

Approximate Trapdoors for Lattices and Smaller Hash-and-Sign Signatures

Yilei Chen*

Nicholas Genise[†]

Pratyay Mukherjee[‡]

August 20, 2019

Abstract

We study a relaxed notion of lattice trapdoor called *approximate trapdoor*, which is defined to be able to invert Ajtai’s one-way function approximately instead of exactly. The primary motivation of our study is to improve the efficiency of the cryptosystems built from lattice trapdoors, including the hash-and-sign signatures.

Our main contribution is to construct an approximate trapdoor by modifying the gadget trapdoor proposed by Micciancio and Peikert [Eurocrypt 2012]. In particular, we show how to use the approximate gadget trapdoor to sample short preimages from a trapdoor-independent distribution. The analysis of the distribution uses a theorem (implicitly used in past works) regarding linear transformations of discrete Gaussians on lattices.

Our approximate gadget trapdoor can be used together with the existing optimization techniques to improve the concrete performance of the hash-and-sign signature in the random oracle model under (Ring-)LWE and (Ring-)SIS assumptions. Our implementation shows that the sizes of the public-key & signature are about half of those in schemes built from exact trapdoors.

*Visa Research.

[†]UCSD. Research conducted at Visa Research.

[‡]Visa Research.

Contents

1	Introduction	1
1.1	Summary of our contributions	1
1.2	Technical overview	2
1.3	Improvement in the efficiency compared to the exact trapdoor	4
2	Preliminaries	6
2.1	Linear Algebra	7
2.2	Lattices background	7
2.3	Gadgets, or G-Lattices	9
2.4	SIS, LWE, and the trapdoor	10
3	The Approximate Trapdoor for Ajtai’s Function	11
3.1	Hardness of the approximate ISIS problem	11
4	Approximate Gadget Trapdoor	13
4.1	Recall the G-trapdoor from MP12	13
4.2	The algorithms of the approximate G-trapdoor	14
4.3	Analyze the preimage and error distributions for a uniformly random target	15
5	Hash-and-Sign Signature Instantiated with the Approximate Trapdoor	18
5.1	Security analysis	19
5.2	Concrete parameters	20
A	The Ring Setting	26
A.1	Approximate Trapdoor Scheme in the Ring Setting	27
A.2	Analogous Lemmas and Theorem	29

1 Introduction

In the past two decades, lattice-based cryptography has emerged as one of the most active areas of research. It has enabled both advanced cryptographic capabilities, such as fully homomorphic encryption [Gen09]; and practical post-quantum secure public-key encryptions and signatures, as observed in the ongoing NIST post-quantum cryptography (PQC) standardization procedure [AAAS⁺19]. A large fraction of the lattice-based cryptosystems uses *lattice trapdoors*. Those cryptosystems include basic primitives like public-key encryption and signature schemes [GGH97, HPS98, HHP⁺03, GPV08], as well as advanced primitives such as identity-based encryption [GPV08, ABB10, CHKP12], attribute-based encryption [GVW13], and graded encodings [GGH15].

In this work, we focus on the trapdoor for the lattice-based one-way function defined by Ajtai [Ajt96], and its application in digital signatures [GPV08]. Given a wide, random matrix \mathbf{A} , and a target vector \mathbf{y} . The *inhomogeneous short integer solution* (ISIS) problem asks to find a short vector \mathbf{x} as a preimage of \mathbf{y} , i.e.

$$\mathbf{A} \cdot \mathbf{x} = \mathbf{y} \pmod{q}.$$

Without the trapdoor, finding a short preimage is proven to be as hard as solving certain lattice problems in the worst case [Ajt96]. A trapdoor for the matrix \mathbf{A} , on the other hand, allows its owner to efficiently produce a short preimage. An explicit construction of the trapdoor for Ajtai’s function was first given in [Ajt99] and later simplified by [AP11, MP12].

Towards the proper use of lattice trapdoors in cryptography, what really gives the trapdoor a punch is the work of Gentry, Peikert and Vaikuntanathan [GPV08]. They show how to sample a short preimage from a *trapdoor-independent distribution*, instead of a distribution which may leak information about the trapdoor (as observed by the attacks [GS02, NR06] on the initial attempts of building lattice-based signatures [GGH97, HHP⁺03]). The trapdoor-independent preimage sampling algorithm allows [GPV08] to securely build a hash-and-sign signature as follows. Let the matrix \mathbf{A} be the public verification key, the trapdoor of \mathbf{A} be the secret signing key. To sign a message m , first hash it to a vector \mathbf{y} , then use the trapdoor to sample a short preimage \mathbf{x} as the signature. The secret signing key is guaranteed to be hidden from the signatures, since the signatures are produced from a trapdoor-independent distribution.

Despite its elegant design, the hash-and-sign signature based on Ajtai’s function suffers from practical inefficiency due to its large key size and signature size. Indeed, all the three lattice-based signature candidates that enter the second round of NIST PQC standardization [AAAS⁺19] are built from two alternative approaches — Falcon [FHK⁺18] is based on the hash-and-sign paradigm over NTRU lattices; Dilithium [DKL⁺18] and qTESLA [ABB⁺19] are based on the rejection sampling approach [Lyu12, BG14]. The suggested parameters for the three candidates lead to competitive performance measures. For example, for 128-bit security, the sizes of the public keys & signatures for all the three candidates are below 5 kB & 4 kB (respectively). By contrast, for the hash-and-sign signature based on Ajtai’s function, the sizes of the public keys & signatures are more than 35 kB & 25 kB according to the implementation benchmarks of [BB13, BFRS18, GPR⁺18].

1.1 Summary of our contributions

In this paper we define a relaxed notion of lattice trapdoor called *approximate trapdoor*, which can be used to solve the ISIS problem *approximately* instead of exactly. The primary motivation is to improve the efficiency of hash-and-sign signatures based on Ajtai’s one-way function. With a relaxation of the correctness requirement, it is possible to generate smaller public matrices, trapdoors,

and preimages for Ajtai’s function, which translate to smaller public-keys, secret-keys, and signatures for the hash-and-sign signature scheme.

Our main technical contribution is to show that the gadget trapdoor proposed by Micciancio and Peikert [MP12] can be modified to an approximate trapdoor. In particular, we show how to use the approximate gadget trapdoor to sample preimages from a trapdoor-independent distribution. The analysis of the distribution uses a theorem (implicitly used in past works) regarding linear transformations of discrete Gaussians on lattices.

Our approximate gadget trapdoor can be used together with all existing optimization techniques, such as using the Hermite normal form and using a bigger base in the gadget, to improve the concrete performance of the hash-and-sign signature in the random oracle model under RingLWE and RingSIS assumptions. Our proof-of-concept implementation shows that the sizes of the public-key & signature can be reduced to 5 kB & 4.45 kB for an estimation of 88-bit security, and 11.25 kB & 9.38 kB for an estimation of 184-bit security. Those are much closer to the sizes of the signatures based on the rejection sampling approach [Lyu12, BG14, DKL⁺18, ABB⁺19]. More details of the comparisons are given in §1.3 and §5.2.

1.2 Technical overview

We now provide an overview of the problem setting and the techniques we use.

Given a public matrix $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$ where $m = O(n \log q)$, and a target \mathbf{y} . We call a vector $\mathbf{x} \in \mathbb{Z}^m$ an *approximate short preimage* of \mathbf{y} if

$$\mathbf{A} \cdot \mathbf{x} = \mathbf{y} + \mathbf{z} \pmod{q}$$

for some $\mathbf{z} \in \mathbb{Z}^n$, and both \mathbf{x} and \mathbf{z} are short. An *approximate trapdoor* for \mathbf{A} is defined to be a string that allows its owner to efficiently find an approximate short preimage given a target \mathbf{y} .

Of course, to make sense of the word “trapdoor”, we first need to argue that solving the approximate version of ISIS is hard without the trapdoor. Under proper settings of parameters, we show the approximate ISIS problem is as hard as the standard ISIS problem, or no easier than LWE. The reductions extensively use the Hermite normal form (HNF) and are pretty straightforward.

The approximate ISIS problem and the approximate trapdoor are natural generalizations of their exact variants. Indeed, both notions have been used in the literature, at least on an informal level. For example, the approximate ISIS problem was used in the work of Bai et al. [BGLS19] to improve the combinatorial algorithms of the exact ISIS problem.

For the approximate trapdoor, it is well-known that an exact trapdoor of a public matrix in the HNF, say a trapdoor for $\mathbf{A} = [\mathbf{I}_n \mid \mathbf{A}']$, can be used as an approximate trapdoor for \mathbf{A}' . Such a method was often used in the implementation of signatures to decrease the sizes of the public key and the signature by a dimension of n . Our goal is thus to further reduce the sizes compared to the HNF approach, while preserving the quality of the trapdoor, i.e. at least not increasing the norm of the preimage.

Approximate gadget trapdoor. Our main contribution is to show that the gadget trapdoor (G-trapdoor) proposed by Micciancio and Peikert [MP12] can be modified to an approximate trapdoor, in a way that further reduces the sizes of the public matrix, the trapdoor, and the preimage.

Recall the core of the G-trapdoor is a specific “gadget” matrix of base b ,

$$\mathbf{G} := \mathbf{I}_n \otimes \mathbf{g}^t := \mathbf{I}_n \otimes (1, b, \dots, b^{k-1}) \in \mathbb{Z}^{n \times (nk)},$$

where $k := \lceil \log_b q \rceil$. The base b is typically chosen to be 2 for simplicity, or a larger value in practical implementations.

Micciancio and Peikert [MP12] show how to generate a random matrix \mathbf{A} together with a matrix \mathbf{D} of small norm such that $\mathbf{A} \cdot \mathbf{D} = \mathbf{G} \pmod{q}$. In particular, \mathbf{A} is designed to be

$$\mathbf{A} = [\bar{\mathbf{A}} | \mathbf{G} - \bar{\mathbf{A}}\mathbf{R}],$$

where \mathbf{R} is a matrix with small entries and is the actual trapdoor. The matrix \mathbf{D} is then equal to $\begin{bmatrix} \mathbf{R} \\ \mathbf{I}_{nk} \end{bmatrix}$. Since the kernel of the \mathbf{G} matrix has a public short basis, one can first solve the ISIS problem under the public matrix \mathbf{G} , then use \mathbf{D} to solve the ISIS problem under the public matrix \mathbf{A} .

We observe that if we drop a few (say l) entries corresponding to the small powers of b from the gadget matrix \mathbf{G} , i.e. let the following \mathbf{F} matrix be a modified gadget matrix

$$\mathbf{F} := \mathbf{I}_n \otimes \mathbf{f}^t := \mathbf{I}_n \otimes (b^l, \dots, b^{k-1}) \in \mathbb{Z}^{n \times n(k-l)},$$

then we are still able to solve the ISIS problem w.r.t. the public matrix \mathbf{F} up to a b^l -approximation of the solution (i.e., the norm of the error vector is proportional to b^l). Replacing \mathbf{G} by \mathbf{F} in \mathbf{A} gives

$$\mathbf{A} = [\bar{\mathbf{A}} | \mathbf{F} - \bar{\mathbf{A}}\mathbf{R}]. \tag{1}$$

Then the dimensions of the trapdoor \mathbf{R} and the public matrix \mathbf{A} can be reduced.

Sampling from a trapdoor-independent distribution. Given a public matrix \mathbf{A} together with its approximate G-trapdoor \mathbf{R} , finding an arbitrary approximate short preimage of a given target \mathbf{u} is quite straightforward, but sampling the preimage from a trapdoor-independent distribution turns out to be non-trivial. As mentioned, the ability to sample from a trapdoor-independent distribution is fundamental to most of the trapdoor applications including digital signatures.

We provide an algorithm that samples an approximate short preimage from a trapdoor-independent distribution. The algorithm itself is a fairly simple generalization of the perturbation-based discrete Gaussian sampler from [MP12], but the analyses of the preimage distribution from [MP12] are not easy to generalize. Our analyses of the preimage distribution and the approximation error distribution extensively use a linear transformation theorem on lattice distributions (cf. Lemma 2.8, or Theorem 2.7, implicitly used in [MP12, MP13, BPMW16, DGPY19]).

The details of the analyses are quite technical. Here let us mention a qualitative difference of the result we have proved compared to the ones from [GPV08, MP12]. While our algorithm works for *all* the target image $\mathbf{u} \in \mathbb{Z}_q^n$ and does not cause any blow up in the standard deviation of the distribution, the analysis of trapdoor-independence only applies to a target image \mathbf{u} sampled *uniformly* from \mathbb{Z}_q^n , as oppose to the analysis for the exact trapdoor in [GPV08, MP12] which is able to spell out the distribution of the preimages of *all* the $\mathbf{u} \in \mathbb{Z}_q^n$. Still, sampling the approximate preimages for uniform targets from a trapdoor-independent distribution suffices for replacing the use of exact lattice trapdoors in the digital signature scheme.

To briefly explain the reason behind the qualitative difference, we observe that the methods we have tried to handle *all* the target images require significant increases in the smoothing parameters of the lattice intersections required in the linear transformation theorem (Lemma 2.8, or Theorem 2.7). In other words, the norm of the resulting preimage increases significantly rendering the result meaningless.

Parameters	Exact G-trapdoor	Approximate G-trapdoor
m	$n(2 + k)$	$n(2 + (k - l))$
σ	$\sqrt{b^2 + 1} \cdot \omega(\sqrt{\log n})$	$\sqrt{b^2 + 1} \cdot \omega(\sqrt{\log n})$
s	$C \cdot \tau \cdot (\sqrt{m} + 2\sqrt{n}) \cdot \sigma$	$C \cdot \tau \cdot (\sqrt{m} + 2\sqrt{n}) \cdot \sigma$
ν	0	$b^l \cdot \sigma$

Figure 1: A brief comparison of the parameters. The parameters in the table are derived under a fixed lattice dimension n , a fixed modulus $q \geq \sqrt{n}$, and a fixed base b . Let $k = \lceil \log_b q \rceil$. Let l denote the number of entries removed from \mathbf{g} ($1 \leq l < k$). Then we list m as the dimension of the public matrix and the preimage; σ as the width of the gadget preimage distribution; s as the width of the final preimage distribution (where $C > 0$ is a universal constant; τ as the width, or subgaussian parameter, of the distribution of the entries in the trapdoor); ν as the length bound of the error for each entry in the image.

1.3 Improvement in the efficiency compared to the exact trapdoor

We now explain the efficiency gain of using our approximate trapdoor compared to the exact trapdoor and the other existing optimization techniques, with a focus on the signature application. Our goal is to set the parameters to achieve the following “win-win-win” scenario:

1. Save on the preimage size (bandwidth).
2. Save on the size for the public matrix \mathbf{A} .
3. Retain, or even gain, concrete security, which is related to the discrete Gaussian width of the preimage and the norm of the error term.

Let us start with an understanding of the dependency of the savings on the variable l , i.e, the number of entries dropped from the gadget \mathbf{g} . In Figure 1 we provide a comparison of the parameters between the exact G-trapdoor of MP12 and the approximate G-trapdoor samplers in this paper. In both cases the public matrices are instantiated in the pseudorandom mode. For the approximate trapdoor, the dimension of the trapdoor decreases from nk to $n(k - l)$. The dimension m of the public matrix and the preimage decreases. The width s of the preimage distribution also decreases slightly following the decreasing of m . However, the norm of the error factor in the image grows with l . So in the concrete instantiation of the hash-and-sign signature discussed later, we need to coordinate the value of l with the norms of the preimage and the error, which will determine the security estimation together.

Our algorithm inherits the $O(\log q)$ -space, $O(n \log q)$ -time G-preimage sample subroutine from [MP12, GM18]. So the saving of space and time in the sampling of the perturbation is proportional to the saving in the dimension m .

Concrete parameters for the signatures. We give a proof-of-concept implementation of the hash-and-sign signature based on our approximate trapdoor. The security is analyzed in the random oracle model, assuming the hardness of RingLWE for the pseudorandomness of the public key and RingSIS for the unforgeability of the signature. Here we provide a short summary and leave more details in Section 5.2.

Params	Exact	Approx	Approx	Exact	Approx	Approx	Exact	Approx	Approx
n	512	512	512	512	512	512	512	512	512
$\lceil \log_2 q \rceil$	24	24	24	16	16	16	16	16	16
b	2	2	2	2	2	2	4	4	4
l	0	12	15	0	7	9	0	2	4
τ	40	40	40	2.6	2.6	2.6	2.6	2.6	2.6
s	38317.0	29615.3	26726.3	2170.7	1756.3	1618.2	3114.2	2833.3	2505.6
m	13312	7168	5632	9216	5632	4608	5120	4096	3072
$\ \mathbf{x}\ _2$	4441737.7	2521387.0	2035008.5	211100.9	133305.5	109339.1	223740.1	183004.9	138145.7
$\ \mathbf{z}\ _2$	0	374014.0	2118987.6	0	11897.9	46428.4	0	1402.3	19807.1
PK	37.50	19.50	15.00	17.00	10.00	8.00	9.00	7.00	5.00
Sig	25.68	13.53	10.51	13.16	7.83	6.30	7.62	5.94	4.45
LWE	100.0	100.0	100.0	104.7	104.7	104.7	104.7	104.7	104.7
AISIS	80.2	85.8	81.1	83.7	89.0	88.1	82.8	85.5	87.8

Figure 2: Summary of the concrete parameters. The size of PK and Sig are measured in kB. $\|\mathbf{x}\|_2$, $\|\mathbf{z}\|_2$ are the upper-bounds of the norms of the preimage and the error term. LWE and AISIS refer to the estimations of security levels for the pseudorandomness of the PK and finding a short approximate preimage.

