta^2)$. In order to apply Lemma 3.17, we first need a bound on the smoothing parameter of the lattice $\mathbb{Z}$. From Lemma 3.19 with $L = \mathbb{Z}$ and $\epsilon = \exp(-c^2)$ where $c = \sqrt{n}$, we get $\eta_\epsilon(\mathbb{Z}) \leq \sqrt{n} \leq \gamma$. By Lemma 3.17 with $L = \mathbb{Z}$ and $\mathbf{t} = 1/2$, we get that $k_1 - k_2$ is e^{1-n}-statistically close to $\mathcal{D}_{\mathbb{Z},2(\gamma^2+\beta^2)}$ as $\eta_\epsilon(\mathbb{Z}) \leq \gamma \leq \sqrt{\gamma^2 + \beta^2}$, which completes the proof. □

This gives the following result:

Proposition 5.10. *For $n \in \mathbb{N}$ and real numbers $\beta = \beta(n), \gamma = \gamma(n) > 0$, there exists a polynomial-time reduction from* $1/2$-hCLWE$_{\beta/\sqrt{2},\gamma/\sqrt{2},n}$ *to* hCLWE$_{\beta,\gamma,n}$.

6 Scheme 2: Bimodal Encryption

In this section we modify the ”pancake” scheme from Section 4 to achieve perfect correctness. Note that the decryption error in this scheme can be at least polynomial since the pancakes have polynomial width in the secret direction. This is due to the fact that the hCLWE assumption can be broken whenever the error distribution has exponentially small width as was shown in [BRST21]. A random normal vector therefore ”hits” a pancake with probability $1/\text{poly}(n)$. If we encrypt a 1 with such a vector, decryption fails. A standard approach to amplify the decryption error is sending multiple independent ciphertexts of the same message [DNR04]. This amplification increases the size of the ciphertext and the security error since a potential adversary only needs to be successful in decrypting one of the ciphertexts. Instead, we modify the encryption process of the bit 1. We introduce the following two changes:

- The public key consists of two matrices. A matrix $\mathbf{A}_0$ whose columns are independent hCLWE samples and a matrix $\mathbf{A}_1$ whose columns are independent 1/2-hCLWE samples. The samples from both matrices are obtained from the same secret direction $\mathbf{w}$.
- To encrypt a 0, take the matrix $\mathbf{A}_0$ and perform the same encryption as in the first scheme. To encrypt a 1, do exactly the same but with the matrix $\mathbf{A}_1$.

In Section 4 we have already seen that the decryption of Enc(0) is $1/\text{poly}(n)$-close to $0 \mod 1$. We show that in our modified scheme the decryption of Enc(1) is $1/\text{poly}(n)$ to $1/2$ so the scheme has perfect correctness. Security of the scheme follows by Proposition 4.8 and the triangle inequality.

6.1 The encryption scheme

The scheme is parametrized by $\gamma > 0$, $\beta > 0$, $n \in \mathbb{Z}$,$r > 0$ and $m \in \mathbb{Z} \setminus 2\mathbb{Z}$ an odd integer.

- The secret key is a uniformly random unit vector $\mathbf{w} \in \mathbb{R}^n$.
- The public key is a pair of matrices $(\mathbf{A}_0, \mathbf{A}_1) \in \mathbb{R}^{n\times m} \times \mathbb{R}^{n\times m}$. The columns of $\mathbf{A}_0$ are independent hCLWE samples and the columns of $\mathbf{A}_1$ are independent 1/2-hCLWE samples.
- To encrypt a bit $b \in \{0, 1\}$, compute
$$\mathbf{c} := \text{round}_r(\mathbf{A}_b\mathbf{t})$$
where $\mathbf{t} \leftarrow \{-1/\sqrt{m}, 1/\sqrt{m}\}^m$ is sampled uniformly at random. Check if all of the entries of $\mathbf{c}$ correspond to a bucket of width less than $1/(5\sqrt{nm}\gamma')$. If yes, output $\mathbf{c}$. If no, output b.
- To decrypt a ciphertext $\mathbf{c}$, take any $\mathbf{z}$ such that $\text{round}_r(\mathbf{z}) = \mathbf{c}$, compute
$$\gamma'\sqrt{m} \cdot \langle \mathbf{w}, \mathbf{z}\rangle \mod 1$$
and check if it is closer to 0 or closer to 1/2. In the former case output 0 in the latter case output 1.

Theorem 6.1. *Set the parameters of the scheme to* $\gamma = \sqrt{n}$, $\beta = (40000n^{5/2}\log(n)^2)^{-1}$, $r = (40000n^3\log(n))^{5/3}$ *and* $m = 10^8 n^2 \log(n)^2$. *Assuming* $(0, 1/2)$*-hCLWE*(s, ε) *we have that for all but a fraction of* $2^{-\Omega(n)}$ *choices of the public key the scheme has perfect correctness and security error at most* $1/2 + 1/n^2 + 3\varepsilon$.

We prove correctness and security of the scheme separately in the next two subsections.

6.2 Correctness

We call a public key good if the norm of the noise vector is less than $m\beta'$ in both matrices. By Corollary 3.11 this holds except with probability $2^{-\Omega(n)}$. During the construction of the public key it can be efficiently tested if a public key is good by checking if the absolute value of the generated noise value is small enough.

Claim 6.2. *If the public key is good, the scheme has perfect correctness.*

Proof. A preimage of a ciphertext $\mathbf{c}$ is of the form $\mathbf{z} = \mathbf{A}_b\mathbf{t} + \mathbf{e}_r$, where $\mathbf{e}_r$ denotes the rounding error. To decrypt one computes

$$\begin{aligned}
\gamma'\sqrt{m}\langle \mathbf{w}, \mathbf{z}\rangle &= \gamma'\sqrt{m}\langle \mathbf{w}, \mathbf{A}_b\mathbf{t} + \mathbf{e}_r\rangle \\
&= \gamma'\sqrt{m}(1/\gamma'\mathbf{k} - \mathbf{e}_b + b/(2\gamma')\cdot \mathbf{1})\mathbf{t} + \gamma'\sqrt{m}\mathbf{w}\mathbf{e}_r \\
&= (\mathbf{k} - \gamma'\mathbf{e}_b + b/2 \cdot \mathbf{1})\mathbf{1} + \gamma'\sqrt{m}\mathbf{w}\mathbf{e}_r \\
&= mb/2 - \gamma' e_b \cdot \mathbf{1} + \gamma'\sqrt{m}\mathbf{w}\mathbf{e}_r \mod 1 \\
&= b/2 - \gamma' e_b \cdot \mathbf{1} + \gamma'\sqrt{m}\mathbf{w}\mathbf{e}_r \mod 1,
\end{aligned}$$

for some integer vector $\mathbf{k} \in \mathbb{Z}^m$. Here $\mathbf{e}_b$ is the noise vector of the corresponding hCLWE or 1/2-hCLWE samples and $\mathbf{1}$ is the m-dimensional vector of all 1's. The second equality holds since $\langle \mathbf{w}, \mathbf{A_0}\rangle$ is a vector of multiples of $1/\gamma'$ minus the noise value and $\langle \mathbf{w}, \mathbf{A_1}\rangle$ is a vector of multiples of $1/\gamma'$ minus the noise value plus $1/2$. The last equality follows from the fact that m is an odd integer.

In order to show that $\gamma'\sqrt{m}\langle \mathbf{w}, \mathbf{z}\rangle$ is close to $b \cdot 1/2$ we bound the above expression by using Fact 3.5 and $\|\mathbf{w}\| = 1$. We get that

$$\left|\gamma'\sqrt{m}\langle \mathbf{w}, \mathbf{z}\rangle - b\cdot 1/2 \mod 1\right| \le \sum_{i\in[m]} |(\mathbf{e}_b)_i| + \gamma'\sqrt{m}\|\mathbf{e}_r\|.$$

Since our public key is good we have that $|(\mathbf{e}_b)_i| \le \sqrt{n}\beta'$ so $\sum_{i\in[m]} |(\mathbf{e}_b)_i| \le m\sqrt{n}\beta'$. We also know that each entry of e_r has absolute value less than $1/(5\sqrt{nm}\gamma')$ since the encryption process only outputs a ciphertext if this is the case. It follows that $\gamma'\sqrt{m}\|\mathbf{e}_r\| < 1/5$. By the choice of parameters we have

$$\left|\gamma'\sqrt{m}\langle \mathbf{w}, \mathbf{z}\rangle \mod 1\right| = b \cdot 1/2 \pm o(1/5 + 1/n),$$

which is closer to 0 if $b = 0$ and closer to 1/2 if $b = 1$. □

6.3 Security

There are two sources of security error in this scheme:

1. If at least one of the entries of the ciphertext corresponds to a bucket of width larger than $1/(5\sqrt{nm}\gamma')$, the encryption algorithm outputs the plaintext in the clear.

