

ON GAUSSIAN SAMPLING, SMOOTHING PARAMETER

• • •

APPLICATION TO LATTICE SIGNATURES

THOMAS ESPITAU^{*}, ALEXANDRE WALLET^{*}, AND YANG YU[†]

ABSTRACT. We present a general framework for polynomial-time lattice Gaussian sampling. It revolves around a systematic study of the discrete Gaussian measure and its samplers under *extensions* of lattices; we first show that given lattices $\Lambda' \subset \Lambda$ we can sample efficiently in Λ if we know how to do so in Λ' and the quotient Λ/Λ' , *regardless* of the primitivity of Λ' . As a direct application, we tackle the problem of domain extension and restriction for sampling and propose a sampler tailored for lattice *filtrations*, which can be seen as a broad generalization of the celebrated Klein’s sampler. Then, we demonstrate how to sample using a change of bases, or even switching the ambient space, even when the target lattice is not represented as full-rank in the ambient space. We show how to correct the induced distortion with the “convolution-like” technique of Peikert (Crypto 2010) (which we encompass as a byproduct). Since our framework aims at modularity and leverage the combinations of smaller samplers to build new ones, we also propose ad-hoc samplers for the so-called *root lattices* A_n, D_n, E_n as base cases, extending the state-of-the-art for root lattice sampling, which was limited to \mathbf{Z}^n . We also show how our framework blends with the so-called *king* construction and provides a sampler for the remarkable Leech and Barnes-Wall lattices.

As a by-product, we obtain novel, quasi-linear samplers for prime and smooth conductor (as $2^\ell 3^k$) cyclotomic rings, achieving essentially optimal Gaussian width. In a practice-oriented application, we showcase the impact of our work on hash-and-sign signatures over NTRU lattices. In the best case, we can gain around 200 bytes (which corresponds to an improvement greater than 20%) on the signature size. We also improve the new gadget-based constructions (Yu, Jia, Wang, Crypto 2023) and gain up to 110 bytes for the resulting signatures.

Lastly, we sprinkle our exposition with several new estimates for the smoothing parameter of lattices, stemming from our algorithmic constructions and by novel methods based on series reversion.

CONTENTS

1. Introduction	4
2. Algebraic and computational background	10
2.1. Euclidean lattices	10
2.2. Discrete Gaussian distributions	12
2.3. On root lattices.	12
2.4. Identifications between cyclotomic ideals and root lattices	14
3. Algebraic extensions and sampling	15
3.1. Gaussian measures over short sequences of groups	15
4. Generic applications of the short-sequence sampler	21
4.1. Domain extension and restriction	21
4.2. A filtration sampler	22
4.3. Recovering some known samplers	25
5. The linear sampler	27
5.1. Smoothing parameters and linear transformations	27
5.2. Sampling by linear transformation	27
5.3. Example of elementary instantiations	29
6. Application: sampling in tensor lattices	30
6.1. Tensors and related bounds.	30
6.2. A sampling algorithm for tensor lattices.	31
7. An exact calculation of the smoothing parameter.	32
8. Sampling in remarkable lattices	35
8.1. Sampling in root lattices.	35
8.2. The <i>king</i> sampler	44
9. Application I: Improved Samplers for Mitaka	47
9.1. Hybrid sampling and representation of cyclotomic numbers.	47
9.2. Sampling over cyclotomic fields of conductor $2^\ell \cdot 3^k$	49

ON GAUSSIAN SAMPLING, SMOOTHING PARAMETER	3
9.3. Sampling over prime cyclotomic fields	50
10. Application II: New Compact Lattice Gadgets	51
10.1. The Yu-Jia-Wang compact gadget framework	52
10.2. Compact gadget from the E_8 lattice	53
Towards improving LIP based schemes.	53
References	54