Let us first remark that different implementation results of the hash-and-sign signatures [BB13, BFRS18, GPR⁺18] possibly use different ways of measuring sizes and security, and not all the details are recoverable from the papers. So we also include our reference implementation of the exact trapdoor as a fair comparison. For an estimation of 88-bit security, our reference implementation for the exact trapdoor under the modulus $q \approx 2^{24}$ and base $b = 2$ matches the parameters reported in [BB13].

We also use smaller moduli and bigger bases to reduce the size and increase the security level. The parameters in Figure 2 suggest that for the 3 choices of q and b , using the approximate gadget trapdoor by setting $l = \lceil (\log_b q)/2 \rceil$ saves about half of the sizes in the public key and signatures comparing to using the exact trapdoor, with even a slight increase in the security estimation.

Our implementation shows that the sizes of the public-key & signature can be reduced to 5 kB & 4.45 kB for an estimation of 88-bit security, and 11.25 kB & 9.38 kB for an estimation of 184-bit security. Those are much closer to the sizes of the signatures based on the rejection sampling approach [Lyu12, BG14, DKL⁺18, ABB⁺19]. As a reference, the sizes of the public-key & signature for qTESLA [ABB⁺19] are 4.03 kB & 3.05 kB for an estimation of 128-bit security, and 8.03 kB & 6.03 kB for an estimation of 192-bit security. The sizes for Dilithium [DKL⁺18] are even smaller. Let us remark that our implementation has not used many low-level optimizations like Dilithium [DKL⁺18] and qTESLA [ABB⁺19]. So it is reasonable to expect we have more room to improve after adding lower-level optimizations. The parameters for Falcon [FHK⁺18] are the smallest due to the use of NTRU lattices, so they are rather incomparable with the ones based on RingLWE. As a side note, we do not know how to construct approximate trapdoors for NTRU lattices, and we leave it as an interesting question to investigate in future.

Remarks on the comparison and compatibility with other possible optimizations. There are many folklore optimizations regarding trapdoors for Ajtai’s one-way function. We discuss the comparison and compatibility of a few of them with our construction. Throughout these comparisons we are concerned with the “win-win-win” scenario mentioned in the beginning.

First is the approximate trapdoor from the HNF optimization: $\mathbf{A} = [\mathbf{I}|\mathbf{A}']$. This barely achieves the

“win-win-win” scenario with a slight savings on the public key and the signature. Our construction can be used in the pseudorandom-mode of the gadget trapdoor which has automatically included the HNF optimization, and saves around 50% in addition.

Our method can also be used together with any base in the gadget, including a large base of size $b \approx \sqrt{q}$ (the resulting gadget is $\mathbf{g} = [1, \sqrt{q}]$), as was used in [dPLS18] when the modulus is large. This construction suffers from a large Gaussian width (\sqrt{q}), which hurts concrete security and is infeasible in the smaller modulus regime we implement in Section 5. Specifically for the smaller moduli, the signature’s Gaussian width is larger than the modulus, as was confirmed both on paper and in our experiments. So we use a moderately large base b .

One may also try to construct a short integer matrix \mathbf{S} for $\mathbf{A} := [(\mathbf{I}, \bar{\mathbf{A}}')|\mathbf{F} - (\mathbf{I}, \bar{\mathbf{A}}')\mathbf{R}]$ (corresponds to the pseudorandom public key in Eqn. (1)) such that $\mathbf{AS} = \mathbf{G}$, and hope this view provides a better approximate trapdoor. From here, the hash and sign signature scheme is to return $\mathbf{S}\mathbf{z} + \mathbf{p}$ where \mathbf{p} is a perturbation and \mathbf{z} is a G -lattice sample. However, such a matrix \mathbf{S} requires a b^l term. So this method does save on the public key, but does not improve the signature size and, most importantly, increases the Gaussian width by a factor of b^l . The increase of the width decreases the concrete security of the underlying SIS problem. In fact, to achieve the same saving in the public key, one can instead directly increase the base from b to b^l in the gadget.

Using approximate trapdoors in the advanced lattice cryptosystems. Finally, let us briefly mention the possible applications of the approximate trapdoors in the cryptosystems built from the dual-Regev approach [GPV08, ABB10, CHKP12, GVW13] and the GGH15 approach [GGH15, BVWW16, CC17, GKW17, WZ17, CVW18].

To use approximate trapdoors in the schemes based on the dual-Regev approach, we need to sample the LWE secret term with a small norm instead of from the uniform distribution to maintain the correctness of the schemes. The security of these schemes requires the corresponding extension of the Bonsai techniques in the approximate setting. We leave the details of the extension to future works.

For the schemes based on the GGH15-approach, the correctness of the schemes holds without any changes. The security also holds, except for the schemes in [CVW18] which requires the extension of the Bonsai techniques. Let us remark that the saving in the dimension m is of significant importance to the applications built on the GGH15 graded encoding scheme (implemented in [HHSS17, CGM⁺18]). In those applications, the modulus q is proportional to m^d (where $d \in \mathbb{N}$ is the number of “levels” of the graded encodings; larger d supports richer functionalities). So reducing the dimension m would dramatically reduce the overall parameter.

2 Preliminaries

Notations and terminology. In cryptography, the security parameter (denoted as λ) is a variable that is used to parameterize the computational complexity of the cryptographic algorithm or protocol, and the adversary’s probability of breaking security. An algorithm is “efficient” if it runs in (probabilistic) polynomial time over λ .

When a variable v is drawn uniformly random from the set S we denote as $v \leftarrow U(S)$. We use \approx_s and \approx_c as the abbreviations for statistically close and computationally indistinguishable. For two

distributions D_1, D_2 over the same support \mathcal{X} , we denote $D_1 \stackrel{\varepsilon}{\approx} D_2$ to denote that each $x \in \mathcal{X}$ has $D_1(x) \in [1 \pm \varepsilon]D_2(x)$ and $D_2(x) \in [1 \pm \varepsilon]D_1(x)$.

Let $\mathbb{R}, \mathbb{Z}, \mathbb{N}$ be the set of real numbers, integers and positive integers. Denote $\mathbb{Z}/q\mathbb{Z}$ by \mathbb{Z}_q . For $n \in \mathbb{N}$, $[n] := \{1, \dots, n\}$. A vector in \mathbb{R}^n (represented in column form by default) is written as a bold lower-case letter, e.g. \mathbf{v} . For a vector \mathbf{v} , the i^{th} component of \mathbf{v} will be denoted by v_i . For an integer base $b > 1$, we call a positive integer's "b-ary" decomposition the vector $(q_0, q_1, \dots, q_{k-1}) \in \{0, \dots, b-1\}^k$ where $k := \lceil \log_b q \rceil$, and $q = \sum q_i b^i$.

A matrix is written as a bold capital letter, e.g. \mathbf{A} . The i^{th} column vector of \mathbf{A} is denoted \mathbf{a}_i . The length of a vector is the ℓ_p -norm $\|\mathbf{v}\|_p := (\sum v_i^p)^{1/p}$, or the infinity norm given by its largest entry $\|\mathbf{v}\|_\infty := \max_i \{|v_i|\}$. The length of a matrix is the norm of its longest column: $\|\mathbf{A}\|_p := \max_i \|\mathbf{a}_i\|_p$. By default we use ℓ_2 -norm unless explicitly mentioned. When a vector or matrix is called "small" or "short", we refer to its norm but not its dimension, unless explicitly mentioned. The thresholds of "small" or "short" will be precisely parameterized in the article when necessary.

2.1 Linear Algebra

Let $\{\mathbf{e}_i\}_{i=1}^n$ be the canonical basis for \mathbb{R}^n , with entries $\delta(j, k)$ where $\delta(j, k) = 1$ when $j = k$ and 0 otherwise. For any set $S \subseteq \mathbb{R}^n$, its span (denoted as $\text{span}(S)$) is the smallest subspace of \mathbb{R}^n containing S . For a matrix, $\mathbf{M} \in \mathbb{R}^{n \times m}$, its span is the span of its column vectors, written as $\text{span}(\mathbf{M})$. We write matrix transpose as \mathbf{M}^t . Let $\tilde{\mathbf{B}}$ denote the Gram-Schmidt orthogonalization of \mathbf{B} . The GSO of an ordered basis $\mathbf{B} = [\mathbf{b}_1, \dots, \mathbf{b}_k]$ is assumed to be from left to right, $\tilde{\mathbf{b}}_1 = \mathbf{b}_1$, unless stated otherwise.

Recall \mathbf{M} 's singular value decomposition (SVD), i.e. $\mathbf{M} = \mathbf{VDW} \in \mathbb{R}^{n \times m}$ where $\mathbf{V} \in \mathbb{R}^{n \times n}$ along with $\mathbf{W} \in \mathbb{R}^{m \times m}$ are unitary, and $\mathbf{D} \in \mathbb{R}^{n \times m}$ is a triangular matrix containing \mathbf{M} 's singular values. Further, let $q = \min\{n, m\}$ and $\mathbf{D}_q = \text{diag}(s_1, \dots, s_q)$ be the diagonal matrix containing \mathbf{M} 's singular values $\{s_1 \geq \dots \geq s_q \geq 0\}$. Then, $\mathbf{D} = \mathbf{D}_q$ when $n = m$, $\mathbf{D} = [\mathbf{D}_q \mathbf{0}]$ when $m > n$, and $\mathbf{D} = \begin{bmatrix} \mathbf{D}_q \\ \mathbf{0} \end{bmatrix}$ in the case $m < n$.

A symmetric matrix $\Sigma \in \mathbb{R}^{n \times n}$ is *positive semi-definite* if for all $\mathbf{x} \in \mathbb{R}^n$, we have $\mathbf{x}^t \Sigma \mathbf{x} \geq 0$. It is *positive definite*, $\Sigma > 0$, if it is positive semi-definite and $\mathbf{x}^t \Sigma \mathbf{x} = 0$ implies $\mathbf{x} = \mathbf{0}$. We say $\Sigma_1 > \Sigma_2$ (\geq) if $\Sigma_1 - \Sigma_2$ is positive-(semi)definite. This forms a partial ordering on the set of positive semi-definite matrices, and we denote $\Sigma \geq \alpha \mathbf{I}$ often as $\Sigma \geq \alpha$ for constants $\alpha \in \mathbb{R}^+$. For any positive semi-definite matrix Σ , we write $\sqrt{\Sigma}$ to be any full rank matrix \mathbf{T} such that $\Sigma = \mathbf{T}\mathbf{T}^t$. We say \mathbf{T} is a *square root* of Σ . For two positive semi-definite matrices, Σ_1 and Σ_2 , we denote the positive semi-definite matrix formed by their block diagonal concatenation as $\Sigma_1 \oplus \Sigma_2$. Let \mathbf{M}^* denote Hermitian transpose. The (Moore-Penrose) *pseudoinverse* for matrix \mathbf{M} with SVD $\mathbf{M} = \mathbf{VDW}$ is $\mathbf{M}^+ = \mathbf{WD}^+ \mathbf{V}^*$ where \mathbf{D}^+ is given by transposing \mathbf{D} and inverting \mathbf{M} 's nonzero singular values. For example, $\mathbf{T} = s\mathbf{I}$ and $\mathbf{T}^+ = s^{-1}\mathbf{I}$ for a covariance $\Sigma = s^2\mathbf{I}$. (An analogous $\mathbf{T}^+ = \mathbf{T}^{-1}$ is given for the non-spherical, full-rank case $\Sigma > 0$ using Σ 's diagonalization.)

2.2 Lattices background

An n -dimensional lattice Λ of rank $k \leq n$ is a discrete additive subgroup of \mathbb{R}^n . Given k linearly independent basis vectors $\mathbf{B} = \{\mathbf{b}_1, \dots, \mathbf{b}_k \in \mathbb{R}^n\}$, the lattice generated by \mathbf{B} is

$$\Lambda(\mathbf{B}) = \Lambda(\mathbf{b}_1, \dots, \mathbf{b}_k) = \left\{ \sum_{i=1}^k x_i \cdot \mathbf{b}_i, x_i \in \mathbb{Z} \right\}.$$

Given $n, m \in \mathbb{N}$ and a modulus $q \geq 2$, we often use q -ary lattices and their cosets, denoted as

for $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$, denote $\Lambda^\perp(\mathbf{A})$ or $\Lambda_q^\perp(\mathbf{A})$ as $\{\mathbf{x} \in \mathbb{Z}^m : \mathbf{A} \cdot \mathbf{x} = \mathbf{0} \pmod{q}\}$;

for $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$, $\mathbf{w} \in \mathbb{Z}_q^n$, denote $\Lambda_{\mathbf{w}}^\perp(\mathbf{A})$ as $\{\mathbf{x} \in \mathbb{Z}^m : \mathbf{A} \cdot \mathbf{x} = \mathbf{w} \pmod{q}\}$.

Gaussians on lattices. For any $s > 0$ define the Gaussian function on \mathbb{R}^n with parameter s :

$$\forall \mathbf{x} \in \mathbb{R}^n, \rho_s(\mathbf{x}) = e^{-\pi \|\mathbf{x}\|^2 / s^2}.$$

For any $\mathbf{c} \in \mathbb{R}^n$, real $s > 0$, and n -dimensional lattice Λ , define the discrete Gaussian distribution $D_{\Lambda + \mathbf{c}, s}$ as:

$$\forall \mathbf{x} \in \Lambda + \mathbf{c}, D_{\Lambda + \mathbf{c}, s}(\mathbf{x}) = \frac{\rho_s(\mathbf{x})}{\rho_s(\Lambda + \mathbf{c})}.$$

The subscripts s and \mathbf{c} are taken to be 1 and $\mathbf{0}$ (respectively) when omitted.

For any positive semidefinite $\Sigma = \mathbf{T} \cdot \mathbf{T}^t$, define the non-spherical Gaussian function as

$$\forall \mathbf{x} \in \text{span}(\mathbf{T}) = \text{span}(\Sigma), \rho_{\mathbf{T}}(\mathbf{x}) = e^{-\pi \mathbf{x}^t \Sigma \mathbf{x}},$$

and $\rho_{\mathbf{T}}(\mathbf{x}) = 0$ for all $\mathbf{x} \notin \text{span}(\Sigma)$. Note that $\rho_{\mathbf{T}}(\cdot)$ only depends on Σ but not the specific choice of the \mathbf{T} , so we may write $\rho_{\mathbf{T}}(\cdot)$ as $\rho_{\sqrt{\Sigma}}(\cdot)$.

For any $\mathbf{c} \in \mathbb{R}^n$, any positive semidefinite Σ , and n -dimensional lattice Λ such that $(\Lambda + \mathbf{c}) \cap \text{span}(\Sigma)$ is non-empty, define the discrete Gaussian distribution $D_{\Lambda + \mathbf{c}, \sqrt{\Sigma}}$ as:

$$\forall \mathbf{x} \in \Lambda + \mathbf{c}, D_{\Lambda + \mathbf{c}, \sqrt{\Sigma}}(\mathbf{x}) = \frac{\rho_{\sqrt{\Sigma}}(\mathbf{x})}{\rho_{\sqrt{\Sigma}}(\Lambda + \mathbf{c})}.$$

Smoothing parameter. We recall the definition of smoothing parameter and some useful facts.

Definition 2.1 (Smoothing parameter [MR07]). For any lattice Λ and positive real $\epsilon > 0$, the smoothing parameter $\eta_\epsilon(\Lambda)$ is the smallest real $s > 0$ such that $\rho_{1/s}(\Lambda^* \setminus \{\mathbf{0}\}) \leq \epsilon$.

Notice that for two lattices of the same rank $\Lambda_1 \subseteq \Lambda_2$, the denser lattice always has the smaller smoothing parameter, i.e. $\eta_\epsilon(\Lambda_2) \leq \eta_\epsilon(\Lambda_1)$.

We will need a generalization of the smoothing parameter to the non-spherical Gaussian.

Definition 2.2. For a positive semi-definite $\Sigma = \mathbf{T} \mathbf{T}^t$, an $\epsilon > 0$, and a lattice Λ with $\text{span}(\Lambda) \subseteq \text{span}(\Sigma)$, we say $\eta_\epsilon(\Lambda) \leq \sqrt{\Sigma}$ if $\eta_\epsilon(\mathbf{T}^+ \Lambda) \leq 1$.