2. If the above event does not happen, the ciphertexts of 0 and of 1 are $1/2{+}2\varepsilon$-indistinguishable.

Claim 6.3. *Let $\mathbf{A}_b \in \mathbb{R}^{n\times m}$ be a matrix whose columns consist either of independent* hCLWE-*samples or of independent* 1/2-hCLWE *samples. Let $\mathbf{t} \leftarrow \{-1/\sqrt{m}, 1/\sqrt{m}\}^m$ be sampled uniformly at random. Assuming* $\mathrm{hCLWE}(s,\varepsilon)$ *and* $1/2\text{-hCLWE}(s,\varepsilon)$*, where s is the complexity of rounding, the probability that any entry of the vector $\mathbf{c} := \mathrm{round}_r(\mathbf{A}_b\mathbf{t})$ corresponds to a bucket of width larger than $1/(5\sqrt{m}\gamma')$ is at most $1/n^2 + \varepsilon$.*

Proof. First consider a matrix $\mathbf{A}$ with i.i.d. entries from $\mathcal{N}(0,1)$. Since $\|\mathbf{t}\| = 1$ we get that $\mathbf{At}$ is a vector with i.i.d. entries in $\mathcal{N}(0,1)$. By Proposition 4.1 we know that the number of intervals of length larger than $1/(5\sqrt{nm}\gamma')$ is at most $10\sqrt{nm}\gamma'/\sqrt{\ln(r/(5\sqrt{nm}\gamma'))} + 2$, so the probability that any entry lands in such a bucket is at most

$$\frac{10n\sqrt{nm}\gamma'}{r\sqrt{\ln(r/(5\sqrt{nm}\gamma'))}} + \frac{2n}{r} \le \frac{\gamma' n\sqrt{nm} + 2n}{r} \le \frac{1}{n^2}.$$

The claim follows from the fact that the matrices $\mathbf{A}_0$ and $\mathbf{A}_1$ are ε-indistinguishable from $\mathbf{A}$ and the rounding function being efficiently computable. □

Remark 6.4. *Note that we can avoid the above event by rejection sampling the public key. Since $\mathbf{t}$ is a unit vector, the absolute value of the inner product of any vector $\mathbf{a}$ with $\mathbf{t}$ is bounded by the norm of $\mathbf{a}$. This means that we can avoid the event that an entry of the ciphertext $\mathbf{c}$ corresponds to a wide bucket by rejection sampling the matrices $\mathbf{A}_0, \mathbf{A}_1$: As long as the rows of these matrices have small enough norm, the entries of the vector $\mathbf{A}_b\mathbf{t}$ will not land in a wide bucket for both $b \in \{0,1\}$. We omit a formal analysis of this optimization because the main security issue is not the rounding error but the probability of distinguishing ciphertexts of* 0 *and* 1 *as is shown by the next claim.*

Claim 6.5. *The distributions* $(\mathbf{N}_0, \mathbf{N}_1, \mathrm{Enc}(\mathbf{N}_0, 0))$ *and* $(\mathbf{N}_0, \mathbf{N}_1, \mathrm{Enc}(\mathbf{N}_1, 1))$ *are* 1/2*-statistically close for matrices $\mathbf{N}_0, \mathbf{N}_1$ of independent standard Gaussians.*

Proof. By Proposition 4.8 we have

$$\Delta((\mathbf{N}_0, \mathbf{N}_1, \mathrm{Enc}(\mathbf{N}_b, b)), (\mathbf{N}_0, \mathbf{N}_1, \mathbf{g})) \le 1/4,$$

where $\mathbf{g}$ is a vector with i.i.d. entries sampled uniformly from $\{1, 2, \ldots, r\}$ and $b \in \{0,1\}$. By the triangle inequality we follow that

$$\Delta((\mathbf{N}_0, \mathbf{N}_1, \mathrm{Enc}(\mathbf{N}_0, 0)), (\mathbf{N}_0, \mathbf{N}_1, \mathrm{Enc}(\mathbf{N}_1, 1))) \le 1/2.$$

□

Corollary 6.6. *Assuming* $(0, 1/2)\text{-hCLWE}(s,\varepsilon)$*,* $(\mathbf{A}_0, \mathbf{A}_1, \mathrm{Enc}(\mathbf{A}_0, 0))$ *and* $(\mathbf{A}_0, \mathbf{A}_1, \mathrm{Enc}(\mathbf{A}_1, 1))$ *are* $(s - \mathrm{poly}(n), 1/2 + 2\varepsilon)$*-indistinguishable where $\mathbf{A}_0, \mathbf{A}_1$ are the public key matrices.*

Proof. Let $\mathbf{N_0}, \mathbf{N_1}$ be standard normal matrices. By $(0, 1/2)\text{-hCLWE}(s,\varepsilon)$, $(\mathbf{A}_0, \mathbf{A}_1\mathrm{Enc}(\mathbf{A}_b, b))$ and $(\mathbf{N}_0, \mathbf{N}_1, \mathrm{Enc}(\mathbf{N}_b, b))$ are $(s - \mathrm{poly}(n), \varepsilon)$-indistinguishable for both $b = 0$ and $b = 1$. By Claim 6.5, $(\mathbf{N}_0, \mathbf{N}_1, \mathrm{Enc}(\mathbf{N}_0, 0))$ and $(\mathbf{N}_0, \mathbf{N}_1, \mathrm{Enc}(\mathbf{N}_1, 1))$ are $(\infty, 1/2)$-indistinguishable. The corollary follows from the triangle inequality. □

7 Scheme 3: Discretized Encryption

In this section we describe an encryption scheme based on CLWE that has negligible soundness error and perfect correctness for all but a fraction of $1/\text{poly}(n)$ many public keys. The scheme is inspired by the encryption scheme in [AD97] which also achieves negligible soundness error but only polynomial decryption error. We reduce this decryption error by applying their techniques to the bimodal encryption scheme from Section 6 which is based on $(0, 1/2)$-hCLWE. Alternatively, it could be applied to the baguette encryption scheme presented in Section 8 which would yield a scheme based on hCLWE. An important concept from [AD97] is the parallelepiped technique which enables us to transform continuous Gaussian samples into uniform ones. We first describe the technique before we present the encryption scheme and prove its correctness and security.

7.1 The parallelepiped technique and $\mathbb{Z}_q$

We will make use of the parallelepiped technique introduced by Ataj and Dwork in [AD97]. Let $\mathbf{B} = (\mathbf{b}_1, \ldots, \mathbf{b}_n) \in \mathbb{R}^{n\times n}$ be an arbitrary matrix of rank n. We denote by $\mathcal{P}(\mathbf{B})$ the n-dimensional parallelepiped that is defined by the columns of $\mathbf{B}$, i.e.

$$\mathcal{P}(\mathbf{B}) := \left\{ \sum_{i\in[n]} \lambda_i \mathbf{b}_i : 0 \leq \lambda_i < 1 \text{ for all } i \in [n] \right\}.$$

We denote by $\mathcal{P}_q(\mathbf{B})$ the set we obtain by partitioning $\mathcal{P}(\mathbf{B})$ into q^n smaller parallelpipeds of equal volume and then rounding each vector to the lower left corner of the corresponding smaller parallelepiped, i.e.

$$\mathcal{P}_q(\mathbf{B}) := \left\{ \mathbf{B}\lfloor q\mathbf{B}^{-1}\mathbf{c}\rfloor q^{-1} : \mathbf{c} \in \mathcal{P}(\mathbf{B}) \right\}.$$

We will later need the following facts:

Fact 7.1. *The distance between a vector $c \in \mathcal{P}(\mathbf{B})$ and its rounded image $\mathbf{c}' := \mathbf{B}\lfloor q\mathbf{B}^{-1}\mathbf{c}\rfloor q^{-1}$ is at most $q^{-1} n \max_i \|\mathbf{b}_i\|$.*

The above fact follows from the observation that the maximum length of an edge of the small parallelepiped is $\max_i \|\mathbf{b}_i\|/q$ and by computing the distance from the "lower left" corner to the "upper right" corner of the parallelepiped.

Fact 7.2. *Let $\mathbf{B} = (\mathbf{b}_1, \ldots, \mathbf{b}_n) \in \mathbb{R}^{n\times n}$ be an arbitrary matrix of rank n. Then $(\mathcal{P}_q(\mathbf{B}), +)$ is a group isomorphic to $\mathbb{Z}_q^n$.*

This can be seen by the following argument: We obtain $\mathcal{P}_q(\mathbf{B})$ by partitioning each vector $\mathbf{b}_i$ into q equal parts. Labelling the parts by $\{0, 1, 2, \ldots, q-1\}$ in the natural way gives an isomorphism between the q parts of $\mathbf{b}_i$ and $\mathbb{Z}_q$ for any $i \in [n]$. Fact 7.2 follows by taking the direct product of the labellings of the $\mathbf{b}_i$.

In the construction of our public key we essentially map continuous Gaussian vectors into $\mathcal{P}(\mathbf{B})$. We will need the next lemma to show that this mapping transforms them into uniformly random vectors. We denote by $\eta_\varepsilon(\mathbf{B})$ the smoothing parameter of the lattice with basis $\mathbf{B}$.