When the covariance matrix $\Sigma > 0$ and the lattice Λ are full-rank, $\sqrt{\Sigma} \geq \eta_\epsilon(\Lambda)$ is equivalent to the minimum eigenvalue of Σ , $\lambda_{\min}(\Sigma)$, being at least $\eta_\epsilon^2(\Lambda)$.

Lemma 2.3 (Smoothing parameter bound from [GPV08]). For any n -dimensional lattice $\Lambda(\mathbf{B})$ and for any $\omega(\sqrt{\log n})$ function, there is a negligible $\epsilon(n)$ for which

$$\eta_\epsilon(\Lambda) \leq \|\tilde{\mathbf{B}}\| \cdot \omega(\sqrt{\log n})$$

Lemma 2.4 ([MR07]). Let Λ be a lattice, $\mathbf{c} \in \text{span}(\Lambda)$. For any $\Sigma \geq 0$, if $\sqrt{\Sigma} \geq \eta_\epsilon(\Lambda)$ for some $\epsilon > 0$, then

$$\rho_{\sqrt{\Sigma}}(\Lambda + \mathbf{c}) \in \left[\frac{1 - \epsilon}{1 + \epsilon}, 1 \right] \cdot \rho_{\sqrt{\Sigma}}(\Lambda)$$

The following is a generalization of [GPV08, Corollary 2.8] for non-spherical Gaussian.

Corollary 2.5 (Smooth over the cosets). Let Λ, Λ' be n -dimensional lattices s.t. $\Lambda' \subseteq \Lambda$. Then for any $\epsilon > 0$, $\sqrt{\Sigma} \geq \eta_\epsilon(\Lambda')$, and $\mathbf{c} \in \text{span}(\Lambda)$, we have

$$\Delta(D_{\Lambda + \mathbf{c}, \sqrt{\Sigma}} \bmod \Lambda', U(\Lambda \bmod \Lambda')) < 2\epsilon$$

Lemma 2.6 ([PR06, MR07]). Let \mathbf{B} be a basis of an n -dimensional lattice Λ , and let $s \geq \|\tilde{\mathbf{B}}\| \cdot \omega(\log n)$, then $\Pr_{\mathbf{x} \leftarrow D_{\Lambda, s}}[\|\mathbf{x}\| \geq s \cdot \sqrt{n} \vee \mathbf{x} = \mathbf{0}] \leq \text{negl}(n)$.

Linear Transformations of Discrete Gaussians. We will use the following general theorem, implicitly used in [MP12, MP13, BPMW16], regarding the linear transformation, \mathbf{T} , of a discrete Gaussian. It states that as long as the original discrete Gaussian is smooth enough in the kernel of \mathbf{T} , then the distribution transformed by \mathbf{T} is statistically close to another discrete Gaussian.

Theorem 2.7 ([Mic]). For any positive definite Σ , vector \mathbf{c} , lattice coset $A := \Lambda + \mathbf{a} \subset \mathbf{c} + \text{span}(\Sigma)$, and linear transformation \mathbf{T} , if the lattice $\Lambda_{\mathbf{T}} = \Lambda \cap \ker(\mathbf{T})$ satisfies $\text{span}(\Lambda_{\mathbf{T}}) = \ker(\mathbf{T})$ and $\eta_\epsilon(\Lambda_{\mathbf{T}}) \leq \sqrt{\Sigma}$, then

$$\mathbf{T}(D_{A, \mathbf{c}, \sqrt{\Sigma}}) \stackrel{\bar{\epsilon}}{\approx} D_{\mathbf{T}A, \mathbf{T}\mathbf{c}, \mathbf{T}\sqrt{\Sigma}}$$

where $\bar{\epsilon} = 2\epsilon/(1 - \epsilon)$.

We remark that if \mathbf{T} is injective (i.e. $\ker(\mathbf{T})$ is trivial), then $\mathbf{T}(D_{A, \mathbf{c}, \sqrt{\Sigma}}) = D_{\mathbf{T}A, \mathbf{T}\mathbf{c}, \mathbf{T}\sqrt{\Sigma}}$.

Let us also remark that at the time of writing this article, the following lemma (which is a special case of Theorem 2.7) has already been proven in [DGPY19]. This lemma is suitable for all of our proofs using a non-injective linear transformation of a discrete gaussian.

In what follows, the max-log distance between two distributions with the same support S is $\Delta_{ML}(\mathcal{X}, \mathcal{Y}) = \max_{s \in S} |\log \mathcal{X}(s) - \log \mathcal{Y}(s)|$ [MW18].

Lemma 2.8 (Lemma 3, [DGPY19]). Let $\mathbf{T} \in \mathbb{Z}^{n \times m}$ such that $\mathbf{T}\mathbb{Z}^m = \mathbb{Z}^n$ and $\Lambda^\perp(\mathbf{T}) = \{\mathbf{x} \in \mathbb{Z}^m : \mathbf{T}\mathbf{x} = \mathbf{0} \in \mathbb{Z}^n\}$. Let $\Sigma = \mathbf{T}\mathbf{T}^t$. For $\epsilon \in (0, 1/2)$, $\hat{\epsilon} = \epsilon + O(\epsilon^2)$, $r \geq \eta_\epsilon(\Lambda^\perp(\mathbf{T}))$, the max-log distance between $\mathbf{T} \cdot D_{\mathbb{Z}^m, r}$ and $D_{\mathbb{Z}^n, r\sqrt{\Sigma}}$ is at most $4\hat{\epsilon}$.

2.3 Gadgets, or G-Lattices

Let $\mathbf{G} = \mathbf{I}_n \otimes \mathbf{g}^t \in \mathbb{Z}_q^{n \times nk}$ with $\mathbf{g}^t = (1, b, \dots, b^{k-1})$, $k = \lceil \log_b q \rceil$. \mathbf{G} is commonly referred to the gadget matrix. The gadget matrix's q -ary lattice, $\Lambda_q^\perp(\mathbf{G})$, is the direct sum of n copies of the lattice $\Lambda_q^\perp(\mathbf{g}^t)$. Further, $\Lambda_q^\perp(\mathbf{g}^t)$ has a simple basis,

$$\mathbf{B}_q = \begin{bmatrix} b & & & q_0 \\ -1 & \ddots & & \vdots \\ & \ddots & b & q_{k-2} \\ & & -1 & q_{k-1} \end{bmatrix}$$

where $(q_0, \dots, q_{k-1}) \in \{0, 1, \dots, b-1\}^k$ is the b -ary decomposition of the modulus, q . When $q = b^k$, we cheat by having $q_0 = q_1 = \dots = q_{k-2} = 0$ and $q_{k-1} = b$. Either way, the integer cosets of $\Lambda_q^\perp(\mathbf{g}^t)$ can be viewed as the syndromes of \mathbf{g}^t as a check matrix, in the terminology of coding theory. These cosets are expressed as $\Lambda_u^\perp(\mathbf{g}^t) = \{\mathbf{x} \in \mathbb{Z}^k : \mathbf{g}^t \mathbf{x} = u \pmod{q}\} = \Lambda_q^\perp(\mathbf{g}^t) + \mathbf{u}$ where \mathbf{u} can be any coset representative. A simple coset representative of $\Lambda_u^\perp(\mathbf{g}^t)$ is the b -ary decomposition of u . The integer cosets of $\Lambda_q^\perp(\mathbf{G})$ are expressed through the direct-sum construction, $\Lambda_{\mathbf{u}}^\perp(\mathbf{G}) = \Lambda_{u_1}^\perp(\mathbf{g}^t) \oplus \dots \oplus \Lambda_{u_n}^\perp(\mathbf{g}^t)$ where $\mathbf{u} = (u_1, \dots, u_n) \in \mathbb{Z}_q^n$. We call \mathbf{G} a gadget matrix since the following problems, SIS and LWE, are easily solved on the matrix \mathbf{G} [MP12].

2.4 SIS, LWE, and the trapdoor

We first recall the short integer solution (SIS) problem.

Definition 2.9 (SIS [Ajt96]). *For any $n, m, q \in \mathbb{Z}$ and $\beta \in \mathbb{R}$, define the short integer solution problem $\text{SIS}_{n,m,q,\beta}$ as follows: Given $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$, find a non-zero vector $\mathbf{x} \in \mathbb{Z}^m$ such that $\|\mathbf{x}\| \leq \beta$, and*

$$\mathbf{A}\mathbf{x} = \mathbf{0} \pmod{q}.$$

Definition 2.10 (ISIS). *For any $n, m, q \in \mathbb{Z}$ and $\beta \in \mathbb{R}$, define the inhomogeneous short integer solution problem $\text{ISIS}_{n,m,q,\beta}$ as follows: Given $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$, $\mathbf{y} \in \mathbb{Z}_q^n$, find $\mathbf{x} \in \mathbb{Z}^m$ such that $\|\mathbf{x}\| \leq \beta$, and*

$$\mathbf{A}\mathbf{x} = \mathbf{y} \pmod{q}.$$

Lemma 2.11 (Hardness of (I)SIS based on the lattice problems in the worst case [Ajt96, MR07, GPV08]). *For any $m = \text{poly}(n)$, any $\beta > 0$, and any sufficiently large $q \geq \beta \cdot \text{poly}(n)$, solving $\text{SIS}_{n,m,q,\beta}$ or $\text{ISIS}_{n,m,q,\beta}$ (where \mathbf{y} is sampled uniformly from \mathbb{Z}_q^n) with non-negligible probability is as hard as solving GapSVP_γ and SIVP_γ on arbitrary n -dimensional lattices with overwhelming probability, for some approximation factor $\gamma = \beta \cdot \text{poly}(n)$.*

All the (I)SIS problems and their variants admit the Hermite normal form (HNF), where the public matrix \mathbf{A} is of the form $[\mathbf{I}_n \mid \mathbf{A}']$ where $\mathbf{A}' \in \mathbb{Z}_q^{n \times (m-n)}$. The HNF variant of (I)SIS is as hard as the standard (I)SIS. This can be seen by rewriting $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$ as $\mathbf{A} =: [\mathbf{A}_1 \mid \mathbf{A}_2] = \mathbf{A}_1 \cdot [\mathbf{I}_n \mid \mathbf{A}_1^{-1} \cdot \mathbf{A}_2]$ (we always work with n, q such that $\mathbf{A}_1 \leftarrow U(\mathbb{Z}_q^{n \times n})$ is invertible with non-negligible probability).

Learning with errors. We recall the decisional learning with errors (LWE) problem.

Definition 2.12 (Decisional learning with errors [Reg09]). *For $n, m \in \mathbb{N}$ and modulus $q \geq 2$, distributions for secret vectors, public matrices, and error vectors $\theta, \pi, \chi \subseteq \mathbb{Z}_q$. An LWE sample is obtained from sampling $\mathbf{s} \leftarrow \theta^n$, $\mathbf{A} \leftarrow \pi^{n \times m}$, $\mathbf{e} \leftarrow \chi^m$, and outputting $(\mathbf{A}, \mathbf{y}^t := \mathbf{s}^t \mathbf{A} + \mathbf{e}^t \pmod{q})$.*

We say that an algorithm solves $\text{LWE}_{n,m,q,\theta,\pi,\chi}$ if it distinguishes the LWE sample from a random sample distributed as $\pi^{n \times m} \times U(\mathbb{Z}_q^m)$ with probability greater than $1/2$ plus non-negligible.

Lemma 2.13 (Hardness of LWE based on the lattice problems in the worst case [Reg09, Pei09, BLP⁺13, PRS17]). *Given $n \in \mathbb{N}$, for any $m = \text{poly}(n)$, $q \leq 2^{\text{poly}(n)}$. Let $\theta = \pi = U(\mathbb{Z}_q)$, $\chi = D_{\mathbb{Z},s}$ where $s \geq 2\sqrt{n}$. If there exists an efficient (possibly quantum) algorithm that breaks $\text{LWE}_{n,m,q,\theta,\pi,\chi}$, then there exists an efficient (possibly quantum) algorithm for solving GapSVP_γ and SIVP_γ on arbitrary n -dimensional lattices with overwhelming probability, for some approximation factor $\gamma = \tilde{O}(nq/s)$.*

The next lemma shows that LWE with the secret sampled from the error distribution is as hard as the standard LWE.

Lemma 2.14 ([ACPS09, BLP⁺13]). *For n, m, q, s chosen as was in Lemma 2.13, $\text{LWE}_{n,m',q,D_{\mathbb{Z},s},U(\mathbb{Z}_q),D_{\mathbb{Z},s}}$ is as hard as $\text{LWE}_{n,m,q,U(\mathbb{Z}_q),U(\mathbb{Z}_q),D_{\mathbb{Z},s}}$ for $m' \leq m - (16n + 4 \log \log q)$.*

Trapdoor. A *trapdoor* for a public matrix $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$ is a string that allows its owner to efficiently solve both the (I)SIS and LWE problems w.r.t. \mathbf{A} .

3 The Approximate Trapdoor for Ajtai's Function

Given a matrix $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$, define an *approximate trapdoor* of \mathbf{A} as anything that allows us to efficiently solve the approximate version of the ISIS problem w.r.t. \mathbf{A} . We first define the approximate ISIS problem.

Definition 3.1 (Approximate ISIS). *For any $n, m, q \in \mathbb{N}$ and $\alpha, \beta \in \mathbb{R}$, define the approximate inhomogeneous short integer solution problem $\text{Approx.ISIS}_{n,m,q,\alpha,\beta}$ as follows: Given $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$, $\mathbf{y} \in \mathbb{Z}_q^n$, find a vector $\mathbf{x} \in \mathbb{Z}^m$ such that $\|\mathbf{x}\| \leq \beta$, and there is a vector $\mathbf{z} \in \mathbb{Z}^n$ satisfying*

$$\|\mathbf{z}\| \leq \alpha \quad \text{and} \quad \mathbf{A}\mathbf{x} = \mathbf{y} + \mathbf{z} \pmod{q}.$$

Let us remark that the approximate ISIS is only non-trivial when the bounds α, β are relatively small compared to the modulus q . Also, our definition chooses to allow the zero vector to be a valid solution, which means when $\|\mathbf{y}\| \leq \alpha$, the zero vector is trivially a solution. Such a choice in the definition does not cause a problem in the application, since the interesting case in the application is to handle all the $\mathbf{y} \in \mathbb{Z}_q^n$, or \mathbf{y} sampled uniformly random from \mathbb{Z}_q^n .

Definition 3.2 (Approximate trapdoor). *A string τ is called an (α, β) -approximate trapdoor for a matrix $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$ if there is a probabilistic polynomial time algorithm (in $n, m, \log q$) that given τ, \mathbf{A} and any $\mathbf{y} \in \mathbb{Z}_q^n$, outputs a non-zero vector $\mathbf{x} \in \mathbb{Z}^m$ such that $\|\mathbf{x}\| \leq \beta$, and there is a vector $\mathbf{z} \in \mathbb{Z}^n$ satisfying*

$$\|\mathbf{z}\| \leq \alpha \quad \text{and} \quad \mathbf{A}\mathbf{x} = \mathbf{y} + \mathbf{z} \pmod{q}.$$

3.1 Hardness of the approximate ISIS problem

To make sense of the approximate trapdoor, we argue that for those who do not have the trapdoor, the approximate ISIS problem is a candidate one-way function under proper settings of parameters.

First, we observe a rather obvious reduction that bases the hardness of solving approximate ISIS (given an arbitrary target) on the hardness of decisional LWE with low-norm secret (e.g. when the secret is sampled from the error distribution). In the theorem statement below, when the norm symbol is applied on a distribution D , i.e. $\|D\|$, it denotes the lowest value $v \in \mathbb{R}^+$ such that $\Pr_{d \leftarrow D}[\|d\| < v] > 1 - \text{negl}(\lambda)$.

Theorem 3.3. *For $n, m, q \in \mathbb{Z}$, $\alpha, \beta \in \mathbb{R}^+$, θ, χ be distributions over \mathbb{Z} such that $q > 4(\|\theta\| \cdot (\alpha + 1) + \|\theta^n\| \cdot \alpha \cdot \sqrt{n} + \|\chi^m\| \cdot \beta \cdot \sqrt{m})$. Then $\text{LWE}_{n,m,q,\theta,U(\mathbb{Z}_q),\chi} \leq_p \text{Approx.ISIS}_{n,m,q,\alpha,\beta}$.*

Proof. Suppose there is a polynomial time adversary A that breaks $\text{Approx.ISIS}_{n,m,q,\alpha,\beta}$, we build a polynomial time adversary B that breaks decisional LWE.