Lemma 7.3 ([MR07, Lemma 4.1])**.** *Let $\mathbf{B} \in \mathbb{R}^{n\times n}$ be a square matrix of rank n. For any $\varepsilon > 0$ and any $s > \eta_\varepsilon(\mathbf{B})$ the statistical distance between $\mathcal{N}_n(0, s^2) \mod \mathbf{B}$ and the uniform distribution over $\mathcal{P}(\mathbf{B})$ is at most $\varepsilon/2$.*

7.2 The encryption scheme

The scheme is parametrized by $\gamma > 0$; $\beta > 0$; $n, m, q \in \mathbb{Z} \setminus 2\mathbb{Z}$ odd integers. We set n to be an odd integer only to clarify the description and the analysis, m and q however are always required to be odd.

- The secret key is a uniformly random unit vector $\mathbf{w} \in \mathbb{R}^n$.
- The public key is a tuple of matrices $(\mathbf{A}_0, \mathbf{A}_1, \mathbf{B}) \in \mathbb{R}^{n\times m} \times \mathbb{R}^{n\times m} \times \mathbb{R}^{n\times n}$. We obtain the public key as follows: We let $\mathbf{B} \in \mathbb{R}^{n\times n}$ be a matrix whose columns consist of hCLWE samples, such that the smallest singular value of $\mathbf{B}$ is larger than $1/m$. To generate one column of $\mathbf{A}_0$ we produce n^2 many samples $\mathbf{a}_i$ from the hCLWE distribution $\mathcal{H}_{\mathbf{w}.\beta,\gamma,n}$ and compute

  $$\mathbf{a} := B\text{-round}\left(\sum_{i\in[n^2]} \mathbf{a}_i \mod \mathbf{B}\right),$$

  where B-round $= B\text{-round}_q : \mathbb{R}^n \to \mathcal{P}_q(\mathbf{B})$ is defined as $B\text{-round}_q(a) = \mathbf{B}\lfloor q\mathbf{B}^{-1}a\rceil/q$. Repeating this process m times gives the columns of $\mathbf{A}_0$. To generate the columns of $\mathbf{A}_1$ we do the same with 1/2-hCLWE samples from $\mathcal{H}^{\frac{1}{2}}_{\mathbf{w},\beta,\gamma,n}$.
- To encrypt a bit $b \in \{0, 1\}$, compute

  $$\mathbf{c} := \mathbf{A}_b\mathbf{t} \mod \mathbf{B},$$

  where $\mathbf{t} \leftarrow \{-1, 1\}^m$ is sampled uniformly at random.
- To decrypt a ciphertext $\mathbf{c}$, compute

  $$\gamma'\langle\mathbf{w}, \mathbf{c}\rangle \mod 1$$

  and check if it is closer to 0 or closer to 1/2. In the former case output 0 in the latter case output 1.

Remark 7.4. *In the next section we will see that we require n to be an odd integer only because we need that the inner product of $\mathbf{w}$ with the sum of n^2 many* 1/2-hCLWE *samples is approximately* $1/2 \mod 1$ *and not* 0*. One can slightly change the scheme for even values of n: obtain one column of the matrices $\mathbf{A}_b$ by using the sum of $n^2 + 1$ samples instead of n. In the rest of the section we will assume that n is odd without loss of generality.*

Theorem 7.5. *Set the parameters of the scheme to $\gamma = \sqrt{n}, m = n^2, \beta = 1/n^{10}, q = n^7$. Assuming* $(0, 1/2)$-hCLWE(s, ε) *we get that for all but a fraction of $1/n^{3/2} + O(\varepsilon)$ choices of the public key the scheme has perfect correctness and negligible soundness error.*

We prove correctness and soundness of the scheme separately in the next two subsections.

7.3 Correctness

We show that for all but a fraction of at most $1/n^{3/2} + \varepsilon$ choices of the public key decryption is always correct. We denote by $\{\mathbf{b}_1, \dots, \mathbf{b}_n\}$ the columns of $\mathbf{B}$, by $\{\mathbf{a}_1^0, \dots, \mathbf{a}_{n^2m}^0\}$ the hCLWE samples used to construct $\mathbf{A}_0$ and by $\{\mathbf{a}_1^1, \dots, \mathbf{a}_{n^2m}^1\}$ the 1/2-hCLWE samples used to construct $\mathbf{A}_1$. We define $\mathbf{e} := \gamma'\mathbf{w}\mathbf{B} \mod 1$ which is the noise vector of the hCLWE samples $\mathbf{b}_i$. For $b \in \{0,1\}$ we define

$$\mathbf{e}_b := \gamma'\mathbf{w}\left(\sum_{i=1}^{n^2}\mathbf{a}_i^b, \sum_{i=n^2+1}^{2n^2}\mathbf{a}_i^b, \dots, \sum_{i=(m-1)n^2+1}^{mn^2}\mathbf{a}_i^b\right) - b\cdot(1/2, 1/2, \dots, 1/2)^T \mod 1.$$

If $b = 0$ this is the vector where each entry is the sum of the n noise values corresponding to the hCLWE samples that we add during the construction of $\mathbf{A}_0$. If $b = 1$ this is the noise vector we get during the construction of $\mathbf{A}_1$. We call a public key $(\mathbf{A}_0, \mathbf{A}_1, \mathbf{B})$ *good* if the following holds:

1. $\|\mathbf{e}_0\|, \|\mathbf{e}_1\| \leq mn\beta'$;
2. $\|\mathbf{e}\| \leq n\beta'$;
3. For all $i \in [nm]$ the entries of $\mathbf{a}_i^0, \mathbf{a}_i^1$ lie in the interval $\left[-n^{3/2}, n^{3/2}\right]$;
4. For all $i \in [n]$ the entries of $\mathbf{b}_i$ lie in the interval $[-n, n]$;
5. the smallest singular value of $\mathbf{B}$ is larger than $1/m$.

Note that all of these conditions can be efficiently tested during the key generation.

Claim 7.6. *If the* $(0, 1/2)$-hCLWE(s, ε) *assumption holds, a public key* $(\mathbf{A}_0, \mathbf{A}_1, \mathbf{B})$ *is good except with probability* $1/n^{3/2} + O(\varepsilon)$.

Proof. By Corollary 3.11 conditions 1 and 2 hold except with negligible probability. Next we consider a hybrid where the matrices $\mathbf{A}_0, \mathbf{A}_1, \mathbf{B}$ are replaced by $\tilde{\mathbf{A}}_0, \tilde{\mathbf{A}}_1, \tilde{\mathbf{B}}$ which are not obtained from $(0, 1/2)$-hCLWE samples but from i.i.d. random Gaussian samples and bound the probability that the rest of the conditions for a good public key hold. By Fact 3.10 condition 3 and 4 hold except with negligible probability. To bound the probability that the smallest singular value $s_n(\tilde{\mathbf{B}})$ is less than $1/m$ we use Fact 3.7. We get that

$$\Pr\left[s_n(\tilde{\mathbf{B}}) \leq 1/m\right] \leq \frac{\sqrt{n}}{m} = \frac{1}{n^{3/2}}.$$

The claim follows from the observation that if the above probability bounds differed for our matrices $\mathbf{A}_0, \mathbf{A}_1, \mathbf{B}$ by more than ε, we could efficiently distinguish between a random normal matrix and matrices that consist of $(0, 1/2)$-hCLWE samples and therefore break the $(0, 1/2)$-hCLWE(s, ε) assumption by checking the absolute values of the matrices $\mathbf{A}_0, \mathbf{A}_1$ and computing the Eigenvalues of the matrix $\mathbf{B}^T\mathbf{B}$. □

Claim 7.7. *If the public key* $(\mathbf{A}_0, \mathbf{A}_1, \mathbf{B})$ *is good, decryption is correct with probability* 1.

Proof. An encryption of a bit b is of the form

$$\begin{aligned}\mathbf{c} &= \mathbf{A}_b\mathbf{t} \mod \mathbf{B}\\ &= \mathbf{A}_b\mathbf{t} - \mathbf{Bs}\\ &= (\tilde{\mathbf{A}}_b - \mathbf{BZ} + \mathbf{E}_q)\mathbf{t} - \mathbf{Bs}\end{aligned}$$

for some $\mathbf{s} \in \mathbb{Z}^n, \mathbf{Z} \in \mathbb{Z}^{n\times m}$. Here $\mathbf{E}_q$ is a matrix whose entries are rounding errors and $\tilde{\mathbf{A}}_b := \mathbf{A}_b + \mathbf{BZ} - \mathbf{E}_q$ is a matrix whose columns are the sum of n^2 hCLWE samples if $b = 0$ or the sum of n^2 many 1/2-hCLWE samples if $b = 1$. In other words $\tilde{\mathbf{A}}_b$ is the matrix we get in the construction of the matrix $\mathbf{A}_b$ before rounding and before mapping to the parallelepiped $\mathcal{P}(\mathbf{B})$. To decrypt one computes

$$\begin{aligned}\gamma'\langle \mathbf{w}, \mathbf{c}\rangle &= \gamma'\langle \mathbf{w}, (\tilde{\mathbf{A}}_b - \mathbf{BZ} + \mathbf{E}_q)\mathbf{t} - \mathbf{Bs}\rangle \mod 1\\ &= (b/2 \cdot \mathbf{1} + \gamma'\mathbf{e}_b - \gamma'\mathbf{e}^{\mathbf{T}}\mathbf{Z} + \gamma'\mathbf{w}^T\mathbf{E}_q)\mathbf{t} - \gamma'\langle \mathbf{e}, \mathbf{s}\rangle \mod 1\\ &= b/2 + \gamma'(\mathbf{e}_b - \mathbf{e}^{\mathbf{T}}\mathbf{Z} + \mathbf{w}^T\mathbf{E}_q)\mathbf{t} - \gamma'\langle \mathbf{e}, \mathbf{s}\rangle \mod 1,\end{aligned}$$

where $\mathbf{1}$ is the m-dimensional vector of all 1's. The last equality follows from the fact that m is odd.

In order to show that $\langle \mathbf{w}, \mathbf{c}\rangle$ is close to $b \cdot 1/2$ we bound the above expression by repeatedly using Fact 3.5 and $\|\mathbf{w}\| = 1$. We get that

$$\left|\gamma'\langle \mathbf{w}, \mathbf{c}\rangle - b\cdot 1/2 \mod 1\right| \le \gamma'(\|\mathbf{e}_b\| + \|\mathbf{e}\| \cdot \|\mathbf{Z}\| + \|\mathbf{E}_q\|) \cdot \|\mathbf{t}\| + \gamma'\|\mathbf{e}\| \cdot \|\mathbf{s}\|$$

Since our public key is good we have that $\|\mathbf{e}_b\| \le mn\beta'$ and $\|\mathbf{e}\| \le n\beta'$. We have $\|\mathbf{t}\| = \sqrt{m}$ since its entries have absolute value 1. By Facts 3.6 and 7.1 we have $\|\mathbf{E}_q\| \le \sqrt{m \cdot \max_i \|(\mathbf{E}_q)_i\|^2} \le \sqrt{m}nq^{-1} \max_i \|\mathbf{b}_i\| \le \sqrt{m}n^{5/2}/q$, where the last inequality follows from the fact that our public key is good and so the absolute values of the entries of $\mathbf{B}$ are at most n. It remains to bound the norms of $\mathbf{Z}$ and $\mathbf{s}$.

Claim 7.8. *Let $1/\alpha$ be the smallest singular value of $\mathbf{B}$. We have $\|\mathbf{s}\| \le \alpha n^{5/2}(m+1)$ and $\|\mathbf{Z}\| \le 2\alpha n^3\sqrt{m}$.*

Plugging in these values we get

$$\left|\gamma'\langle \mathbf{w}, \mathbf{c}\rangle - b/2 \mod 1\right| \le \beta'\gamma'(nm\sqrt{m} + 2n^4\alpha m + n^{7/2}\alpha(m+1)) + \gamma' q^{-1}n^{5/2}m$$

Since our public key is good and by the choice of our parameters we have that $\alpha \le m$, $\beta = 1/n^{10}$, $\gamma = \sqrt{n}$ and $q = n^7$. It follows that $\gamma'\langle \mathbf{w}, \mathbf{c}\rangle = b \cdot 1/2 + o(1/n) \mod 1$. □

Proof of Claim 7.8. By definition we have $\mathbf{s} = \mathbf{B}^{-1}(\mathbf{A}_b\mathbf{t} - \mathbf{r})$ and $\mathbf{Z} = \mathbf{B}^{-1}(\tilde{\mathbf{A}'_b} - \mathbf{R})$, where $\mathbf{r}$ and the columns of $\mathbf{R}$ are vectors in $\mathcal{P}(\mathbf{B})$. Since the smallest singular value of $\mathbf{B}$ is $1/\alpha$ we have that the largest singular value of $\mathbf{B}^{-1}$ is α so $\|\mathbf{B}^{-1}\| = \alpha$. Furthermore, since our public key is good the entries of the matrices are not too large so

$$|(\mathbf{A}_b\mathbf{t} - \mathbf{r})_i| \le (m+1)\sum_j |(\mathbf{B})_{ij}| \le n^2(m+1)$$

and hence $\|\mathbf{A}_b\mathbf{t} - \mathbf{r}\| \leq \sqrt{n^5(m+1)^2} = n^{5/2}(m+1)$. By the same argument we can bound

$$\left|(\tilde{\mathbf{A}}_b - \mathbf{R})_{ij}\right| \leq \left|(\tilde{\mathbf{A}}_b)_{ij}\right| + \sum_j |(\mathbf{B})_{ij}| = n^2(\sqrt{n}+1)$$

and get $\|\tilde{\mathbf{A}}_b - \mathbf{R}\| \leq 2n^3\sqrt{m}$ by Fact 3.6. □

7.4 Security

We show that encryptions of 0 and 1 are indistinguishable under the $(0, 1/2)$-hCLWE assumption by showing that the following distributions are indistinguishable for $b \in \{0, 1\}$:

1. Real_b: $(\mathbf{A}_0, \mathbf{A}_1, \mathbf{B}, \mathbf{A}_b t \mod \mathbf{B})$ is a public key of the encryption scheme.
2. Hybrid_b: $(\mathbf{A}_0, \mathbf{A}_1, \mathbf{B}, \mathbf{A}_b t \mod \mathbf{B})$ is a tuple where the entries of $\mathbf{B}$ are independently sampled from $\mathcal{N}(0, 1)$, the columns of $\mathbf{A}_0$ and $\mathbf{A}_1$ are uniformly random vectors in $\mathcal{P}_q(\mathbf{B})$.
3. Ideal: $(\mathbf{A}_0, \mathbf{A}_1, \mathbf{B}, \mathbf{r})$ is the same as above but with $\mathbf{r}$ a uniformly random vector in $\mathcal{P}_q(\mathbf{B})$.

Real_b and Hybrid_b are computationally indistinguishable under the $(0, 1/2)$-hCLWE assumption. Hybrid_b and Ideal are statistically indistinguishable by the leftover hash lemma. In the rest of the section we formally prove the above statements. We start by showing the first claim.

Claim 7.9. *Under the* $(0, 1/2)$-hCLWE(s, ε) *assumption the distributions* Real_b *and* Hybrid_b *are* $(s - \text{poly}(n), 2^{-n+1} + \varepsilon)$*-indistinguishable.*

Proof. Assume that there is a distinguisher D that decides if $(\mathbf{A}_0, \mathbf{A}_1, \mathbf{B}, \mathbf{A}_b t \mod \mathbf{B})$ is from Real_b or from Hybrid_b with probability δ. We construct an algorithm D' that distinguishes between $(0, 1/2)$-hCLWE samples and random samples with probability $\delta - 2^{-n+1}$ as follows:

1. Given poly(n) many $(0, 1/2)$-hCLWE samples $\{(\mathbf{y}_i, z_i)\}_{i \in [\text{poly}(n)]}$, define a matrix $\mathbf{B}$ by choosing n samples with $z_i = 0$ such that the corresponding vectors $\mathbf{y}_i$ are linearly independent. These vectors are the columns of $\mathbf{B}$.
2. Repeat the following procedure m times: choose n^2 samples of the form $\{(\hat{\mathbf{y}}_i, 0)\}_{i \in [n^2]}$ and compute
$$\mathbf{y}_0 = B\text{-round}\left(\sum_{i \in [n^2]} \hat{\mathbf{y}}_i \mod \mathbf{B}\right)$$
and choose n samples of the form $\{(\tilde{\mathbf{y}}_i, 1/2)\}_{i \in [n^2]}$ and compute
$$\mathbf{y}_1 = B\text{-round}\left(\sum_{i \in [n^2]} \tilde{\mathbf{y}}_i \mod \mathbf{B}\right),$$
where $B\text{-round} = B\text{-round}_q : \mathbb{R}^n \to \mathcal{P}_q(\mathbf{B})$ is defined as $B\text{-round}_q(\mathbf{a}) = \mathbf{B}\lfloor q\mathbf{B}^{-1}\mathbf{a}\rceil / q$.
3. Let $\mathbf{A}_0$ be the matrix with m columns generated as above from the samples with $z_i = 0$ and $\mathbf{A}_1$ be the matrix with m columns generated as above from the samples with $z_i = 1/2$. Give $(\mathbf{A}_0, \mathbf{A}_1, \mathbf{B}, \mathbf{A}_b t \mod \mathbf{B})$ to the distinguisher D.

Note that in the case where the samples $\{(\mathbf{y}_i, z_i)\}_{i\in[\mathrm{poly}(n)]}$ are $(0,1/2)$-hCLWE samples, $\mathbf{A}_0, \mathbf{A}_1, \mathbf{B}$ is a public key of our scheme. It remains to prove that given samples $\{(\mathbf{y}_i, z_i)\}_{i\in[\mathrm{poly}(n)]}$, where the $\mathbf{y}_i$ are normal random vectors and the z_i are uniform in $\{0,1/2\}$, the resulting matrices $\mathbf{A}_0, \mathbf{A}_1$ are statistically close to uniform matrices in $\mathcal{P}_q(\mathbf{B})$. Lemma 7.3 says that if we sample a vector from a Gaussian distribution with standard deviation larger than $\eta_{2^{-n}}(\mathbf{B})$ and map it into $\mathcal{P}_q(\mathbf{B})$, the resulting vector is statistically close to uniform in $\mathcal{P}_q(\mathbf{B})$.