Let $r = \lfloor \alpha \rfloor + 1$. Given an LWE challenge $(\mathbf{A}, \mathbf{w}) \in \mathbb{Z}_q^{n \times m} \times \mathbb{Z}_q^m$, where \mathbf{w} is either an LWE sample or sampled uniformly from \mathbb{Z}_q^m . B picks a vector $\mathbf{y} := (r, 0, \dots, 0)^t \in \mathbb{Z}_q^n$, sends \mathbf{A} and \mathbf{y} to the adversary A as an approximate ISIS challenge. A replies with $\mathbf{x} \in \mathbb{Z}^m$ such that $\|\mathbf{x}\| \leq \beta$, and there is a vector $\mathbf{z} \in \mathbb{Z}^n$ satisfying

$$\|\mathbf{z}\| \leq \alpha \quad \text{and} \quad \mathbf{A}\mathbf{x} = \mathbf{y} + \mathbf{z} \pmod{q}.$$

Note that $\mathbf{x} \neq \mathbf{0}$ since $\|\mathbf{y}\| > \alpha$.

B then computes $v := \langle \mathbf{w}, \mathbf{x} \rangle$. If $\mathbf{w}^t = \mathbf{s}^t \mathbf{A} + \mathbf{e}^t$ for $\mathbf{s} \leftarrow \theta^n$, $\mathbf{e} \leftarrow \chi^m$, then

$$v = (\mathbf{s}^t \mathbf{A} + \mathbf{e}^t) \mathbf{x} = \mathbf{s}^t (\mathbf{y} + \mathbf{z}) + \mathbf{e}^t \mathbf{x} \Rightarrow \|v\| \leq \|\theta\| \cdot r + \|\theta^n\| \cdot \alpha \cdot \sqrt{n} + \|\chi^m\| \cdot \beta \cdot \sqrt{m} < q/4.$$

Otherwise v distributes uniformly random over \mathbb{Z}_q . So B can compare v with the threshold value and wins the decisional LWE challenge with probability $1/2$ plus non-negligible. \square

Alternatively, we can also prove that the approximate ISIS problem is as hard as the standard ISIS. The reductions go through the HNFs of the ISIS and the approximate ISIS problems. All the reductions in the following theorem works for uniformly random target vectors.

Theorem 3.4. $\text{ISIS}_{n,n+m,q,\beta} \geq_p \text{Approx.ISIS}_{n,m,q,\alpha+\beta,\beta}$; $\text{ISIS}_{n,n+m,q,\alpha+\beta} \leq_p \text{Approx.ISIS}_{n,m,q,\alpha,\beta}$.

Proof. We will show $\text{ISIS} = \text{HNF.ISIS} = \text{HNF.Approx.ISIS} = \text{Approx.ISIS}$ under proper settings of parameters.

Recall that $\text{ISIS}_{n,m,q,\beta} = \text{HNF.ISIS}_{n,m,q,\beta}$ as explained in the preliminary. Also, $\text{HNF.ISIS}_{n,m,q,\beta} \geq_p \text{HNF.Approx.ISIS}_{n,m,q,\alpha,\beta}$ for any $\alpha \geq 0$ by definition. It remains to show the rest of the connections.

Lemma 3.5. $\text{HNF.ISIS}_{n,m,q,\alpha+\beta} \leq_p \text{HNF.Approx.ISIS}_{n,m,q,\alpha,\beta}$.

Proof. Suppose there is a polynomial time algorithm A that solves $\text{HNF.Approx.ISIS}_{n,m,q,\alpha,\beta}$, we build a polynomial time algorithm B that solves $\text{HNF.ISIS}_{n,m,q,\alpha+\beta}$. Given an HNF.ISIS instance $[\mathbf{I}_n \mid \mathbf{A}] \in \mathbb{Z}_q^{n \times m}$, \mathbf{y} , B passes the same instance to A , gets back a vector \mathbf{x} such that

$$[\mathbf{I}_n \mid \mathbf{A}] \cdot \mathbf{x} = \mathbf{y} + \mathbf{z} \pmod{q},$$

where $\|\mathbf{x}\| \leq \beta$, $\|\mathbf{z}\| \leq \alpha$. Now write $\mathbf{x} =: [\mathbf{x}_1^t \mid \mathbf{x}_2^t]^t$ where $\mathbf{x}_1 \in \mathbb{Z}^n$, $\mathbf{x}_2 \in \mathbb{Z}^m$. Then $\mathbf{x}' := [(\mathbf{x}_1 - \mathbf{z})^t \mid \mathbf{x}_2^t]^t$ satisfies

$$[\mathbf{I}_n \mid \mathbf{A}] \cdot \mathbf{x}' = \mathbf{y} \pmod{q},$$

and $\|\mathbf{x}'\| \leq \alpha + \beta$. So \mathbf{x}' is a valid solution to HNF.ISIS . \square

Lemma 3.6. $\text{HNF.Approx.ISIS}_{n,n+m,q,\alpha,\beta} \leq_p \text{Approx.ISIS}_{n,m,q,\alpha,\beta}$.

Proof. Suppose there is a polynomial time algorithm A that solves $\text{Approx.ISIS}_{n,m,q,\alpha,\beta}$, we build a polynomial time algorithm B that solves $\text{HNF.Approx.ISIS}_{n,n+m,q,\alpha,\beta}$. Given $[\mathbf{I}_n \mid \mathbf{A}] \in \mathbb{Z}_q^{n \times (n+m)}$, $\mathbf{y} \in \mathbb{Z}_q^n$ as an HNF.Approx.ISIS instance, B passes $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$, \mathbf{y} to A , gets back a short vector $\mathbf{x} \in \mathbb{Z}^m$. Then $[\mathbf{0}_n^t \mid \mathbf{x}^t]^t$ is a valid solution to the HNF.Approx.ISIS instance. \square

Lemma 3.7. $\text{HNF.Approx.ISIS}_{n,n+m,q,\alpha,\beta} \geq_p \text{Approx.ISIS}_{n,m,q,\alpha+\beta,\beta}$.

Proof. Suppose there is a polynomial time algorithm A that solves $\text{HNF.Approx.ISIS}_{n,n+m,q,\alpha,\beta}$, we build a polynomial time algorithm B that solves $\text{Approx.ISIS}_{n,m,q,\alpha+\beta,\beta}$. Given an Approx.ISIS instance $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$, $\mathbf{y} \in \mathbb{Z}^n$, B passes $[\mathbf{I}_n \mid \mathbf{A}] \in \mathbb{Z}_q^{n \times (n+m)}$, \mathbf{y} as an HNF.Approx.ISIS instance to A , gets back an answer $\mathbf{x} \in \mathbb{Z}^{m+n}$ such that

$$[\mathbf{I}_n \mid \mathbf{A}] \cdot \mathbf{x} = \mathbf{y} + \mathbf{z} \pmod{q}, \quad (2)$$

where $\|\mathbf{x}\| \leq \beta$, $\|\mathbf{z}\| \leq \alpha$.

Now write $\mathbf{x} =: [\mathbf{x}_1^t \mid \mathbf{x}_2^t]^t$ where $\mathbf{x}_1 \in \mathbb{Z}^n$, $\mathbf{x}_2 \in \mathbb{Z}^m$. Rewriting Eqn. (2) gives

$$\mathbf{A} \cdot \mathbf{x}_2 = \mathbf{y} + \mathbf{z} - \mathbf{x}_1 \pmod{q},$$

so \mathbf{x}_2 is a valid solution to $\text{Approx.ISIS}_{n,m,q,\alpha+\beta,\beta}$. □

Theorem 3.4 then follows the lemmas above. □

The following statement immediately follows the proof of Lemma 3.7.

Corollary 3.8. *An (α, β) -approximate trapdoor for $[\mathbf{I} \mid \mathbf{A}]$ is an $(\alpha + \beta, \beta)$ -approximate trapdoor for \mathbf{A} .*

4 Approximate Gadget Trapdoor

We present an instantiation of an approximate trapdoor based on the gadget-based trapdoor generation and preimage sampling algorithms of Micciancio and Peikert [MP12] (without the tag matrices). In short, we show how to generate a pseudorandom \mathbf{A} with entries modulo q along with an approximate trapdoor \mathbf{R} with small integer entries.

In the rest of this section, we first recall the exact G-trapdoor from [MP12], then present the approximate trapdoor generation algorithm and the approximate preimage sampling algorithm. Finally we analyze the preimage distribution. The analyses make extensive use of Lemma 2.8, or Theorem 2.7 (linear transformations of discrete Gaussians).

4.1 Recall the G-trapdoor from MP12

Let $b \geq 2$ be the base for the G-lattice. Let q be the modulus, $k = \lceil \log_b q \rceil$. b is typically chosen to be 2 for simplicity, but often a higher base b is used for efficiency trade-offs in lattice-based schemes.

Recall the gadget-lattice trapdoor technique from [MP12]: the public matrix is

$$\mathbf{A} = [\bar{\mathbf{A}} \mid \mathbf{G} - \bar{\mathbf{A}}\mathbf{R}]$$

where \mathbf{G} is the commonly used gadget matrix, $\mathbf{G} := \mathbf{I}_n \otimes \mathbf{g}_k^t$, $\mathbf{g}_k^t := (1, b, \dots, b^{k-1})$, and \mathbf{R} is a secret, trapdoor matrix with small, random entries. \mathbf{A} is either statistically close to uniformly random or pseudorandom, depending on the structure of $\bar{\mathbf{A}}$ and the choice of χ (in the pseudorandom case $\chi \subseteq \mathbb{Z}$ is chosen to be a distribution such that $\text{LWE}_{n,n,q,\chi,U(\mathbb{Z}_q),\chi}$ is hard). In this paper we focus on the pseudorandom case since the resulting public matrix \mathbf{A} and preimage have smaller dimensions.

In order to sample a short element in $\Lambda_{\mathbf{u}}^\perp(\mathbf{A})$, we use the trapdoor to map short coset representatives of $\Lambda_q^\perp(\mathbf{G})$ to short coset representatives of $\Lambda_q^\perp(\mathbf{A})$ by the relation

$$\mathbf{A} \begin{bmatrix} \mathbf{R} \\ \mathbf{I} \end{bmatrix} = \mathbf{G}.$$

Using the trapdoor as a linear transformation alone leaks information about the trapdoor. Therefore, we perturb the sample to statistically hide the trapdoor. Let Σ_p be a positive definite matrix defined as $\Sigma_p := s^2\mathbf{I} - \sigma^2 \begin{bmatrix} \mathbf{R}\mathbf{R}^t & \mathbf{R}^t \\ \mathbf{R} & \mathbf{I} \end{bmatrix}$ where σ is at least $\eta_\varepsilon(\Lambda_q^\perp(\mathbf{G}))$. The perturbation can be computed offline as $\mathbf{p} \leftarrow D_{\mathbb{Z}^m, \sqrt{\Sigma_p}}$. We then sample a G-lattice vector in a coset dependent on \mathbf{p} as $\mathbf{z} \leftarrow D_{\Lambda_q^\perp(\mathbf{G}), \sigma}$ and $\mathbf{v} = \mathbf{u} - \mathbf{A}\mathbf{p} \in \mathbb{Z}_q^n$. Finally, the preimage is set to be

$$\mathbf{y} := \mathbf{p} + \begin{bmatrix} \mathbf{R} \\ \mathbf{I} \end{bmatrix} \mathbf{z}.$$

4.2 The algorithms of the approximate G-trapdoor

As mentioned in the introduction, the main idea of obtaining an approximate trapdoor is to adapt the MP12 algorithms with a gadget matrix without the lower-order entries. Let $0 < l < k$ be the number of lower-order entries dropped from the gadget vector $\mathbf{g} \in \mathbb{Z}_q^k$. Define the resulting approximate gadget vector as $\mathbf{f} := (b^l, b^{l+1}, \dots, b^{k-1})^t \in \mathbb{Z}_q^{(k-l)}$. Let $w = n(k-l)$ be the number of columns of the approximate gadget $\mathbf{F} := \mathbf{I}_n \otimes \mathbf{f}^t \in \mathbb{Z}^{n \times w}$. Then the number of columns of \mathbf{A} will be $m := 2n + w$.

Once we replace the gadget matrix \mathbf{G} with its truncated version, \mathbf{F} , our approximate trapdoor generation and approximate preimage sampling algorithms match the original gadget-based algorithms. The generation and preimage algorithms are given as Algorithms 2 and 3, respectively. Algorithm 1 represents our approximate \mathbf{F} -sampling algorithm. It simply runs the G-lattice preimage sampling algorithm and drops the first l entries from the preimage. The covariance of the perturbation in Algorithm 3 is chosen as

$$\Sigma_p := s^2\mathbf{I}_m - \sigma^2 \begin{bmatrix} \mathbf{R}\mathbf{R}^t & \mathbf{R} \\ \mathbf{R}^t & \mathbf{I} \end{bmatrix}.$$

The results of this section are summarized in the following theorem.

Theorem 4.1. *There exists probabilistic, polynomial time algorithms $\text{APPROX.TRAPGEN}(\cdot)$ and $\text{APPROX.SAMPLEPRE}(\cdot, \cdot, \cdot, \cdot)$ satisfying the following.*

1. $\text{APPROX.TRAPGEN}(n)$ takes as input a security parameter n and returns a matrix-approximate trapdoor pair $(\mathbf{A}, \mathbf{R}) \in \mathbb{Z}_q^{n \times m} \times \mathbb{Z}^{2n \times n(k-l)}$.
2. Let \mathbf{A} be generated with an approximate trapdoor as above and let $\text{APPROX.A}^{-1}(\cdot)$ denote the approximate preimage sampling algorithm, $\text{APPROX.SAMPLEPRE}(\mathbf{A}, \mathbf{R}, s, \cdot)$. The following two distributions are statistically indistinguishable:

$$\{(\mathbf{A}, \mathbf{y}, \mathbf{u}, \mathbf{e}) : \mathbf{u} \leftarrow U(\mathbb{Z}_q^n), \mathbf{y} \leftarrow \text{APPROX.A}^{-1}(\mathbf{u}), \mathbf{e} = \mathbf{u} - \mathbf{A}\mathbf{y} \pmod{q}\}$$

<hr/> Algorithm 1: GSAMP.CUT(v, σ) <hr/> Input: $v \in \mathbb{Z}_q, \sigma \in \mathbb{R}^+$ Output: $\mathbf{z} \in \mathbb{Z}^{k-l}$ <ol style="list-style-type: none"> 1 Sample $\mathbf{x} \in \mathbb{Z}^k$ from $D_{\Lambda_v^\perp(\mathbf{g}^t), \sigma}$ 2 Let \mathbf{z} be the last $k - l$ entries of \mathbf{x} 3 return \mathbf{z}. <hr/>	<hr/> Algorithm 3: APPROX.SAMPLEPRE($\mathbf{A}, \mathbf{R}, \mathbf{u}, s$). <hr/> Input: $(\mathbf{A}, \mathbf{R}, \mathbf{u}, s)$ as in Theorem 4.1. Output: An approximate preimage of \mathbf{u} for $\mathbf{A}, \mathbf{y} \in \mathbb{Z}^m$. <ol style="list-style-type: none"> 1 Sample a perturbation $\mathbf{p} \leftarrow D_{\mathbb{Z}^m, \sqrt{\Sigma_p}}$. 2 Form $\mathbf{v} = \mathbf{u} - \mathbf{A}\mathbf{p} \in \mathbb{Z}_q^n$. 3 Sample the approximate gadget preimage $\mathbf{z} \in \mathbb{Z}^{n(k-l)}$ as $\mathbf{z} \leftarrow \text{GSAMP.CUT}(\mathbf{v}, \sigma)$. 4 Form $\mathbf{y} := \mathbf{p} + \begin{bmatrix} \mathbf{R} \\ \mathbf{I} \end{bmatrix} \mathbf{z} \in \mathbb{Z}^m$. 5 return \mathbf{y}. <hr/>
<hr/> Algorithm 2: APPROX.TRAPGEN(λ) $_\chi$ <hr/> Input: Security parameter λ Output: matrix-approximate trapdoor pair (\mathbf{A}, \mathbf{R}) . <ol style="list-style-type: none"> 1 Sample a uniformly random $\hat{\mathbf{A}} \leftarrow U(\mathbb{Z}_q^{n \times n})$. 2 Let $\bar{\mathbf{A}} := [\mathbf{I}_n, \hat{\mathbf{A}}]$. 3 Sample the approximate trapdoor $\mathbf{R} \leftarrow \chi^{2n \times w}$. 4 Form $\mathbf{A} := [\bar{\mathbf{A}} \mathbf{F} - \bar{\mathbf{A}}\mathbf{R}] \in \mathbb{Z}_q^{n \times m}$. return (\mathbf{A}, \mathbf{R}). <hr/>	

Figure 3: Pseudocode for the approximate trapdoor sampling algorithm in Subsection 4.3. We abuse notation and let $\text{GSAMP.CUT}(\mathbf{v}, \sigma)$ denote n independent calls to Algorithm 1 on each entries of $\mathbf{v} \in \mathbb{Z}_q^n$, and then concatenate the output vectors. The distribution $\chi \subseteq \mathbb{Z}$ is chosen so that $\text{LWE}_{n,n,q,\chi,U(\mathbb{Z}_q),\chi}$ is hard.

and

$$\{(\mathbf{A}, \mathbf{y}, \mathbf{u}, \mathbf{e}) : \mathbf{y} \leftarrow D_{\mathbb{Z}^m, s}, \mathbf{e} \leftarrow D_{\mathbb{Z}^n, \sigma \sqrt{(b^{2l}-1)/(b^2-1)}} \pmod{q}, \mathbf{u} = \mathbf{A}\mathbf{y} + \mathbf{e} \pmod{q}\}$$

for any $\sigma \geq \sqrt{b^2 + 1} \cdot \omega(\sqrt{\log n})$ and $s \geq \sqrt{\sigma^2 s_1^2(\mathbf{R}) + \omega(\sqrt{\log n})}$. Furthermore, in the second distribution, \mathbf{A} is computationally indistinguishable from random assuming $\text{LWE}_{n,n,q,\chi,U(\mathbb{Z}_q),\chi}$.