Now we only need an upper bound on the smoothing parameter in order to prove that the columns of $\mathbf{A}_0$ and $\mathbf{A}_1$ are sampled from a Gaussian with sufficiently large variance.

By Corollary 3.11 the length of a vector with entries independently sampled from $\mathcal{N}(0,1)$ is at most n except with probability $\sqrt{n}e^{-n}$. Hence, the smoothing parameter of $\mathbf{B}$ is at most $n^{3/2}$ by Lemma 3.20 except with probability $\sqrt{n}e^{-n}$. The entries of $\mathbf{A}_0$ and $\mathbf{A}_1$ are sampled from $\mathcal{N}(0,n^2)$. Since $n^2 > n^{3/2}$ we follow from Lemma 7.3 that $\mathbf{A}_0$ and $\mathbf{A}_1$ are 2^{-n+1}-statistically close to uniformly random matrices in $\mathcal{P}_q(\mathbf{B})$. □

Next we show that Hybrid_b is statistically close to Ideal, which completes the proof of soundness. By Fact 7.2 we know that $(\mathcal{P}_q(\mathbf{B}),+)$ is a group isomorphic to $\mathbb{Z}_q^n$ for any full rank $n\times n$ matrix $\mathbf{B}$. It is therefore sufficient to prove statistical closeness of the tuples $(\hat{\mathbf{A}}_0, \hat{\mathbf{A}}_1, \hat{\mathbf{A}}_b\mathbf{t} \mod q)$ and $(\hat{\mathbf{A}}_0, \hat{\mathbf{A}}_1, \hat{\mathbf{r}})$, where $\hat{\mathbf{A}}_0, \hat{\mathbf{A}}_1$ are matrices with i.i.d. uniform entries in $\mathbb{Z}_q$ and $\hat{\mathbf{r}}$ is a uniform vector in $\mathbb{Z}_q^n$. This can be done using the classical leftover hash lemma [IZ89]. To this end we need to show that multiplication of a $\{-1,1\}^m$ vector by a uniform matrix $\mathbf{H}\in\mathbb{Z}_q^{m\times n}$ is a universal family of hash functions, i.e.:

Claim 7.10. *For q odd, $\mathbf{x},\mathbf{y}\in\{-1,1\}^m$ such that $\mathbf{x}\neq\mathbf{y}$ we have*

$$\Pr_{\mathbf{H}\leftarrow\mathbb{Z}_q^{m\times n}}[\mathbf{H}\mathbf{x} = \mathbf{H}\mathbf{y} \mod q] = \frac{1}{q^n}.$$

Proof. Since $\mathbf{x}\neq\mathbf{y}$ we know that they differ in at least one coordinate. Without loss of generality assume that $x_i = 1$ and $y_i = -1$. Choose all of $\mathbf{H}$ except for the i-th column $\mathbf{h}_i$. We have that $\mathbf{H}\mathbf{x} = \mathbf{b}+\mathbf{h}_i \mod q$ and $\mathbf{H}\mathbf{y} = \mathbf{c}-\mathbf{h}_i \mod q$ for some fixed $\mathbf{b}$ and $\mathbf{c}$. This means that $\mathbf{H}\mathbf{x} = \mathbf{H}\mathbf{y}$ if and only if $\mathbf{b}+\mathbf{h}_i = \mathbf{c}-\mathbf{h}_i$ which is equivalent to $2\mathbf{h}_i = \mathbf{c}-\mathbf{b}$. Since q is odd we can divide by 2 and get that $\mathbf{H}\mathbf{x}+\mathbf{H}\mathbf{y}$ if and only if $\mathbf{h}_i = 2^{-1}(\mathbf{c}-\mathbf{b}) \mod q$. This holds for exactly one choice of $\mathbf{h}_i$ in $\mathbb{Z}_q^n$ which concludes the proof. □

The following lemma is a special case of the leftover hash lemma [IZ89, Reg05] :

Lemma 7.11. *Let q be an odd integer. Let $\mathbf{H}\in\mathbb{Z}_q^{n\times m}$ be a matrix with columns chosen uniformly at random from $\mathbb{Z}_q^n$ and $\mathbf{t}\leftarrow\{-1,1\}^m$ a uniformly random vector. Then the statistical distance of the uniform distribution on $\mathbb{Z}_q^n$ and the distribution given by multiplying $\mathbf{H}$ with $\mathbf{t}$ is at most $(q^n/2^m)^{1/4}$ with probability $1-(q^n/2^m)^{1/4}$.*

By our choice of parameters we have $m = n^2$ and $q = n^7$. We follow that the statistical distance of Hybrid_0 and Hybrid_1 to Ideal is $(n^{7n}/2^{n^2})^{1/4}\leq 2^{-n}$ for large enough values of n. Hence, Hybrid_0 is at least 2^{-n+1}-close to Hybrid_1. Together with Claim 7.9 this yields that an encryption of 0 is $2^{-n+2}+2\varepsilon$-indistinguishable from an encryption of 1.

8 Scheme 4: Baguette Encryption

In this section we present a second approach that reduces the decryption error of the pancake scheme. The security error remains constant but could be reduced by the parallelepiped technique presented in Section 7. Instead of publishing samples that have a pancake distribution in only one secret direction, we publish samples that have a pancake distribution in multiple secret directions, i.e. samples from the hCLWE(ℓ) distribution. This is a distribution defined in [BRST21] to which the authors give a reduction from hCLWE. To decrypt we take the inner products of the ciphertext with all secret directions. If the ciphertext is an encryption of 0 all of the results are polynomially close to an integer. If the ciphertext is an encryption of 1, at least one of the results is not close to an integer with high probability since taken modulo 1 they are uniformly random values in $[0, 1)$. Before presenting the encryption scheme we formally define the hCLWE(ℓ) distribution.

8.1 The hCLWE(ℓ) distribution

Both the hCLWE(ℓ), distribution and the corresponding decision problem were introduced in [BRST21]. This problem is the extension of hCLWE to the case of ℓ hidden orthogonal directions.

Definition 8.1 (hCLWE(ℓ) Distribution)**.** *For a matrix $\mathbf{W} = (\mathbf{w}_1|\dots|\mathbf{w}_\ell) \in \mathbb{R}^{n\times\ell}$ such that $\mathbf{W}^T\mathbf{W} = \mathbf{I}_\ell$, real numbers $\beta, \gamma > 0$, $n \in \mathbb{N}$ and $\ell \in \mathbb{N}$ with $0 \leq \ell \leq n$, samples $\mathbf{y} \in \mathbb{R}^n$ for the* hCLWE(ℓ) *distribution $\mathcal{H}_{\mathbf{W},\beta,\gamma,n,\ell}$ are generated as follows:*

1. *Sample $k_1, \dots, k_\ell \in \mathbb{Z}$ independently with distribution $\mathcal{D}_{\mathbb{Z},\gamma^2+\beta^2}$.*
2. *Sample $e_1, \dots, e_\ell \leftarrow \mathcal{N}(0, \beta'^2)$ independently where $\beta'^2 := \beta^2/(\gamma^2 + \beta^2)$.*
3. *Sample $\mathbf{v}$ as $\mathcal{N}_{n-\ell}(0, 1)$ from the subspace orthogonal to $\mathbf{W}$.*
4. *Output $\mathbf{y} := \mathbf{v} + \sum_{i=1}^{\ell}(k_i/\gamma' + e_i)\mathbf{w}_i$ where $\gamma' := (\gamma^2 + \beta^2)/\gamma$.*

For $\ell = 0$ we get the normal distribution with covariance matrix $\mathbf{I}_n$ and for $\ell = 1$ we recover the hCLWE distribution. We refer to the columns of $\mathbf{W}$ as the hidden directions. Note that they are orthonormal vectors.

Definition 8.2 (hCLWE(ℓ) Distinguishing Problem)**.** *For real numbers $\beta, \gamma > 0$, $n \in \mathbb{N}$ and $\ell \in \mathbb{N}$ with $0 \leq \ell \leq n$, the (average-case) distinguishing problem* $\text{hCLWE}_{\beta,\gamma,n}(\ell)$ *asks to distinguish between $\mathcal{H}_{\mathbf{W},\beta,\gamma,n,\ell}$ for a uniform matrix $\mathbf{W} \in \mathbb{R}^{n\times\ell}$ such that $\mathbf{W}^T\mathbf{W} = \mathbf{I}_\ell$, and $\mathcal{N}_n(\mathbf{0}, 1)$.*

The hCLWE(ℓ)(s, ϵ) assumption postulates that the hCLWE(ℓ) distinguishing problem cannot be solved in size s with advantage ϵ. As shown in [BRST21] (Lemma 9.3.), if $n - \ell = \Omega(n^k)$ for some constant $k > 0$, there is an efficient reduction from $\text{hCLWE}_{\beta,\gamma,n-\ell+1}$ to $\text{hCLWE}_{\beta,\gamma,n}(\ell)$.

8.2 Encryption scheme

We now give an encryption scheme that builds on the pancake scheme from Section 4. It achieves negligible decryption error using more hidden directions instead of the $(0, 1/2)$-hCLWE distribution.