Let us remark that the preimage sampling algorithm works for *any* target $\mathbf{u} \in \mathbb{Z}_q^n$, but we are only able to prove the approximate preimage sample-coset pair (\mathbf{y}, \mathbf{u}) hides the approximate trapdoor, \mathbf{R} , over a *uniform* target $\mathbf{u} \in \mathbb{Z}_q^n$. This is unlike the exact gadget-based trapdoor setting in [MP12] which proves the trapdoor is hidden for each fixed \mathbf{u} . In the overview of the proof idea, we will briefly explain where the proof breaks down when we try to spell out the preimage distributions of every $\mathbf{u} \in \mathbb{Z}_q^n$.

4.3 Analyze the preimage and error distributions for a uniformly random target

This subsection is dedicated to proving Theorem 4.1. For the convenience of explanation, in this subsection we redefine the gadget \mathbf{G} by permuting the columns so that the columns of smaller entries are all on the left, i.e.

$$\mathbf{G} := [\mathbf{M}|\mathbf{F}] := [\mathbf{I}_n \otimes (1, b, \dots, b^{l-1})|\mathbf{F}]$$

Let $\mathbf{x} = (\mathbf{x}_1, \mathbf{x}_2) \in \mathbb{Z}^{nl} \times \mathbb{Z}^{n(k-l)}$ denote the short preimage of $\mathbf{v} := \mathbf{u} - \mathbf{A}\mathbf{p} \pmod{q}$ under the full gadget matrix \mathbf{G} , i.e. $\mathbf{G}\mathbf{x} = \mathbf{v} \pmod{q}$. The main idea of the proof is to first show that the joint distribution of (\mathbf{p}, \mathbf{x}) produced in Algorithm 3 is statistically close to $D_{\Lambda_{\mathbf{u}}^{\perp}[\mathbf{A}, \mathbf{G}], \sqrt{\Sigma_{\mathbf{p}} \oplus \sigma^2}}$ for any $\mathbf{u} \in \mathbb{Z}_q^n$ (this is a stronger theorem than what we need). And then apply the linear transformation theorem on (\mathbf{p}, \mathbf{x}) to obtain the distributions of the preimage \mathbf{y} and the error term \mathbf{e} .

However, applying the linear transformation theorem directly on the lattice coset $\Lambda_{\mathbf{u}}^{\perp}[\mathbf{A}, \mathbf{G}]$ leads to a technical problem. That is, the intermediate lattice intersections $\Lambda_{\mathbf{T}}$ required in Lemma 2.8 or Theorem 2.7 have large smoothing parameters, which means even if we go through that route, the Gaussian width of the resulting preimage would blow up significantly.

To get around this problem, we work only with a uniformly random target \mathbf{u} instead of an arbitrary target. In this case we prove that (\mathbf{p}, \mathbf{x}) is statistically close to $D_{\mathbb{Z}^{m+nk}, \sqrt{\Sigma_{\mathbf{p}} \oplus \sigma^2}}$. Now the support of (\mathbf{p}, \mathbf{x}) is the integer lattice \mathbb{Z}^{m+nk} . Working with the integer lattice is important for two reasons. First, it allows us to treat \mathbf{x}_1 and \mathbf{x}_2 as statistically independent samples; and second, it gives us short vectors in the kernels summoned when using Lemma 2.8 or Theorem 2.7. From here we are able to prove the final distributions of \mathbf{y} and \mathbf{e} are trapdoor-independent.

Formally, let $\varepsilon = \text{negl}(\lambda) > 0$. We first prove three lemmas.

Lemma 4.2. *For any $\sigma \geq \eta_{\varepsilon}(\Lambda^{\perp}(\mathbf{G}))$, the random process of first choosing $\mathbf{v} \leftarrow U(\mathbb{Z}_q^n)$ then returning $\mathbf{x} \leftarrow D_{\Lambda_{\mathbf{v}}^{\perp}(\mathbf{G}), \sigma}$ is statistically close to sampling $\mathbf{x} \leftarrow D_{\mathbb{Z}^{nk}, \sigma}$.*

Proof. The proof follows directly from $\det(\Lambda_q^{\perp}(\mathbf{G})) = q^n$ and Corollary 2.5. Alternatively, one can use two applications of the fact $\rho_r(\Gamma + \mathbf{c}) \in (1 \pm \varepsilon)\sigma^n / \det(\Gamma)$ for any $r \geq \eta_{\varepsilon}(\Gamma)$. The latter yields $\Pr\{\text{Process returns } \mathbf{x}\} \in \left(\frac{1-\varepsilon}{1+\varepsilon}, \frac{1+\varepsilon}{1-\varepsilon}\right) \cdot D_{\mathbb{Z}^{nk}, \sigma}(\mathbf{x})$. \square

Lemma 4.3. *The following random processes are statistically close for any $\sigma \geq \sqrt{b^2 + 1} \cdot \omega(\sqrt{\log n}) \geq \eta_{\varepsilon}(\mathbf{g}^t)$: sample $\mathbf{x}_1 \leftarrow D_{\mathbb{Z}^l, \sigma}$ and return $e = [1, b, \dots, b^{l-1}]\mathbf{x}_1$; or, return $e \leftarrow D_{\mathbb{Z}, \sigma \sqrt{(b^{2l}-1)/(b^2-1)}}$.*

Proof. We use Lemma 2.8 or Theorem 2.7 where $[1, b, \dots, b^{l-1}]$ is the linear transformation. Notice that the kernel of $[1, b, \dots, b^{l-1}]$ is the linear span of $[\mathbf{b}_1, \dots, \mathbf{b}_{l-1}]$ where

$$\mathbf{b}_1 = (b, -1, 0, \dots, 0), \mathbf{b}_2 = (0, b, -1, 0, \dots, 0), \dots, \mathbf{b}_{l-1} = (0, \dots, 0, b, -1) \in \mathbb{Z}^l.$$

The support of $\mathbf{x}_1, \mathbb{Z}^l$, contains the $(l-1)$ -dimensional lattice, $\Gamma = \mathbb{Z}^l \cap \text{Ker}([1, b, \dots, b^{l-1}])$, spanned by $[\mathbf{b}_1, \dots, \mathbf{b}_{l-1}]$. Further, $\sigma \geq \eta_{\varepsilon}(\mathbf{g}^t)$ implies σ is larger than the smoothing parameter of Γ since $\|\mathbf{b}_i\| \leq \sqrt{b^2 + 1}$ for $i = 1, \dots, l-1$. Finally by routine calculation on the Gaussian width (and support), we have $e = [1, b, \dots, b^{l-1}]\mathbf{x}_1 \approx_s D_{\mathbb{Z}, \sigma \sqrt{(b^{2l}-1)/(b^2-1)}}$. \square

Let $\mathbf{R}' := \begin{bmatrix} \mathbf{R} \\ \mathbf{I}_{n(k-l)} \end{bmatrix}$. Next, we analyze the distribution given by the linear transformation representing the convolution step:

$$\mathbf{y} = \mathbf{p} + \mathbf{R}'\mathbf{x}_2 = [\mathbf{I}_m | \mathbf{R}'] \begin{pmatrix} \mathbf{p} \\ \mathbf{x}_2 \end{pmatrix}$$

for $(\mathbf{p}, \mathbf{x}_2) \leftarrow D_{\mathbb{Z}^{m+n(k-l)}, \sqrt{\Sigma_{\mathbf{p}} \oplus \sigma^2 \mathbf{I}_{n(k-l)}}}$. Let $\mathbf{L} := [\mathbf{I}_m | \mathbf{R}']$ in Lemma 4.4 and its proof below.

Lemma 4.4. *For any $\sqrt{\Sigma_{\mathbf{p}}} \geq \eta_{\varepsilon}(\mathbb{Z}^{n(k-l)})$, and $\sigma \geq \eta_{\varepsilon}(\mathbb{Z}^{n(k-l)})$, $\mathbf{L}D_{\mathbb{Z}^{m+n(k-l)}, \sqrt{\Sigma_{\mathbf{p}} \oplus \sigma^2 \mathbf{I}_{n(k-l)}}}$ is statistically close to $D_{\mathbb{Z}^m, s}$.*

Proof. The range and covariance are immediate. Next, we use Theorem 2.7 (Lemma 2.8). The kernel of \mathbf{L} is given by all vectors (\mathbf{a}, \mathbf{b}) where $\mathbf{b} \in \mathbb{R}^{n(k-l)}$ and $\mathbf{a} = -\mathbf{R}'\mathbf{b}$. The integer lattice $\mathbb{Z}^{m+n(k-l)}$ contains all such integer vectors so $\Lambda_{\mathbf{L}} := \mathbb{Z}^{m+n(k-l)} \cap \ker(\mathbf{L})$ spans \mathbf{L} 's kernel.

Now we determine the smoothing parameter of $\Lambda_{\mathbf{L}}$. Rotate the space by

$$\mathbf{Q} := \begin{bmatrix} \mathbf{I} & \mathbf{R} \\ \mathbf{0} & \mathbf{I}_{n(k-l)} \end{bmatrix}.$$

This rotation yields $\mathbf{Q}\Lambda_{\mathbf{L}} = \{\mathbf{0}\} \oplus \mathbb{Z}^{n(k-l)}$ and since rotating a covariance does not change its eigenvalues, we have $\sqrt{\Sigma_p \oplus \sigma^2 \mathbf{I}_{n(k-l)}} \geq \eta_\epsilon(\mathbb{Z}^{n(k-l)}) = \eta_\epsilon(\Lambda_{\mathbf{L}})$. This implies $\mathbf{L}D_{\mathbb{Z}^{m+n(k-l)}, \sqrt{\Sigma_p \oplus \sigma^2 \mathbf{I}_{n(k-l)}}}$ is statistically close to $D_{\mathbb{Z}^m, s}$. \square

We are now ready to prove Theorem 4.1.

Proof. (of Theorem 4.1) The proof's overview is given via the following. Let

- $\mathbf{p} \leftarrow D_{\mathbb{Z}^m, \sqrt{\Sigma_p}}$ be a perturbation,
- $\mathbf{u} \in \mathbb{Z}_q^n$ be the input target coset,
- $\mathbf{v} = \mathbf{u} - \mathbf{A}\mathbf{p} \in \mathbb{Z}_q^n$ be the G-lattice coset,
- $\mathbf{x} = (\mathbf{x}_1, \mathbf{x}_2) \leftarrow D_{\mathbb{Z}^{nk}, \sigma}$ (G-lattice randomized over uniform coset \mathbf{v} and $\sigma \geq \eta_\epsilon(\mathbf{g}^t)$, Lemma 4.2)
- $\mathbf{e} \leftarrow D_{\mathbb{Z}^{nl}, \sigma \sqrt{(b^{2l}-1)/(b^2-1)}}$ be the concatenation of the errors, e , in Lemma 4.3,
- and $\mathbf{y} \leftarrow D_{\mathbb{Z}^m, s}$ as in Lemma 4.4.

The proof is best summarized via the sequence of hybrids below:

$$\begin{aligned} \mathbf{u} &= \mathbf{v} + \mathbf{A}\mathbf{p} \\ &\approx_s \mathbf{G}\mathbf{x} + \mathbf{A}\mathbf{p} \\ &= \mathbf{M}\mathbf{x}_1 + \mathbf{F}\mathbf{x}_2 + \mathbf{A}\mathbf{p} \\ &\approx_s \mathbf{e} + \mathbf{F}\mathbf{x}_2 + \mathbf{A}\mathbf{p} \\ &= \mathbf{e} + \mathbf{A}\mathbf{R}'\mathbf{x}_2 + \mathbf{A}\mathbf{p} \\ &= \mathbf{e} + \mathbf{A}\mathbf{L} \begin{pmatrix} \mathbf{p} \\ \mathbf{x}_2 \end{pmatrix} \\ &\approx_s \mathbf{e} + \mathbf{A}\mathbf{y}. \end{aligned}$$

The first \approx_s is through swapping the order of sampling \mathbf{u} and \mathbf{v} uniformly at random, then using the fact that $\sigma \geq \eta_\epsilon(\mathbf{G})$ (Lemma 4.2). The next \approx_s is given by Lemma 4.3. Finally, the last \approx_s is given by concatenating $(\mathbf{p}, \mathbf{x}_2) \leftarrow D_{\mathbb{Z}^{m+n(k-l)}, \sqrt{\Sigma_p \oplus \sigma}}$ and using Lemma 4.4.

We remark that the key in the equivalences above is that we can separate \mathbf{x} into two statistically independent samples, \mathbf{x}_1 and \mathbf{x}_2 , concatenate \mathbf{p} and \mathbf{x}_2 , then perform two instances of Theorem 2.7 (Lemma 2.8) on the statistically independent samples $\mathbf{L}(\mathbf{p}, \mathbf{x}_2)$ and $\mathbf{M}\mathbf{x}_1$. The statistical independence of \mathbf{x}_1 and \mathbf{x}_2 is due to the orthogonality of \mathbb{Z}^{nk} and the same cannot be said if $\mathbf{x} \sim D_{\Lambda_{\mathbf{L}}^\perp(\mathbf{G}), \sigma}$

for a fixed \mathbf{v} (via a fixed \mathbf{u}). This difference highlights why we must argue security for a uniformly random input coset \mathbf{u} (and \mathbf{v}).

Real distribution: The real distribution of $\{(\mathbf{A}, \mathbf{y}, \mathbf{u}, \mathbf{e})\}$ is:

$$\mathbf{A}, \mathbf{u} \leftarrow U(\mathbb{Z}_q^n), \mathbf{p} \leftarrow D_{\mathbb{Z}^m, \sqrt{\Sigma_p}}, \mathbf{v} := \mathbf{u} - \mathbf{A}\mathbf{p}, \mathbf{x} = (\mathbf{x}_1, \mathbf{x}_2) \leftarrow D_{\Lambda_{\mathbf{v}}^\perp(\mathbf{G}), \sigma}, \mathbf{e} = \mathbf{M}\mathbf{x}_1, \text{ and } \mathbf{y} = \mathbf{L}(\mathbf{p}, \mathbf{x}_2).$$

Hybrid 1: Here we swap the order of sampling \mathbf{u} and \mathbf{v} by first sampling $\mathbf{v} \leftarrow U(\mathbb{Z}_q^n)$ and setting $\mathbf{u} := \mathbf{v} + \mathbf{A}\mathbf{p}$: $\mathbf{A}, \mathbf{v} \leftarrow U(\mathbb{Z}_q^n), \mathbf{p} \leftarrow D_{\mathbb{Z}^m, \sqrt{\Sigma_p}}, \mathbf{u} = \mathbf{v} + \mathbf{A}\mathbf{p}$. We keep \mathbf{x}, \mathbf{e} , and \mathbf{y} unchanged: $\mathbf{x} = (\mathbf{x}_1, \mathbf{x}_2) \leftarrow D_{\Lambda_{\mathbf{v}}^\perp(\mathbf{G}), \sigma}, \mathbf{e} = \mathbf{M}\mathbf{x}_1$, and $\mathbf{y} = \mathbf{L}(\mathbf{p}, \mathbf{x}_2)$. Then, the real distribution and Hybrid 1 are the same.