The scheme is parametrized by $\gamma > 0$; $\beta > 0$; $r > 0$, $n, \ell, m \in \mathbb{N}$ and a parameter $a > 0$ for which we will only consider two possible values, namely, $a = 1/n$ and $a = 1/100$.

- The secret key is a uniformly random matrix $\mathbf{W} \in \mathbb{R}^{n \times \ell}$ such that $\mathbf{W}^T\mathbf{W} = \mathbf{I}_\ell$.
- The public key is a matrix $\mathbf{A} \in \mathbb{R}^{n \times m}$ whose columns are independently sampled from $\mathcal{H}_{\mathbf{W},\beta,\gamma,n,\ell}$.
- To encrypt 0, choose a vector $\mathbf{t} \in \{-1/\sqrt{m}, +1/\sqrt{m}\}^m$ uniformly at random and output
$$\mathbf{c} := \mathrm{round}_r(\mathbf{At}).$$
Check if all entries of $\mathbf{c}$ correspond to buckets of width less than $1/(4a\sqrt{n}\sqrt{m}\gamma')$. If yes, output $\mathbf{c}$. Otherwise, output 0.
- To encrypt 1, choose a vector $\mathbf{c} \leftarrow \{1, 2, \ldots, r\}^n$ uniformly at random. Check if all entries of $\mathbf{c}$ correspond to buckets of width less than $1/(4a\sqrt{n}\sqrt{m}\gamma')$. If yes, output $\mathbf{c}$. Otherwise, output 1.
- To decrypt a ciphertext $\mathbf{c}$, take any $\mathbf{z}$ such that $\mathrm{round}_r(\mathbf{z}) = \mathbf{c}$, compute
$$\gamma'\sqrt{m}\mathbf{W}^T\mathbf{z} \mod 1$$
and check if all ℓ entries are in $(-1/2a, 1/2a)$. If yes, output 0, else output 1.

Theorem 8.3. *Set the parameters of the scheme to $\gamma = \sqrt{n}$, $\beta = (16 \cdot 10^4 n^3 \log(n))^{-1}$, $\ell = \log n$, $m = 10^8 n^2 \log(n)^2$, $r = (40001 n^3 \log(n))^{5/3}$ and $a = 1/n$. Assuming* hCLWE(s, ε)*, the scheme has negligible decryption error and security error at most $1/4 + 4\varepsilon$.*

We prove correctness and security of the scheme separately in the next two subsections.

We are also interested in using this scheme to prove that hCLWE and hCLWE(ℓ) are in SZK (statistical zero knowledge), what is shown in Section 9 for the following choice of parameters:

$$\begin{aligned} a &= 100 \\ \beta'\gamma'\ln\gamma' &< \frac{1}{4 \cdot 10^4 K n \log n} \\ \gamma' &> 1 \\ m &= (Kn\log n \ln\gamma')^2 \\ r &= m^{10}(\gamma')^{5/3} \end{aligned} \tag{4}$$

where $K = 4 \cdot 9 \cdot 10 \cdot e \cdot 2 \cdot 5$.

8.3 Correctness

Claim 8.4. *The probability that* $\mathrm{Dec}(\mathbf{W}, \mathrm{Enc}(\mathbf{A}, 0)) = 0$ *over the joint choice of the public key and encryption randomness is at least*

$$1 - \ell\sqrt{\frac{2\beta'^2\gamma'^2 m}{\pi}}\frac{e^{-\frac{(1/4a)^2}{2\beta'^2\gamma'^2 m}}}{1/4a}.$$

In particular,

- *for the choice of parameters made in Theorem 8.3, it is at least $1-e^{-n}$, i.e., the error is a negligible function.*
- *for the choice of parameters suggested in Equation 4, the probability is at least $1-e^{-5000}$.*

Proof. For correctness we only need to consider the case when all entries of $\mathbf{c}$ correspond to buckets of width less than $1/(4n^{3/2}\sqrt{m}\gamma')$. We write $\langle \mathbf{w}_i, \mathbf{z}\rangle = \langle \mathbf{w}_i, \mathbf{At}\rangle + \langle \mathbf{w}_i, \mathbf{z}-\mathbf{At}\rangle$ and bound each inner product separately. We start by bounding the first inner product.

For each $i \in \{1,\dots,\ell\}$, $\gamma'\sqrt{m}\sum_{j=1}^m e_{ij}t_j$ follows a $\mathcal{N}(0,\beta'^2\gamma'^2 m)$ distribution. By a union bound and 3.10,

$$\begin{aligned}\Pr\left[\forall i\colon \left|\gamma'\sqrt{m}\sum_{j=1}^m e_{ij}t_j\right| \le \frac{1}{4a}\right] =& 1 - \Pr\left[\exists i\colon \left|\gamma'\sqrt{m}\sum_{j=1}^m e_{ij}t_j\right| \le \frac{1}{4a}\right] \\ \ge& 1 - \ell\sqrt{\frac{2\beta'^2\gamma'^2 m}{\pi}}\frac{e^{-\frac{(1/4a)^2}{2\beta'^2\gamma'^2 m}}}{1/4a}\end{aligned}$$

By definition of the encryption scheme $\|\mathbf{z}-\mathbf{At}\| < \sqrt{n}/(4a\sqrt{n}\sqrt{m}\gamma')$, so

$$\left|\gamma'\sqrt{m}\langle \mathbf{w}_i, \mathbf{z}-\mathbf{At}\rangle\right| \le \|\mathbf{z}-\mathbf{At}\|\,\gamma'\sqrt{m} < \sqrt{n}\frac{1}{4a\sqrt{n}\sqrt{m}\gamma'}\gamma'\sqrt{m} \le \frac{1}{4a}.$$

Thus $\mathrm{Dec}(\mathbf{W}, \mathrm{Enc}(\mathbf{A},0)) = 0$. □

Claim 8.5. *If $n \ge 4$, the probability that $\mathrm{Dec}(\mathbf{w}, \mathrm{Enc}(\mathbf{A},1)) = 1$ is at least $1-(3/2a)^\ell - \exp(-\gamma'^2 m)$. In particular,*

- *for the choice of parameters made in Theorem 8.3, the probability is at least $1-(3/2n)^{\log n} - \exp(-n^3)$, i.e., the error is negligible.*
- *for the choice of parameters suggested in Equation 4, the probability is at least $1-(3/200)^\ell - \exp(-n^2)$.*

Proof. Encryptions of 1 can be seen as sampling $\mathbf{g} \leftarrow \mathcal{N}_n(\mathbf{0}, \mathbf{I}_n)$, rounding it and checking the width of its entries as described in the definition of the encryption scheme. For correctness we only need to consider the case when all entries correspond to buckets of width less than $1/(4a\sqrt{n}\sqrt{m}\gamma')$. To decrypt we take any $\mathbf{z}$ such that $\mathrm{round}_r(\mathbf{z}) = \mathrm{round}_r(\mathbf{g})$. We write $\langle \mathbf{w}_i, \mathbf{z}\rangle = \langle \mathbf{w}_i, \mathbf{g}\rangle + \langle \mathbf{w}_i, \mathbf{z}-\mathbf{g}\rangle$ and consider each inner product separately.

From the bound on the width of the buckets it follows that

$$\left|\gamma'\sqrt{m}\langle \mathbf{w}_i, \mathbf{z}-\mathbf{g}\rangle\right| \le \|\mathbf{z}-\mathbf{g}\|\,\gamma'\sqrt{m} < \sqrt{n}\frac{1}{4a\sqrt{n}\sqrt{m}\gamma'}\gamma'\sqrt{m} \le \frac{1}{4a}.$$

By the reverse triangle inequality, it follows that

$$\left|\left|\gamma'\sqrt{m}\langle \mathbf{w}_i, \mathbf{z}\rangle \mod 1\right| - \left|\gamma'\sqrt{m}\langle \mathbf{w}_i, \mathbf{g}\rangle \mod 1\right|\right| \le \left|\gamma'\sqrt{m}\langle \mathbf{w}_i, \mathbf{z}-\mathbf{g}\rangle \mod 1\right|.$$

Combining these two inequalities, we obtain that

$$-\frac{1}{4a} + \left|\gamma'\sqrt{m}\langle \mathbf{w}_i, \mathbf{g}\rangle \mod 1\right| < \left|\gamma'\sqrt{m}\langle \mathbf{w}_i, \mathbf{z}\rangle \mod 1\right|.$$

We now show that at least one entry satisfies $|\gamma'\sqrt{m}\langle \mathbf{w}_i, \mathbf{g}\rangle \mod 1| \geq 3/4a$ with probability at least $1 - (3/2a)^\ell - \exp(-\gamma'^2 m)$.