Hybrid 2: Instead of sampling a uniform $\mathbf{v} \in \mathbb{Z}_q^n$ and a \mathbf{G} -lattice sample $\mathbf{x} = (\mathbf{x}_1, \mathbf{x}_2) \leftarrow D_{\Lambda_{\mathbf{v}}^\perp(\mathbf{G}), \sigma}$, we sample $\mathbf{x} \leftarrow D_{\mathbb{Z}^{nk}, \sigma}$ and $\mathbf{v} = \mathbf{G}\mathbf{x} \in \mathbb{Z}_q^n$. The rest remains the same:

$$\mathbf{A}, \mathbf{x} \leftarrow D_{\mathbb{Z}^{nk}, \sigma}, \mathbf{v} = \mathbf{G}\mathbf{x}, \mathbf{p} \leftarrow D_{\mathbb{Z}^m, \sqrt{\Sigma_p}}, \mathbf{u} \approx_s \mathbf{G}\mathbf{x} + \mathbf{A}\mathbf{p}, \mathbf{e} = \mathbf{M}\mathbf{x}_1, \text{ and } \mathbf{y} = \mathbf{L}(\mathbf{p}, \mathbf{x}_2). \text{ Lemma 4.2 implies Hybrid 1 and Hybrid 2 are statistically close.}$$

Hybrid 3: We combine \mathbf{p}, \mathbf{x}_2 into the joint distribution $(\mathbf{p}, \mathbf{x}_2) \leftarrow D_{\mathbb{Z}^{m+n(k-l)}, \sqrt{\Sigma_p \oplus \sigma^2 \mathbf{I}}}$:

$$\mathbf{A}, (\mathbf{p}, \mathbf{x}_2) \leftarrow D_{\mathbb{Z}^{m+n(k-l)}, \sqrt{\Sigma_p \oplus \sigma^2 \mathbf{I}}}, \mathbf{e} = \mathbf{M}\mathbf{x}_1, \mathbf{y} = \mathbf{L}(\mathbf{p}, \mathbf{x}_2), \mathbf{v} = \mathbf{G}\mathbf{x}, \text{ and } \mathbf{u} = \mathbf{v} + \mathbf{A}\mathbf{p} = \mathbf{M}\mathbf{x}_1 + \mathbf{F}\mathbf{x}_2 + \mathbf{A}\mathbf{p}.$$

Hybrid 4: Here we apply the linear transformation theorem on \mathbf{e}, \mathbf{y} (statistically independent \mathbf{e} and \mathbf{y}):

$$\mathbf{A}, (\mathbf{p}, \mathbf{x}_2) \leftarrow D_{\mathbb{Z}^{m+n(k-l)}, \sqrt{\Sigma_p \oplus \sigma^2 \mathbf{I}}}, \mathbf{e} \leftarrow D_{\mathbb{Z}^{nl}, \sigma \sqrt{(b^{2l}-1)/(b^2-1)}}, \mathbf{y} \leftarrow D_{\mathbb{Z}^m, s}, \mathbf{v} = \mathbf{G}\mathbf{x}, \text{ and } \mathbf{u} = \mathbf{M}\mathbf{x}_1 + \mathbf{F}\mathbf{x}_2 + \mathbf{A}\mathbf{p} \approx_s \mathbf{A}\mathbf{L}(\mathbf{p}, \mathbf{x}_2) + \mathbf{e} \approx_s \mathbf{A}\mathbf{y} + \mathbf{e}.$$

Lemmas 4.3 and 4.4 imply Hybrids 3 and 4 are statistically close.

Final distribution: Sample $\mathbf{A} \leftarrow U(\mathbb{Z}_q^{n \times m})$ and keep the rest of the vectors from the same distribution as Hybrid 4 (notice that the trapdoor \mathbf{R} of \mathbf{A} is not used to sample $\mathbf{p}, \mathbf{x}, \mathbf{e}$ and \mathbf{y}). The final distribution is computationally indistinguishable from Hybrid 4 assuming $\text{LWE}_{n,n,q,\chi,U(\mathbb{Z}_q),\chi}$. \square

5 Hash-and-Sign Signature Instantiated with the Approximate Trapdoor

We spell out the details of the hash-and-sign signature scheme from [GPV08] instantiated with the approximate \mathbf{G} -trapdoor instead of an exact trapdoor.

Recall the parameters from the last section. We set $k = \lceil \log_b q \rceil$, set l to be the number of entries dropped from the \mathbf{G} -trapdoor such that $1 \leq l < k$ and $m = n(2 + (k - l))$. Let $\sigma, s \in \mathbb{R}^+$ be the discrete Gaussian widths of the distributions over the cosets of $\Lambda_q^\perp(\mathbf{G})$ and $\Lambda_q^\perp(\mathbf{A})$ respectively. Let χ be the distribution of the entries of the trapdoor \mathbf{R} chosen so that $\text{LWE}_{n,n,q,\chi,U(\mathbb{Z}_q),\chi}$ is hard.

Construction 5.1. *Given an approximate trapdoor sampler from Theorem 4.1, a hash function $H = \{H_\lambda : \{0, 1\}^* \rightarrow R_\lambda\}$ modeled as a random oracle, we build a signature scheme as follows.*

- $\text{Gen}(1^\lambda)$: *The key-generation algorithm samples $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$ together with its (α, β) -approximate trapdoor \mathbf{R} from $\text{APPROX.TRAPGEN}(1^\lambda)$. Let the range R_λ of H be \mathbb{Z}_q^n . It outputs \mathbf{A} as the verification key, keeps \mathbf{R} as the secret signing key.*

- $\text{Sig}(\mathbf{R}, m)$: The signing algorithm checks if the message-signature pair (m, \mathbf{x}_m) has been produced before. If so, it outputs \mathbf{x}_m as the signature of m ; if not, computes $\mathbf{u} = H(m)$, and samples an approximate preimage $\mathbf{x}_m \leftarrow \text{APPROX.SAMPLEPRE}(\mathbf{A}, \mathbf{R}, \mathbf{u}, s)$. It outputs \mathbf{x}_m as the signature and stores (m, \mathbf{x}_m) in the list.
- $\text{Ver}(\mathbf{A}, m, \mathbf{x})$: The verification algorithm checks if $\|\mathbf{x}\| \leq \beta$ and $\|\mathbf{A} \cdot \mathbf{x} - H(m)\| \leq \alpha$. If so, it outputs accept; otherwise, it outputs reject.

5.1 Security analysis

In the security analysis we use the following properties on the distributions produced by APPROX.SAMPLEPRE proven in Theorem 4.1:

1. An approximate preimage of a uniformly random image is statistically close to a distribution D_{pre} , which is independent of \mathbf{A} and \mathbf{R} .
2. For a uniformly random image \mathbf{u} , the error vector $\mathbf{e} := \mathbf{u} - \mathbf{A}\mathbf{x} \pmod{q}$ is statistically close to a distribution D_{err} , which is also independent of \mathbf{A} and \mathbf{R} .

To prove that the signature satisfies the strong EU-CMA security, we need an additional “near-collision-resistance” property for Ajtai’s function, which can be based on the standard SIS assumption. Let us remark that without this property, we can still prove the signature scheme satisfies static security based on the hardness of the approximate ISIS problem, which is tighter by a factor of two according to Theorem 3.4.

Lemma 5.2 (The near-collision-resistance of Ajtai’s function). *For any $n, m, q \in \mathbb{N}$ and $\alpha, \beta \in \mathbb{R}$. If there is an efficient adversary A that given $\mathbf{A} \leftarrow U(\mathbb{Z}_q^{n \times m})$, finds $\mathbf{x}_1 \neq \mathbf{x}_2 \in \mathbb{Z}^m$ such that*

$$\|\mathbf{x}_1\| \leq \beta \quad \text{and} \quad \|\mathbf{x}_2\| \leq \beta \quad \text{and} \quad \|\mathbf{A}\mathbf{x}_1 - \mathbf{A}\mathbf{x}_2 \pmod{q}\| \leq 2\alpha$$

Then there is an efficient adversary B that solves $\text{SIS}_{n, n+m, q, 2(\alpha+\beta)}$.

Proof. Suppose B gets an $\text{HNF.SIS}_{n, n+m, q, 2(\alpha+\beta)}$ challenge (which is as hard as $\text{SIS}_{n, n+m, q, 2(\alpha+\beta)}$) with the public matrix $[\mathbf{I}_n \mid \mathbf{A}]$, B sends \mathbf{A} to A , gets back $\mathbf{x}_1 \neq \mathbf{x}_2 \in \mathbb{Z}^m$ such that

$$\|\mathbf{x}_1\| \leq \beta \quad \text{and} \quad \|\mathbf{x}_2\| \leq \beta \quad \text{and} \quad \|\mathbf{y} := \mathbf{A}\mathbf{x}_1 - \mathbf{A}\mathbf{x}_2 \pmod{q}\| \leq 2\alpha$$

B then sets $\mathbf{z} := [-\mathbf{y}^t \mid (\mathbf{x}_1 - \mathbf{x}_2)^t]^t$ as the solution. \mathbf{z} is then non-zero and satisfies $\|\mathbf{z}\| \leq 2(\alpha + \beta)$ and $[\mathbf{I}_n \mid \mathbf{A}]\mathbf{z} = \mathbf{0} \pmod{q}$. \square

Theorem 5.3. *Construction 5.1 is strongly existentially unforgeable under a chosen-message attack in the random oracle model assuming the hardness of $\text{SIS}_{n, n+m, q, 2(\alpha+\beta)}$ and $\text{LWE}_{n, n, q, \chi, U(\mathbb{Z}_q), \chi}$.*

Proof. Suppose there is a polynomial time adversary A that breaks the strong EU-CMA of the signature scheme, we construct a polynomial time adversary B that breaks the near-collision-resistance of Ajtai’s function, which is as hard as $\text{SIS}_{n, n+m, q, 2(\alpha+\beta)}$ due to Lemma 5.2.

To start, B sends Ajtai’s function \mathbf{A} to A as the public key for the signature scheme. Once A makes a random oracle query w.r.t. a message m , B samples $\mathbf{x} \leftarrow D_{\text{pre}}$, computes $\mathbf{u} := \mathbf{A}\mathbf{x} + D_{\text{err}} \pmod{q}$ as the random oracle response on m . B then replies \mathbf{u} to A and stores (m, \mathbf{u}) in the random oracle

Params	Exact	Approx	Approx	Exact	Approx	Approx	Exact	Approx	Approx
n	512	512	512	512	512	512	512	512	512
$\lceil \log_2 q \rceil$	24	24	24	20	20	20	16	16	16
b	2	2	2	2	2	2	2	2	2
l	0	12	15	0	10	12	0	7	9
τ	40	40	40	10	10	10	2.6	2.6	2.6
s	38317.0	29615.3	26726.3	8946.4	6919.8	6416.4	2170.7	1756.3	1618.2
m	13312	7168	5632	11264	6144	5120	9216	5632	4608
$\ \mathbf{x}\ _2$	4441737.7	2521387.0	2035008.5	956758.1	545470.5	464022.0	211100.9	133305.5	109339.1
$\ \mathbf{x}\ _\infty$	184653	111909	94559	38507	25275	24762	8848	6853	6334
$\ \mathbf{z}\ _2$	0	374014.0	2118987.6	0	94916.6	343682.9	0	11897.9	46428.4
$\ \mathbf{z}\ _\infty$	0	46895	346439	0	13265	52789	0	1439	7213
PK	37.50	19.50	15.00	26.25	13.75	11.25	17.00	10.00	8.00
Sig	25.68	13.53	10.51	18.87	10.01	8.29	13.16	7.83	6.30
LWE	100.0	100.0	100.0	102.8	102.8	102.8	104.7	104.7	104.7
ISIS	80.2	85.8	81.1	82.0	87.5	84.3	83.7	89.0	88.1
δ	1.00685	1.00643	1.00678	1.00670	1.00631	1.00653	1.00658	1.00621	1.00628
k	174	193	177	180	199	188	186	204	201

Figure 4: Summary of the concrete parameters, with base $b = 2$, aiming at around 80 to 90-bit security. The sizes of PK and Sig are measured in kB. τ is the Gaussian width of the secret matrix \mathbf{R} . s is the Gaussian width of the preimage. “LWE” refers to the security level of the pseudorandomness of the PK. “ISIS” refers to the security level of breaking approximate ISIS. δ and k are the variables used in the ISIS security estimation.

storage, (m, \mathbf{x}) in the message-signature pair storage. Once A makes a signing query on the message m (wlog assume m has been queried to the random oracle before, since if not B can query it now), B finds (m, \mathbf{x}) in the storage and reply \mathbf{x} as the signature. The signatures and the hash outputs produced by B are indistinguishable from the real ones due to the properties of the distributions D_{pre} and D_{err} , and the assumption that a real public key is indistinguishable from random under $\text{LWE}_{n,n,q,\chi,U(\mathbb{Z}_q),\chi}$.

Without loss of generality, assume that before A tries to forge a signature on m^* , A has queried H on m^* . Denote the pair that B prepares and stores in the random oracle storage as (m^*, \mathbf{u}^*) , and the pair in the signature storage as (m^*, \mathbf{x}^*) . Finally A outputs \mathbf{x} as the forged signature on m^* . So we have $\|\mathbf{A}(\mathbf{x} - \mathbf{x}^*) \pmod{q}\| \leq 2\alpha$. It remains to prove that $\mathbf{x} \neq \mathbf{x}^*$ so as to use them as a near-collision-pair. If m^* has been queried to the signing oracle before, then $\mathbf{x} \neq \mathbf{x}^*$ by the definition of a successful forgery; if m^* has not been queried to the signing oracle before, then \mathbf{x}^* is with high min-entropy by the settings of the parameter, so $\mathbf{x} \neq \mathbf{x}^*$ with overwhelming probability. \square

5.2 Concrete parameters

We provide a proof-of-concept implementation of the signature. Experiments are performed over several groups of parameters using different dimensions n , moduli q , bases b , targeting different security level (mainly around 80 to 90-bit and 170 to 185-bit security). In each group of parameters, we use fixed n, q, b , and compare the use of exact trapdoor (under our reference implementation) versus approximate trapdoor. In Figures 4 and 5 we list 6 groups of parameters.

Methods for security estimation. Let us first explain how we make the security estimations. The concrete security estimation of lattice-based cryptographic primitive is a highly active research area

Params	Exact	Approx	Approx	Exact	Approx	Approx	Exact	Approx	Approx
n	512	512	512	1024	1024	1024	1024	1024	1024
$\lceil \log_2 q \rceil$	16	16	16	18	18	18	18	18	18
b	4	4	4	8	8	8	4	4	4
l	0	2	4	0	2	3	0	4	5
τ	2.6	2.6	2.6	2.8	2.8	2.8	2.8	2.8	2.8
s	3114.2	2833.3	2505.6	8861.1	7824.8	7227.9	5118.8	4297.8	4015.5
m	5120	4096	3072	8192	6144	5120	11264	7168	6144
$\ \mathbf{x}\ _2$	223740.1	183004.9	138145.7	805772.9	604711.5	516446.3	552713.4	369981.2	311153.9
$\ \mathbf{x}\ _\infty$	13320	11868	8948	35348	28823	30435	19274	18283	14927
$\ \mathbf{z}\ _2$	0	1402.3	19807.1	0	7316.5	54379.8	0	29958.0	115616.4
$\ \mathbf{z}\ _\infty$	0	174	2448	0	905	6680	0	3025	12070
PK	9.00	7.00	5.00	15.75	11.25	9.00	22.50	13.50	11.25
Sig	7.62	5.94	4.45	13.70	10.14	8.36	18.74	11.09	9.38
LWE	104.7	104.7	104.7	192.7	192.7	192.7	192.7	192.7	192.7
AISIS	82.8	85.5	87.8	165.3	172.9	174.9	175.8	185.7	183.7
δ	1.00664	1.00645	1.00629	1.0036	1.00347	1.00343	1.00342	1.00326	1.00329
k	183	192	200	462	488	495	498	532	525

Figure 5: Summary of the concrete parameters, with base $b \geq 4$, aiming at around 80 to 90-bit and 170 to 184-bit security.

and more sophisticated methods are proposed recently. Here we use relatively simple methods to estimate the pseudorandomness of the public-key (henceforth “LWE security”), and the hardness of breaking approximate ISIS (henceforth “AISIS security”). Let us remark that our estimations may not reflect the state-of-art, but at least provide a fair comparison of the parameters for the exact trapdoor versus the approximate trapdoor.

LWE security depends on the choices of q , n , and the Gaussian width τ of the trapdoor \mathbf{R} . The estimation of LWE security was done with the online LWE bit security estimator with BKZ as the reduction model¹ [ACD⁺18].

For the approximate ISIS problem, the only direct cryptanalysis result we are aware of is the work of Bai et al. [BGLS19], but it is not clearly applicable to the parameters we are interested. Instead we estimate AISIS through $\text{ISIS}_{n,m,q,\alpha+\beta}$ following the reduction in Lemma 3.5, where α and β are the upper-bounds of l_2 norm of the error \mathbf{z} and preimage \mathbf{x} . We estimate the security level of $\text{ISIS}_{n,m,q,\alpha+\beta}$ based on how many operations BKZ would take to find a vector in the lattice $\Lambda_q^\perp(\mathbf{A})$ of length $\alpha + \beta$. Further, we can throw away columns in \mathbf{A} . We choose to only use $2n$ columns of \mathbf{A} as done in [BFRS18], denoted \mathbf{A}_{2n} , since Minkowski’s theorem² tells us $\Lambda_q^\perp(\mathbf{A}_{2n})$ has a short enough vector. Following [APS15, ACD⁺18], we use sieving as the SVP oracle with time complexity $2^{.292k+16.4}$ in the block size, k . BKZ is expected to return a vector of length $\delta^{2n} \det^{1/2n}$ for a lattice of dimension $2n$. Hence, we found the smallest block size k achieving the needed δ corresponding to forging a signature, $\frac{\alpha+\beta}{\sqrt{q}} = \delta^{2n}$. Finally, we used the heuristic $\delta \approx (\frac{k}{2\pi e} (\pi k)^{1/k})^{1/2(k-1)}$ to determine the relation between k and δ , and we set the total time complexity of BKZ with block-size k , dimension $2n$ as $8 \cdot 2n \cdot \text{time}(SVP) = 8 \cdot 2n \cdot 2^{.292k+16.4}$ [Che13, APS15]³.