As $\mathbf{g}$ is a vector with distribution $\mathcal{N}_n(0,1)$, $\mathbf{W}^T\mathbf{g}$ has distribution $\mathcal{N}_\ell(0,1)$. By the smoothing property of Gaussians modulo $\mathbb{Z}^\ell$ (Lemmas 3.18 and 3.19), the statistical distance between $\gamma'\sqrt{m}\mathbf{W}^T\mathbf{g} \mod \mathbb{Z}^\ell$ and a uniform random variable on $(-1/2,1/2)^\ell$ is at most $\exp(-(\gamma'\sqrt{m}/\sqrt{\ell})^2\ell) = \exp(-\gamma'^2 m)$. This implies that the probability that at least one entry of $\gamma'\sqrt{m}\mathbf{W}^T\mathbf{g} \mod \mathbb{Z}^\ell$ does not belong to $(-3/4a, 3/4a)$ is

$$1 - \Pr[\gamma'\sqrt{m}\mathbf{W}^T\mathbf{g} \bmod \mathbb{Z}^\ell \in (-3/4a, 3/4a)^\ell] \geq 1 - \left(\frac{3}{2a}\right)^\ell - \exp(-\gamma'^2 m)$$

This shows that $\mathrm{Dec}(\mathbf{W}, \mathrm{Enc}(\mathbf{A}, 1)) = 1$ with probability at least $1 - (3/2a)^\ell - \exp(-\gamma'^2 m)$. □

8.4 Security

In order to analyze the security of the scheme we have to take into account the possibility that at least one of the entries of the ciphertext corresponds to a bucket of width larger than $1/(4a\sqrt{n}\sqrt{m}\gamma')$ as the encryption algorithm outputs the plaintext in the clear in that case.

Claim 8.6. *Let r be such that the following inequalities are satisfied*

$$r^{-3/5} \leq \frac{1}{4a\sqrt{n}\sqrt{m}\gamma'} \tag{5}$$

$$\frac{2nr^{-2/5}}{\sqrt{\ln r^{2/5}}} + \frac{2n}{r} \leq \delta(n). \tag{6}$$

Let $\mathbf{A} \in \mathbb{R}^{n\times m}$ be a matrix whose columns consist of independent hCLWE(ℓ) *samples and assume* hCLWE$(\ell)(s, \varepsilon)$ *where s is the complexity of rounding and ε is a function of n. Let $\mathbf{t} \leftarrow \{-1/\sqrt{m}, 1/\sqrt{m}\}^m$ be sampled uniformly at random. The probability that any entry of the vector* $\mathbf{c} := \mathrm{round}_r(\mathbf{A}\mathbf{t})$ *corresponds to a bucket of width larger than $1/(4a\sqrt{n}\sqrt{m}\gamma')$ is at most $\delta(n) + \varepsilon$. For the choice of parameters made in Theorem 8.3 and in Equation 4 both conditions are satisfied for $\delta(n) = \frac{1}{24}$.*

Proof. First consider a matrix $\mathbf{N}$ with i.i.d. entries from $\mathcal{N}(0,1)$. Since $\|\mathbf{t}\| = 1$ we get that $\mathbf{At}$ is a vector with i.i.d. entries in $\mathcal{N}(0,1)$. By Proposition 4.1 and condition 5 we know that the number of intervals of length larger than $1/(4a\sqrt{n}\sqrt{m}\gamma')$ is at most

$$\frac{2r^{3/5}}{\sqrt{\ln r^{2/5}}} + 2.$$

By a union bound and condition 6 the probability that any entry lands in such a bucket is at most

$$\frac{2nr^{-2/5}}{\sqrt{\ln r^{2/5}}} + \frac{2n}{r} \leq \delta(n).$$

The claim follows from the fact that the matrix $\mathbf{A}$ is ε-indistinguishable from $\mathbf{N}$.

For the choice of parameters made in Theorem 8.3, $r^{-3/5} = (40001n^3 \log n)^{-1}$, while $1/(4a\sqrt{n}\sqrt{m}\gamma') = 1/(40000n^2\sqrt{n}\log n(\sqrt{n} + O(1/n^6)))$. This proves that condition 5 holds. Condition 6 holds since $r^{-2/5} = (40001n^3 \log n)^{-2/3}$.

For the choice of parameters made in Equation 4, $r^{-3/5} = m^{-6}(\gamma')^{-1}$ and $r^{-2/5} = m^{-4}(\gamma')^{-2/3}$. This proves that condition 5 and condition 6 hold. □

The next claim follows directly from Proposition 4.8.

Claim 8.7. *If the ciphertexts are not the messages, the distributions* $(\mathbf{N}, \mathrm{Enc}(\mathbf{N}, 0))$ *and* $(\mathbf{N}, \mathrm{Enc}(\mathbf{N}, 1))$ *are* $\sqrt{4e \ln r^n/\sqrt{m}}$*-statistically close for a matrix* $\mathbf{N}$ *of independent standard Gaussians. In particular,*

- *for the choice of parameters made in Theorem 8.3, the distance is at most* $1/\sqrt{50} < 1/4$*.*
- *for the choice of parameters suggested in Equation 4, the distance is at most* $1/3$*.*

Corollary 8.8. *If* $\mathrm{hCLWE}(\ell)(s, \varepsilon)$ *holds, then the distributions* $(\mathbf{A}, \mathrm{Enc}(\mathbf{A}, 0))$ *and* $(\mathbf{A}, \mathrm{Enc}(\mathbf{A}, 1))$ *are* $(s - \mathrm{poly}(n), \sqrt{4e \ln r^n/\sqrt{m}} + 4\varepsilon)$*-indistinguishable where* $\mathbf{A}$ *is the public key matrix. In particular,*

- *for the choice of parameters made in Theorem 8.3, and* $\varepsilon = 1/24$*, we get* $1/4 + 4/24 < 1/2$*.*
- *for the choice of parameters suggested in Equation 4 and* $\varepsilon = 1/24$*, we get* $1/3 + 4/24 = 1/2$*.*

Proof. Let $\mathbf{N} \in \mathbb{R}^{n \times m}$ be a random standard normal matrix. By $\mathrm{hCLWE}(\ell)(s, \varepsilon)$, $(\mathbf{A}, \mathrm{Enc}(\mathbf{A}, b))$ and $(\mathbf{N}, \mathrm{Enc}(\mathbf{N}, b))$ are $(s - \mathrm{poly}(n), \varepsilon)$-indistinguishable for both $b = 0$ and $b = 1$. By Claim 8.7 and Claim 8.6 for $\delta(n) = \varepsilon$, $(\mathbf{N}, \mathrm{Enc}(\mathbf{N}, 0))$ and $(\mathbf{N}, \mathrm{Enc}(\mathbf{N}, 1))$ are $(\infty, \sqrt{4e \ln r^n/\sqrt{m}} + 2\varepsilon)$-indistinguishable. The result follows from the triangle inequality and the bound on the advantage that we get is $\sqrt{4e \ln r^n/\sqrt{m}} + 4\varepsilon$. □

9 hCLWE and hCLWE(ℓ) are in SZK

In this section we prove that hCLWE and hCLWE(ℓ) are in SZK, which is the class of decision problems that admit a statistical zero-knowledge proof [GMR89]. Zero-knowledge is defined with respect to honest verifiers.

We say that a sampling problem is in SZK if there is a polynomial-time honest-verifier statistical zero-knowledge protocol that accepts at least 2/3 of the YES instances and rejects at least 2/3 of the NO instances. The choice of threshold 2/3 is operational.

Our proof consists in a reduction from hCLWE to the statistical difference problem (SD). Sahai and Vadhan proved in [SV03] that SD is complete for SZK.

Definition 9.1 (SD Problem)**.** *The YES instances of the Statistical Difference (SD) problem are pairs of circuits* (C_0, C_1) *such that* $\Delta(C_0, C_1) > 2/3$ *and the NO instances are pairs of circuits* (C_0, C_1) *such that* $\Delta(C_0, C_1) < 1/3$*.*

Here Δ is the statistical (total variation) distance between the output distributions sampled by the circuits when instantiated with a uniformly random seed. That is, if the output space of C_0 and C_1 is some finite set Ω,

$$\Delta(C_0, C_1) = \sup_{A \subseteq \Omega} |\Pr[C_0 \in A] - \Pr[C_1 \in A]| = \frac{1}{2} \sum_{\omega \in \Omega} |\Pr[C_0 = \omega] - \Pr[C_1 = \omega]|$$

Since SD is a complete problem for the SZK class and SZK is a class closed under reductions (see [SV03]), we can study the SZK class by considering reductions to SD instead of interactive proof systems. This approach also removes any reference to zero-knowledge.

In order to show that hCLWE is in SZK, it suffices to define two circuits that satisfy the conditions of Definition 9.1.

Theorem 9.2. *Let K, K' be sufficiently large constants. If $\gamma' > 1$, $\beta'\gamma' \ln \gamma' < 1/(K' n \log n)$ and γ' is polynomially bounded,* $\mathrm{hCLWE}_{\beta,\gamma,n}$ *with $m = (Kn \log n \ln \gamma')^2$ samples is in* SZK.

Proof. Take K and r as in Equation 4, that is, $K = 4 \cdot 9 \cdot 10 \cdot e \cdot 2 \cdot 5$ and $r = m^{10}(\gamma')^{5/3}$. Let $K' = 4 \cdot 10^4 K$. Let $\mathbf{X}$ be either a valid public key $\mathbf{A} \in \mathbb{R}^{n \times m}$ or a matrix $\mathbf{N} \in \mathbb{R}^{n \times m}$ with i.i.d. entries sampled from $\mathcal{N}(0, 1)$. We define two circuits C_0, C_1 that take as input the pair $(\mathbf{t}, \mathbf{u})$ where $\mathbf{t} \in \{-1/\sqrt{m}, 1/\sqrt{m}\}^m$ and $\mathbf{u} \in \{1, 2, \ldots, r\}^n$. C_0 outputs $\mathrm{round}_r(\mathbf{Xt})$, i.e., an encryption of 0 using randomness t, while C_1 outputs $\mathbf{u}$, i.e., an encryption of 1 with randomness $\mathbf{u}$.

If $\mathbf{X} = \mathbf{A}$, by Claim 8.4 and Claim 8.5 and Claim 8.6 for $\epsilon(n) = 1/24 = \delta(n)$, the decryption error is at most $e^{-5000} + 3/200 + \exp(-n^2) + 1/24 + 1/24$. It follows that $\Delta(C_0, C_1) > 2/3$.

If $\mathbf{X} = \mathbf{N}$, then the statistical distance between C_0 and C_1 is at most $1/3$ by Proposition 4.8. □

We now prove an analogous statement for $\mathrm{hCLWE}(\ell)$.

Theorem 9.3. *Let K, K' be sufficiently large constants. If $\gamma' > 1$, $\beta'\gamma' \ln \gamma' < 1/(K' n \log n)$, γ' is polynomially bounded and $1 \leq \ell \leq n$,* $\mathrm{hCLWE}_{\beta,\gamma,n}(\ell)$ *with $m = (Kn \log n \ln \gamma')^2$ samples is in* SZK.

Proof. Take K and r as in Equation 4, that is, $K = 4 \cdot 9 \cdot 10 \cdot e \cdot 2 \cdot 5$ and $r = m^{10}(\gamma')^{5/3}$. Let $K' = 4 \cdot 10^4 K$. Let $\mathbf{X}$ be either a valid public key $\mathbf{A} \in \mathbb{R}^{n \times m}$ or a matrix $\mathbf{N} \in \mathbb{R}^{n \times m}$ with i.i.d. entries sampled from $\mathcal{N}(0, 1)$. We define two circuits C_0, C_1 that take as input the pair (t, u) where $\mathbf{t} \in \{-1/\sqrt{m}, 1/\sqrt{m}\}^m$ and $\mathbf{u} \in \{1, 2, \ldots, r\}^n$. C_0 outputs $\mathrm{round}_r(\mathbf{Xt})$, i.e., an encryption of 0 using randomness t, while C_1 outputs $\mathbf{u}$, i.e., an encryption of 1 with randomness $\mathbf{u}$.

If $\mathbf{X} = \mathbf{A}$, then the statistical distance between C_0 and C_1 is at least $2/3$. By Claim 8.4 and Claim 8.5 and Claim 8.6 for $\epsilon(n) = 1/24 = \delta(n)$, the decryption error is at most $e^{-5000} + (3/200)^\ell + \exp(-n^2) + 1/24 + 1/24$, so

$$\Delta(C_0, C_1) > 1 - e^{-5000} - \left(\frac{3}{200}\right)^\ell - \exp(-n^2) - \frac{1}{12} > \frac{2}{3}.$$

If $\mathbf{X} = \mathbf{N}$, then the statistical distance between C_0 and C_1 is at most $1/3$ by Claim 8.7. □

Acknowledgements

We are grateful to Devika Sharma and Luca Trevisan for their insight and advice.

This work was supported by the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (Grant agreement No. 101019547). The first author was additionally supported by RGC GRF CUHK14209920 and the fourth author was additionally supported by ISF grant No. 1399/17 and project PROMETHEUS (Grant 780701).

References

[AD97] Miklós Ajtai and Cynthia Dwork. A public-key cryptosystem with worst-case/average-case equivalence. In *Proceedings of the Twenty-Ninth Annual ACM Symposium on Theory of Computing*, STOC '97, page 284–293, New York, NY, USA, 1997. Association for Computing Machinery.

[BB20] Matthew S. Brennan and Guy Bresler. Reducibility and statistical-computational gaps from secret leakage. In Jacob D. Abernethy and Shivani Agarwal, editors, *Conference on Learning Theory, COLT 2020, 9-12 July 2020, Virtual Event [Graz, Austria]*, volume 125 of *Proceedings of Machine Learning Research*, pages 648–847. PMLR, 2020.

[BF10] Dan Boneh and David Mandell Freeman. Linearly homomorphic signatures over binary fields and new tools for lattice-based signatures. Cryptology ePrint Archive, Report 2010/453, 2010. https://ia.cr/2010/453.

[BF11] Dan Boneh and David Mandell Freeman. Linearly homomorphic signatures over binary fields and new tools for lattice-based signatures. In *Public Key Cryptography - PKC 2011 - 14th International Conference on Practice and Theory in Public Key Cryptography*, volume 6571 of *Lecture Notes in Computer Science*, page 1. Springer, 2011. Full version in [BF10].

[BR13] Quentin Berthet and Philippe Rigollet. Complexity theoretic lower bounds for sparse principal component detection. In Shai Shalev-Shwartz and Ingo Steinwart, editors, *Proceedings of the 26th Annual Conference on Learning Theory*, volume 30 of *Proceedings of Machine Learning Research*, pages 1046–1066, Princeton, NJ, USA, 12–14 Jun 2013. PMLR.

[BRST21] Joan Bruna, Oded Regev, Min Jae Song, and Yi Tang. Continuous lwe. In *Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing*, STOC 2021, pages 694–707, New York, NY, USA, 2021. Association for Computing Machinery.

[CS15] Jung Hee Cheon and Damien Stehlé. Fully Homomophic Encryption over the Integers Revisited. In *EUROCRYPT 2015 - 34th Annual International Conference on the Theory and Applications of Cryptographic Techniques*, pages 513–536, Sofia, Bulgaria, April 2015.

[DKS17] Ilias Diakonikolas, Daniel M. Kane, and Alistair Stewart. Statistical query lower bounds for robust estimation of high-dimensional gaussians and gaussian mixtures. *2017 IEEE 58th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 73–84, 2017.

[DNR04] Cynthia Dwork, Moni Naor, and Omer Reingold. Immunizing encryption schemes from decryption errors. In *EUROCRYPT*, 2004.

[Ede88] Alan Edelman. Eigenvalues and condition numbers of random matrices. *SIAM Journal on Matrix Analysis and Applications*, 9:543–560, 1988.

[GG98] Oded Goldreich and Shafi Goldwasser. On the limits of non-approximability of lattice problems. In *Proceedings of the Thirtieth Annual ACM Symposium on Theory of Computing*, STOC '98, page 1–9, New York, NY, USA, 1998. Association for Computing Machinery.

[GM84] Shafi Goldwasser and Silvio Micali. Probabilistic encryption. *J. Comput. Syst. Sci.*, 28(2):270–299, 1984.

[GMR89] Shafi Goldwasser, Silvio Micali, and Charles Rackoff. The knowledge complexity of interactive proof systems. *SIAM J. Comput.*, 18(1):186–208, 1989.

[Gor41] R. D. Gordon. Values of Mill's ratio of area to bounding ordinate of the normal probability integral for large values of the argument. *Annals of Math. Stat.*, 12:364–366, 1941.

[GPV08] Craig Gentry, Chris Peikert, and Vinod Vaikuntanathan. Trapdoors for hard lattices and new cryptographic constructions. In Cynthia Dwork, editor, *Proceedings of the 40th Annual ACM Symposium on Theory of Computing, Victoria, British Columbia, Canada, May 17-20, 2008*, pages 197–206. ACM, 2008.

[HR05] Thomas Holenstein and Renato Renner. One-way secret-key agreement and applications to circuit polarization and immunization of public-key encryption. In *Proceedings of the 25th Annual International Conference on Advances in Cryptology*, CRYPTO'05, page 478–493, Berlin, Heidelberg, 2005. Springer-Verlag.

[HWX15] Bruce Hajek, Yihong Wu, and Jiaming Xu. Computational lower bounds for community detection on random graphs. In *Proceedings of The 28th Conference on Learning Theory*, volume 40 of *Proceedings of Machine Learning Research*, pages 899–928, Paris, France, 03–06 Jul 2015. PMLR.

[IZ89] Russell Impagliazzo and David Zuckerman. How to recycle random bits. pages 248–253. IEEE, 1989.

[MR07] Daniele Micciancio and Oded Regev. Worst-case to average-case reductions based on gaussian measures. *SIAM J. Comput.*, 37:267–302, 2007.

[O'D14] Ryan O'Donnell. *Analysis of Boolean Functions.* Cambridge University Press, 2014.

[PRSD17] Chris Peikert, Oded Regev, and Noah Stephens-Davidowitz. Pseudorandomness of ring-lwe for any ring and modulus. In *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing*, STOC 2017, page 461–473, New York, NY, USA, 2017. Association for Computing Machinery.

[Reg05] O. Regev. On lattices, learning with errors, random linear codes, and cryptography. In *STOC*, pages 84–93, 2005. Full version in [Reg09].

[Reg09] O. Regev. On lattices, learning with errors, random linear codes, and cryptography. *J. ACM*, 56(6), 2009.

[SV03] Amit Sahai and Salil Vadhan. A complete problem for statistical zero knowledge. *J. ACM*, 50(2):196–249, mar 2003.