¹<https://bitbucket.org/malb/lwe-estimator>

²For any lattice \mathbf{L} , $\lambda_1 \leq \sqrt{r} \det(\mathbf{L})^{1/r}$ where r is the rank of the lattice.

³We use the “magic eight four number” for BKZ to keep consistency with the LWE online estimator. We have not incorporated the more recent developments in [Duc18] and [ADH⁺19].

The comparison. For an estimation of 80-bit⁴ security, our reference implementation for the exact trapdoor under the modulus $q \approx 2^{24}$ and base $b = 2$ matches the parameters reported in [BB13] (the parameters in the other implementation [BFRS18, GPR⁺18] are possibly measured in different ways). We also use smaller moduli and bigger bases to reduce the size and increase the security level. The parameters in Figures 4 and 5 suggest that for all the choices of q and b , using the approximate gadget trapdoor by setting $l = \lceil (\log_b q)/2 \rceil$ saves about half of the sizes in the public key and signatures comparing to using the exact trapdoor, with even a slight increase in the security estimation.

Our implementation shows that the sizes of the public-key & signature can be reduced to 5 kB & 4.45 kB for an estimation of 88-bit security, and 11.25 kB & 9.38 kB for an estimation of 184-bit security. Those are still larger than, but much closer to the sizes for the signatures based on the rejection sampling approach [Lyu12, BG14, DKL⁺18, ABB⁺19]. As a reference, the sizes of the public-key & signature for qTESLA [ABB⁺19] are 4.03 kB & 3.05 kB for an estimation of 128-bit security, and 8.03 kB & 6.03 kB for an estimation of 192-bit security.

Acknowledgment

We are grateful to Daniele Micciancio for valuable advice and his generous sharing of ideas on the subject of this work. We would also like to thank Léo Ducas, Steven Galbraith, Thomas Prest, Yang Yu, Chuang Gao, Eamonn Postlethwaite, and the anonymous reviewers for their helpful suggestions and comments.

References

- [AAAS⁺19] Gorjan Alagic, Gorjan Alagic, Jacob Alperin-Sheriff, Daniel Apon, David Cooper, Quynh Dang, Yi-Kai Liu, Carl Miller, Dustin Moody, Rene Peralta, et al. *Status Report on the First Round of the NIST Post-Quantum Cryptography Standardization Process*. US Department of Commerce, National Institute of Standards and Technology, 2019.
- [ABB10] Shweta Agrawal, Dan Boneh, and Xavier Boyen. Efficient lattice (H)IBE in the standard model. In *EUROCRYPT*, volume 6110 of *Lecture Notes in Computer Science*, pages 553–572. Springer, 2010.
- [ABB⁺19] Erdem Alkim, Paulo S. L. M. Barreto, Nina Bindel, Patrick Longa, and Jefferson E. Ricardini. The lattice-based digital signature scheme qtesla. *IACR Cryptology ePrint Archive*, 2019:85, 2019.
- [ACD⁺18] Martin R. Albrecht, Benjamin R. Curtis, Amit Deo, Alex Davidson, Rachel Player, Eamonn W. Postlethwaite, Fernando Virdia, and Thomas Wunderer. Estimate all the {LWE, NTRU} schemes! In *Security and Cryptography for Networks - 11th International Conference, SCN 2018, Amalfi, Italy, September 5-7, 2018, Proceedings*, pages 351–367, 2018.

⁴When one applies our security estimate methods to Table 1 of [BB13], one gets 82-bit security under the $\lambda = 97$, $n = 512$, $q = 2^{24}$ column.

- [ACPS09] Benny Applebaum, David Cash, Chris Peikert, and Amit Sahai. Fast cryptographic primitives and circular-secure encryption based on hard learning problems. In Shai Halevi, editor, *Advances in Cryptology - CRYPTO 2009, 29th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 16-20, 2009. Proceedings*, volume 5677 of LNCS, pages 595–618. Springer, 2009.
- [ADH⁺19] Martin R. Albrecht, Léo Ducas, Gottfried Herold, Elena Kirshanova, Eamonn W. Postlethwaite, and Marc Stevens. The general sieve kernel and new records in lattice reduction. In *EUROCRYPT (2)*, volume 11477 of *Lecture Notes in Computer Science*, pages 717–746. Springer, 2019.
- [Ajt96] Miklós Ajtai. Generating hard instances of lattice problems (extended abstract). In *STOC*, pages 99–108, 1996.
- [Ajt99] Miklós Ajtai. Generating hard instances of the short basis problem. In Jirí Wiedermann, Peter van Emde Boas, and Mogens Nielsen, editors, *Automata, Languages and Programming, 26th International Colloquium, ICALP’99, Prague, Czech Republic, July 11-15, 1999, Proceedings*, volume 1644 of LNCS, pages 1–9. Springer, 1999.
- [AP11] Joël Alwen and Chris Peikert. Generating shorter bases for hard random lattices. *Theory of Computing Systems*, 48(3):535–553, 2011.
- [APS15] Martin R. Albrecht, Rachel Player, and Sam Scott. On the concrete hardness of learning with errors. *J. Mathematical Cryptology*, 9(3):169–203, 2015.
- [BB13] Rachid El Bansarkhani and Johannes A. Buchmann. Improvement and efficient implementation of a lattice-based signature scheme. In *Selected Areas in Cryptography*, volume 8282 of *Lecture Notes in Computer Science*, pages 48–67. Springer, 2013.
- [BFRS18] Pauline Bert, Pierre-Alain Fouque, Adeline Roux-Langlois, and Mohamed Sabt. Practical implementation of ring-sis/lwe based signature and IBE. In *PQCrypto*, volume 10786 of *Lecture Notes in Computer Science*, pages 271–291. Springer, 2018.
- [BG14] Shi Bai and Steven D. Galbraith. An improved compression technique for signatures based on learning with errors. In *CT-RSA*, volume 8366 of *Lecture Notes in Computer Science*, pages 28–47. Springer, 2014.
- [BGLS19] Shi Bai, Steven D. Galbraith, Liangze Li, and Daniel Sheffield. Improved combinatorial algorithms for the inhomogeneous short integer solution problem. *J. Cryptology*, 32(1):35–83, 2019.
- [BLP⁺13] Zvika Brakerski, Adeline Langlois, Chris Peikert, Oded Regev, and Damien Stehlé. Classical hardness of learning with errors. In *Proceedings of the forty-fifth annual ACM symposium on Theory of computing*, pages 575–584. ACM, 2013.
- [BPMW16] Florian Bourse, Rafaël Del Pino, Michele Minelli, and Hoeteck Wee. FHE circuit privacy almost for free. In Matthew Robshaw and Jonathan Katz, editors, *Advances in Cryptology - CRYPTO 2016 - 36th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 14-18, 2016, Proceedings, Part II*, volume 9815 of *Lecture Notes in Computer Science*, pages 62–89. Springer, 2016.
- [BVWW16] Zvika Brakerski, Vinod Vaikuntanathan, Hoeteck Wee, and Daniel Wichs. Obfuscating conjunctions under entropic ring LWE. In *ITCS*, pages 147–156. ACM, 2016.

- [CC17] Ran Canetti and Yilei Chen. Constraint-hiding constrained prfs for nc^1 from LWE. In *EUROCRYPT (1)*, volume 10210 of *Lecture Notes in Computer Science*, pages 446–476, 2017.
- [CGM⁺18] Cheng Chen, Nicholas Genise, Daniele Micciancio, Yuriy Polyakov, and Kurt Rohloff. Implementing token-based obfuscation under (ring) LWE. *IACR Cryptology ePrint Archive*, 2018:1222, 2018.
- [Che13] Yuanmi Chen. *Réduction de réseau et sécurité concrète du chiffrement complètement homomorphe*. PhD thesis, Paris 7, 2013.
- [CHKP12] David Cash, Dennis Hofheinz, Eike Kiltz, and Chris Peikert. Bonsai trees, or how to delegate a lattice basis. *Journal of cryptology*, 25(4):601–639, 2012.
- [CVW18] Yilei Chen, Vinod Vaikuntanathan, and Hoeteck Wee. GGH15 beyond permutation branching programs: Proofs, attacks, and candidates. In *CRYPTO (2)*, volume 10992 of *Lecture Notes in Computer Science*, pages 577–607. Springer, 2018.
- [DGPY19] Léo Ducas, Steven Galbraith, Thomas Prest, and Yang Yu. Integral matrix gram root and lattice gaussian sampling without floats. *IACR Cryptology ePrint Archive*, 2019:320, 2019.
- [DKL⁺18] Léo Ducas, Eike Kiltz, Tancrede Lepoint, Vadim Lyubashevsky, Peter Schwabe, Gregor Seiler, and Damien Stehlé. Crystals-dilithium: A lattice-based digital signature scheme. *IACR Trans. Cryptogr. Hardw. Embed. Syst.*, 2018(1):238–268, 2018.
- [DN12] Léo Ducas and Phong Q. Nguyen. Faster gaussian lattice sampling using lazy floating-point arithmetic. In *Advances in Cryptology - ASIACRYPT 2012 - 18th International Conference on the Theory and Application of Cryptology and Information Security, Beijing, China, December 2-6, 2012. Proceedings*, pages 415–432, 2012.
- [DP16] Léo Ducas and Thomas Prest. Fast fourier orthogonalization. In *Proceedings of the ACM on International Symposium on Symbolic and Algebraic Computation, ISSAC 2016, Waterloo, ON, Canada, July 19-22, 2016*, pages 191–198, 2016.
- [dPLS18] Rafaël del Pino, Vadim Lyubashevsky, and Gregor Seiler. Lattice-based group signatures and zero-knowledge proofs of automorphism stability. In *Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security, CCS 2018, Toronto, ON, Canada, October 15-19, 2018*, pages 574–591, 2018.
- [Duc18] Léo Ducas. Shortest vector from lattice sieving: A few dimensions for free. In *EUROCRYPT (1)*, volume 10820 of *Lecture Notes in Computer Science*, pages 125–145. Springer, 2018.
- [FHK⁺18] Pierre-Alain Fouque, Jeffrey Hoffstein, Paul Kirchner, Vadim Lyubashevsky, Thomas Pornin, Thomas Prest, Thomas Ricosset, Gregor Seiler, William Whyte, and Zhenfei Zhang. Falcon: Fast-fourier lattice-based compact signatures over ntru, 2018.
- [Gen09] Craig Gentry. Fully homomorphic encryption using ideal lattices. In *STOC*, pages 169–178, 2009.

- [GGH97] Oded Goldreich, Shafi Goldwasser, and Shai Halevi. Public-key cryptosystems from lattice reduction problems. In *Advances in Cryptology - CRYPTO '97, 17th Annual International Cryptology Conference, Santa Barbara, California, USA, August 17-21, 1997, Proceedings*, pages 112–131, 1997.
- [GGH15] Craig Gentry, Sergey Gorbunov, and Shai Halevi. Graph-induced multilinear maps from lattices. In *Theory of Cryptography - 12th Theory of Cryptography Conference, TCC 2015, Warsaw, Poland, March 23-25, 2015, Proceedings, Part II*, pages 498–527, 2015.
- [GKW17] Rishab Goyal, Venkata Koppula, and Brent Waters. Lockable obfuscation. In *FOCS*, pages 612–621. IEEE Computer Society, 2017.
- [GM18] Nicholas Genise and Daniele Micciancio. Faster gaussian sampling for trapdoor lattices with arbitrary modulus. In *Advances in Cryptology - EUROCRYPT 2018 - 37th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Tel Aviv, Israel, April 29 - May 3, 2018 Proceedings, Part I*, pages 174–203, 2018.
- [GPR⁺18] Kamil Doruk Gür, Yuriy Polyakov, Kurt Rohloff, Gerard W Ryan, and ErKay Savas. Implementation and evaluation of improved gaussian sampling for lattice trapdoors. In *Proceedings of the 6th Workshop on Encrypted Computing & Applied Homomorphic Cryptography*, pages 61–71. ACM, 2018.
- [GPV08] Craig Gentry, Chris Peikert, and Vinod Vaikuntanathan. Trapdoors for hard lattices and new cryptographic constructions. In *STOC*, pages 197–206, 2008.
- [GS02] Craig Gentry and Mike Szydlo. Cryptanalysis of the revised ntru signature scheme. In *International Conference on the Theory and Applications of Cryptographic Techniques*, pages 299–320. Springer, 2002.
- [GVW13] Sergey Gorbunov, Vinod Vaikuntanathan, and Hoeteck Wee. Attribute-based encryption for circuits. In *STOC*, pages 545–554. ACM, 2013.
- [HHP⁺03] Jeffrey Hoffstein, Nick Howgrave-Graham, Jill Pipher, Joseph H. Silverman, and William Whyte. NTRUSIGN: digital signatures using the NTRU lattice. In *CT-RSA*, volume 2612 of *Lecture Notes in Computer Science*, pages 122–140. Springer, 2003.
- [HHSS17] Shai Halevi, Tzipora Halevi, Victor Shoup, and Noah Stephens-Davidowitz. Implementing bp-obfuscation using graph-induced encoding. In *ACM Conference on Computer and Communications Security*, pages 783–798. ACM, 2017.
- [HPS98] Jeffrey Hoffstein, Jill Pipher, and Joseph H Silverman. Ntru: A ring-based public key cryptosystem. In *Algorithmic number theory*, pages 267–288. Springer, 1998.
- [LMPR08] Vadim Lyubashevsky, Daniele Micciancio, Chris Peikert, and Alon Rosen. SWIFFT: A modest proposal for FFT hashing. In *Fast Software Encryption, 15th International Workshop, FSE 2008, Lausanne, Switzerland, February 10-13, 2008, Revised Selected Papers*, pages 54–72, 2008.
- [LPR10] Vadim Lyubashevsky, Chris Peikert, and Oded Regev. On ideal lattices and learning with errors over rings. In Henri Gilbert, editor, *Advances in Cryptology - EUROCRYPT 2010, 29th Annual International Conference on the Theory and Applications of Cryptographic Techniques, French Riviera, May 30 - June 3, 2010. Proceedings*, volume 6110 of *LNCS*, pages 1–23. Springer, 2010.

- [Lyu12] Vadim Lyubashevsky. Lattice signatures without trapdoors. In *EUROCRYPT*, volume 7237 of *Lecture Notes in Computer Science*, pages 738–755. Springer, 2012.
- [Mic] Daniele Micciancio. personal communication.
- [Mic02] Daniele Micciancio. Generalized compact knapsacks, cyclic lattices, and efficient one-way functions from worst-case complexity assumptions. In *43rd Symposium on Foundations of Computer Science (FOCS 2002), 16-19 November 2002, Vancouver, BC, Canada, Proceedings*, pages 356–365, 2002.
- [MP12] Daniele Micciancio and Chris Peikert. Trapdoors for lattices: Simpler, tighter, faster, smaller. In *Advances in Cryptology - EUROCRYPT 2012 - 31st Annual International Conference on the Theory and Applications of Cryptographic Techniques, Cambridge, UK, April 15-19, 2012. Proceedings*, pages 700–718, 2012.
- [MP13] Daniele Micciancio and Chris Peikert. Hardness of sis and lwe with small parameters. In *Advances in Cryptology - CRYPTO 2013*, pages 21–39. Springer, 2013.
- [MR07] Daniele Micciancio and Oded Regev. Worst-case to average-case reductions based on Gaussian measure. *SIAM Journal on Computing*, 37(1):267–302, 2007.
- [MW18] Daniele Micciancio and Michael Walter. On the bit security of cryptographic primitives. In *EUROCRYPT (1)*, volume 10820 of *Lecture Notes in Computer Science*, pages 3–28. Springer, 2018.
- [NR06] Phong Q. Nguyen and Oded Regev. Learning a parallelepiped: Cryptanalysis of GGH and NTRU signatures. In *Advances in Cryptology - EUROCRYPT 2006, 25th Annual International Conference on the Theory and Applications of Cryptographic Techniques, St. Petersburg, Russia, May 28 - June 1, 2006, Proceedings*, pages 271–288, 2006.
- [Pei09] Chris Peikert. Public-key cryptosystems from the worst-case shortest vector problem: extended abstract. In *Proceedings of the 41st Annual ACM Symposium on Theory of Computing, STOC 2009, Bethesda, MD, USA, May 31 - June 2, 2009*, pages 333–342, 2009.
- [PR06] Chris Peikert and Alon Rosen. Efficient collision-resistant hashing from worst-case assumptions on cyclic lattices. In *Theory of Cryptography*, pages 145–166. Springer, 2006.
- [PRS17] Chris Peikert, Oded Regev, and Noah Stephens-Davidowitz. Pseudorandomness of ring-lwe for any ring and modulus. In *STOC*, pages 461–473. ACM, 2017.
- [Reg09] Oded Regev. On lattices, learning with errors, random linear codes, and cryptography. *J. ACM*, 56(6), 2009.
- [WZ17] Daniel Wichs and Giorgos Zirdelis. Obfuscating compute-and-compare programs under LWE. In *FOCS*, pages 600–611. IEEE Computer Society, 2017.

A The Ring Setting

Here we verify the approximate gadget trapdoor construction given in Section 4 holds in the ring setting [Mic02, LPR10].

Rings. Let $\phi_{2n}(x)$ be the $2n$ -th cyclotomic polynomial where n is a power of two, $R := \mathbb{Z}[x]/\phi_{2n}(x)$ be the associated cyclotomic ring, and $q \geq 2$ be an integer modulus so that $R_q := R/qR$ is the polynomial ring with coefficients modulo q . Usually q is taken as a prime with $q \equiv 1 \pmod{2n}$ in order to use the Number Theoretic Transform (NTT) [LMPR08]. The ring dimension, n , ultimately depends on an underlying security parameter, λ .

Coefficient embedding. Geometrically, we are concerned with the *coefficient embedding* where each polynomial $h(x) \in R$ is associated with its ordered vector of coefficients,

$$h(x) \mapsto \mathbf{h} = (h_0, h_1, \dots, h_{n-1}) \in \mathbb{Z}^n.$$

Multiplication by a polynomial $g(x) = \sum_{i=0}^{n-1} g_i x^i \in R$ is given by the anti-cyclic (also called nega-cyclic) matrix

$$\psi(g) = \begin{bmatrix} g_0 & -g_{n-1} & -g_{n-2} & \cdots & -g_1 \\ g_1 & g_0 & -g_{n-1} & \cdots & -g_2 \\ g_2 & g_1 & g_0 & \cdots & -g_3 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ g_{n-1} & g_{n-2} & g_{n-3} & \cdots & g_0 \end{bmatrix} \in \mathbb{Z}^{n \times n}.$$

Therefore, we have $\psi(g) \cdot \mathbf{h}$ as the coefficient embedding of $h(x) \cdot g(x) \in R$. In fact, the map $\psi : R \rightarrow \mathbb{Z}^{n \times n}$ taking a polynomial and mapping it to its anti-cyclic matrix is a (another) ring embedding, from R to $n \times n$ integer matrices. The same holds when we consider the quotient R_q , except now the previous is expressed via operations modulo q . An important feature of the coefficient embedding is that scalars in the ring, $\alpha \in \mathbb{Z} \subset R$ (or $\mathbb{Z}_q \subset R_q$), embed as the scalar matrix, $\psi(\alpha) = \alpha \cdot \mathbf{I}_n \in \mathbb{Z}^{n \times n}$ (or $\mathbb{Z}_q^{n \times n}$). Further, vectors over R or R_q have an entry-wise coefficient embedding: $\mathbf{v} = (v_0, \dots, v_{l-1}) \in R^l \mapsto (v_{0,0}, \dots, v_{0,n-1}, \dots, v_{l-1,0}, \dots, v_{l-1,n-1}) \in \mathbb{Z}^{nl}$. We suggest the reader regularly traverse between the rings (R or R_q) and their coefficient embedding, in \mathbb{Z}^n or \mathbb{Z}_q^n , as well as their representations as linear transformations, while reading the remainder of this section. Lastly, our norm will be the normal Euclidean norm under the the coefficient embedding, $\|\mathbf{x}\| = \|\psi(\mathbf{x})\|_2$.

Ring-SIS and Ring-ApproxSIS. The ring-SIS problem [Mic02] is: given a uniformly random $\mathbf{a} \in R_q^m$, find a short ($\|\mathbf{x}\| \leq \beta$), non-zero vector $\mathbf{x} \in R^m$ such that $\mathbf{a}^t \mathbf{x} = 0 \in R_q$.

Definition A.1 (Ring Approximate ISIS). For any $m, q \in \mathbb{N}$ and $\alpha, \beta \in \mathbb{R}$, define the ring approximate inhomogeneous short integer solution problem $R\text{-ApproxISIS}_{m,q,\alpha,\beta}$ as follows: Given $\mathbf{a} \in R_q^m$, $y \in R_q$, find a vector $\mathbf{x} \in R^m$ such that $\|\mathbf{x}\| \leq \beta$, and there is an element $z \in R$ satisfying

$$\|z\| \leq \alpha \quad \text{and} \quad \mathbf{a}^t \mathbf{x} = y + z \pmod{q}.$$

Further, an approximate trapdoor is any string τ allowing one to efficiently find a short \mathbf{x} such that $\mathbf{a}^t \mathbf{x} \approx y$ given $y \in R_q$.

A.1 Approximate Trapdoor Scheme in the Ring Setting

Here we sketch the approximate trapdoor scheme in the ring setting.

Algorithm 4: RINGGSAMP.CUT(u, σ)

Input: $u \in R_q, \sigma \in \mathbb{R}^+$
Output: $\mathbf{z} \in R^{k-l}$ satisfying $\mathbf{F}\mathbf{z} \approx u \in R_q$
1 **for** $i \in \{0, \dots, n-1\}$ **do**
2 $\mathbf{x}_i \leftarrow \text{GSAMP.CUT}(u_i, \sigma)$
3 **end**
4 Let $\mathbf{x} := (\mathbf{x}_1, \dots, \mathbf{x}_{n-1}) \in \mathbb{Z}^{n(k-l)}$.
5 **return** $\mathbf{z} := \mathbf{P}_f \mathbf{x} \in R^{k-l}$.

Figure 6: The pseudocode for the approximate G-lattice sampler in the ring setting. Note, we store the output as a vector of polynomials.

Approximate G-Sampling. Our approximate gadget is now the vector

$$\mathbf{f} = (b^l, \dots, b^{k-1}) \in R_q^{k-l},$$

and under the coefficient embedding, we have

$$\psi(\mathbf{f}) = \mathbf{F} = [b^l \mathbf{I}_n, \dots, b^{k-1} \mathbf{I}_n] \in \mathbb{Z}_q^{n \times n(k-l)}.$$

Next, let \mathbf{P}_f be the permutation matrix satisfying $\mathbf{F}\mathbf{P}_f = \mathbf{I}_n \otimes (b^l, \dots, b^{k-1}) \in \mathbb{Z}_q^{n \times n(k-l)}$ (mapping the coefficient embedding of \mathbf{f} to the approximate gadget used in Section 4). This reduces the task of approximate G-sampling in the ring setting to approximate G-sampling as done in Section 4, given as the pseudocode of GSAMP.CUT in Algorithm 1. The pseudocode for the approximate G-sampler in the ring setting is given in Figure 6.

Let $\mathbf{z} \leftarrow \text{RINGGSAMP.CUT}(u, \sigma)$, and let $\mathbf{u} \in \mathbb{Z}_q^n$ be u 's coefficient embedding. Then, we have $\mathbf{z} = \mathbf{P}_f \mathbf{x}$ where $\mathbf{x} \in \mathbb{Z}^{n(k-l)}$ satisfies $(\mathbf{I}_n \otimes [b^l, \dots, b^{k-1}])\mathbf{x} \approx \mathbf{u} \pmod{q}$. Notice under the coefficient embedding, we have

$$\mathbf{F}\psi(\mathbf{z}) = (\mathbf{I}_n \otimes [b^l, \dots, b^{k-1}])\mathbf{P}_f^t \mathbf{z} = (\mathbf{I}_n \otimes [b^l, \dots, b^{k-1}])\mathbf{x},$$

and this is the same distribution as in Section 4.

Public key Our trapdoor is a random matrix with polynomial entries, $\mathbf{R} \in R^{2 \times (k-l)}$, whose coefficients are chosen independently from a distribution χ over \mathbb{Z} , usually a discrete Gaussian. Let \mathbf{R} have $\mathbf{r}_1, \mathbf{r}_2$ as its first and second rows, respectively. Then, the public key is

$$\mathbf{a}^t = [(1, a)|\mathbf{f} - a\mathbf{r}_2 - \mathbf{r}_1] \in R_q^{1 \times 2+k-l}$$

where $a \in R_q$ is uniformly sampled.

Preimage sampling. Now given a coset $u \in R_q$, the procedure to sample a short $\mathbf{x} \in R^{2+k-l}$ such that $\mathbf{a}^t \mathbf{x} \approx u \in R_q$ is clear.

1. Sample a perturbation $\mathbf{p} \leftarrow D_{R^{2+k-l}, \sqrt{\Sigma_p}}$ where

$$\Sigma_p := s^2 \mathbf{I} - \sigma^2 \begin{bmatrix} \mathbf{R}\mathbf{R}^t & \mathbf{R}^t \\ \mathbf{R} & \mathbf{I} \end{bmatrix}$$

is now a block matrix with anti-cyclic blocks.

2. Compute the shifted coset, $v = u - \mathbf{a}^t \mathbf{x} \in R_q$ and sample $\mathbf{z} \leftarrow \text{RINGGSAMP.CUT}(v, \sigma)$.
3. Return the approximate preimage sample:

$$\mathbf{y} = \mathbf{p} + \begin{bmatrix} \mathbf{R} \\ \mathbf{I} \end{bmatrix} \mathbf{z} \in R^{2+k-l}.$$

The perturbation sampling techniques previously used in the ring versions of MP12, [GM18, DP16, DN12], immediately apply to the above procedure.

A.2 Analogous Lemmas and Theorem

Here we state the analogous lemmas needed to prove Theorem A.2 and either sketch their proofs or their proofs are immediate from the coefficient embedding.

Theorem A.2. *There exists probabilistic, polynomial time algorithms $\text{RINGAPPROX.TRAPGEN}(\cdot)$ and $\text{RINGAPPROX.SAMPLEPRE}(\cdot, \cdot, \cdot, \cdot)$ satisfying the following.*

1. $\text{RINGAPPROX.TRAPGEN}(\lambda)$ takes as input a security parameter λ and returns a matrix-approximate trapdoor pair $(\mathbf{a}^t, \mathbf{R}) \in R_q^{1 \times (2+k-l)} \times R^{2 \times (k-l)}$.
2. Let \mathbf{a} be generated with an approximate trapdoor as in Section A.1, $\text{RINGAPPROX.a}^{-1}(\cdot)$ denote the approximate preimage sampling algorithm, $\text{RINGAPPROX.SAMPLEPRE}(\mathbf{a}, \mathbf{R}, s, \cdot)$. The following two distributions are statistically indistinguishable:

$$\{(\mathbf{a}, \mathbf{y}, u, e) : u \leftarrow U(R_q), \mathbf{y} \leftarrow \text{RINGAPPROX.a}^{-1}(u), e = u - \mathbf{a}^t \mathbf{y} \pmod{q}\}$$

and

$$\{(\mathbf{a}, \mathbf{y}, \mathbf{a}^t \mathbf{y} + e \in R_q, e \pmod{q}) : \mathbf{y} \leftarrow D_{R^m, s}, e \leftarrow D_{R, \sigma \sqrt{(b^{2l}-1)/(b^2-1)}}\}$$

for any $\sigma \geq \sqrt{b^2 + 1} \omega(\sqrt{\log n})$ and $s \geq \sqrt{\sigma^2 s_1^2(\mathbf{R}) + \omega(\sqrt{\log n})}$. Furthermore, in the second distribution, \mathbf{a} is computationally indistinguishable from random assuming RLWE over power of two cyclotomic rings.

The algorithms $\text{RINGAPPROX.TRAPGEN}(\cdot)$ and $\text{RINGAPPROX.SAMPLEPRE}(\cdot, \cdot, \cdot, \cdot)$ are implicit from the methods sketched below.

Now, consider an integer $n \geq 1$ ring dimension, an integer modulus $q \geq 1$, an integer base $b \geq 1$ as fixed and $k = \lceil \log_b q \rceil$. Then, we have the gadget matrix over the ring as $\mathbf{g} = [1, b, \dots, b^{k-1}] \in R_q^k$. Notice that $\mathbf{g} \in R^k$ embeds as

$$[\mathbf{I}, b\mathbf{I}, \dots, b^{k-1}\mathbf{I}] \in \mathbb{Z}_q^{n \times nk}$$

under the coefficient embedding. Let $\mathbf{G} := \mathbf{I}_n \otimes [1, b, \dots, b^{k-1}] \in \mathbb{Z}_q^{n \times nk}$ be the normal G-matrix, and \mathbf{P} be the permutation matrix satisfying $[\mathbf{I}, b\mathbf{I}, \dots, b^{k-1}\mathbf{I}] = \mathbf{G}\mathbf{P}$. This leads us to the following fact.

Fact A.3. *Let $u \in R_q$, $\Lambda_u^\perp(\mathbf{g}^t) = \{\mathbf{x} \in R^k : \mathbf{g}^t \mathbf{x} = u \in R_q\}$. Then under the coefficient embedding, we have $\Lambda_u^\perp(\mathbf{g}^t) = \mathbf{P}^t \cdot \Lambda_{\mathbf{u}}^\perp(\mathbf{G})$ where $\mathbf{u} \in \mathbb{Z}_q^n$ is the embedding of $u \in R_q$.*

The proof is immediate from the definition of coefficient embedding. As a corollary, we see that for every fixed $\epsilon > 0$, the two smoothing parameters are equal, $\eta_\epsilon(\Lambda_q^\perp(\mathbf{g}^t)) = \eta_\epsilon(\Lambda_q^\perp(\mathbf{G}))$. In addition, Fact A.3, gives us the following lemma.

Lemma A.4. *For any $\sigma \geq \eta_\epsilon(\Lambda^\perp(\mathbf{g}^t))$, the random process of first choosing $u \leftarrow U(R_q)$ then returning $\mathbf{x} \leftarrow D_{\Lambda_u^\perp(\mathbf{g}^t), \sigma}$ is statistically close to sampling $\mathbf{x} \leftarrow D_{R^k, \sigma}$.*

And because we can implement an algorithm to sample $D_{\Lambda_u^\perp(\mathbf{g}^t), \sigma}$ (with support \mathbb{Z}^{nk} or R^k under the embedding) with n independent calls to the normal G-sampler [MP12, GM18] (with support over \mathbb{Z}^k), we have the analogous lemma to Lemma 4.3. In short, the coefficient embedding gives us that the error is distributed as $\mathbf{e} = (\mathbf{I}_n \otimes \mathbf{L}) \cdot D_{\mathbb{Z}^{nk}, \sigma} \approx D_{R, \sigma'}$ with $\sigma' = \sigma \sqrt{(b^{2l} - 1)/(b^2 - 1)}$ and $\mathbf{L} = (1, b, \dots, b^{l-1}, \mathbf{0}) \in \mathbb{Z}^k$.

Next, we argue the analogous lemma to Lemma 4.4 is true. Let $\mathbf{R} \in R^{2 \times (k-l)}$ be the trapdoor to our ring version of the approximate trapdoor scheme. Under the coefficient embedding, $\psi(\mathbf{R})$ is a block matrix with anti-cyclic blocks (we use the trapdoor as a linear transformation). Following the steps of the proof of Lemma 4.2 in the ring setting, we note that we can first permute coordinates under the coefficient embedding to match the coordinates used in the proof of Lemma 4.2. Further, the structure of the trapdoor \mathbf{R} does not affect the proof. Finally, the same rotation \mathbf{Q} is used over R^{m+k-l} , where $m = 2 + k - l$, and R^{m+k-l} is identified with $\mathbb{Z}^{nm+n(k-l)}$ under the coefficient embedding, and the proof follows. We summarize with the following lemma. Without loss of generality, we can again permute the columns of the gadget matrix's embedding to $\mathbf{G}' := [\mathbf{M}|\mathbf{F}]$. This allows us to express truncation and convolution as a simple linear transformation: $\mathbf{y} = \mathbf{L}(\mathbf{p}, \mathbf{x}_2)$ for $\mathbf{L} := [\mathbf{I}_m|\mathbf{R}']$.

Lemma A.5. *For any $\epsilon > 0$, $\sqrt{\Sigma_p} \geq \eta_\epsilon(R^k) = \eta_\epsilon(\mathbb{Z}^{nk})$, and $\sigma \geq \eta_\epsilon(R^k)$ (where Σ_p is now a block matrix with anti-cyclic blocks), $\mathbf{L} \cdot D_{R^{m+k}, \sqrt{\Sigma_p \oplus \sigma^2 \mathbf{I}_k}}$ is statistically close to $D_{R^m, s}$.*

Finally, the proof of Theorem A.2 follows from the previous lemmas and is proved as in Section 4.