1. INTRODUCTION

For the last few decades, lattices have proved themselves to be a cornerstone of modern cryptography, allowing the development of feature-rich schemes, including digital signatures [9, 30, 15], identity-based encryption [18], functional encryption [2], (non-interactive) zero-knowledge proofs [28] and last but not least fully homomorphic encryption [17, 4]. A common denominator of many such schemes revolves around the ability of sampling from the so-called *discrete Gaussian distribution* over a given lattice Λ . Given a center \mathbf{c} in the ambient space $\Lambda_{\mathbf{R}}$ and a “width” s – which is essentially the standard deviation by analogy with the normal distribution – the distribution $\mathcal{D}_{\Lambda, \mathbf{c}, s^2}$ assigns the vector $\mathbf{v} \in \Lambda$ the probability proportional to the Gaussian function $\exp(-\pi\|\mathbf{v} - \mathbf{c}\|^2/s^2)$. Remark that this distribution only depends on the lattice and not on the basis used to represent it. In this sense it does not leak any information about a possible secret basis: this “zero-knowledge” property accounts for its utility in cryptography.

For specific lattices such as \mathbf{Z}^n or lattices stemming from some trapdoor sampling as in [23], *ad-hoc* approaches are commonly used. In comparison, to sample in an *a priori arbitrary* lattice, two polynomial time samplers are well-known and widely used in constructions and beyond: the so-called *Klein* sampler (or *GPV* sampler) [18] and the *Peikert* sampler [26], both having different advantages and drawbacks. The former is a *sequential* sampler: the algorithm performs adaptive iterations of sampling in projected lines, where the choices made in each iteration affect the values used in the next. It is rather costly and imposes to work with the Gram-Schmidt orthogonalization of the input basis. The latter is a naturally parallel sampler, reducing the problem of sampling in Λ to sample the coefficients of the desired sample on the input basis. This “change of basis” induces a distortion, blurred by convolving with a sufficiently wide perturbation. It is faster than Klein’s sampler at the price of slightly worse quality, in the sense that the minimal sampleable width is larger. Note that these two algorithms correspond to the randomization of two famous polynomial time oracles for the (approximate) Closest Vector Problem from Babai [3]: the Klein sampler corresponds to the *nearest plane* algorithm and the Peikert sampler to the *rounding* algorithm. Fine-tuning such algorithms is one of the main tasks for designers of signatures in the hash-then-sign framework of [18, 10].

On hash-and-sign digital signatures. Designing, selecting, and analyzing quantum-resistant schemes is the main goal of the ongoing NIST standardization effort for post-quantum cryptography. In July 2022, NIST eventually announced four post-quantum algorithms to be standardized. For signatures, two of the three selected algorithms are lattice-based, FALCON [30] and DILITHIUM [9], epitomizing two known classes of lattice signatures: hash-and-sign and Fiat-Shamir with aborts. Recently, Espitau et al. designed an alternative approach to FALCON, called MITAKA [15]. As an attractive advantage, MITAKA can be instantiated over arbitrary cyclotomic fields, conveniently allowing it to reach all NIST security levels. MITAKA relies on a so-called hybrid sampler [29], which acts as Klein’s sampler at the level of the NTRU module and calls the Peikert sampler to sample within this ring. For power-of-two cyclotomics, this approach is sufficient, as the sampling of the ring of integers amounts to sampling in a square lattice \mathbf{Z}^n . However, for the other cyclotomic rings considered in [15], this induces a non-negligible quality loss, thus a slight degradation in security. Recently, the scheme Hawk [11] presented a new hash-and-sign signature tied to lattices, relying on a more recent—yet natural—cryptographic assumption, the so-called *lattice isomorphism problem*. This latter scheme raises the non-trivial question of the possibility of sampling efficiently in lattices with remarkable packing properties, such as the Leech or the Barnes-Walls lattices. Interestingly, the same question was also independently triggered in the new design of lattice gadgets of [32] to improve the efficiency and compactness of trapdoors. All in all, the quest for new efficient samplers, especially for remarkable lattices, has interesting consequences in the realm of cryptographic design.

Contributions. In this work, we aim at going beyond this Klein/Peikert dichotomy for polynomial time sampling. We showcase a general framework based on a systematic study of the discrete Gaussian distribution under *extensions*: algebraic extensions through short exact sequences and metric extensions through linear transformations. This framework allows us to build new samplers over extensions or restrictions of domains in which we already know how to sample. Our abstract samplers correspond to effective versions of general bounds on the smoothing parameter of lattices: this correspondence is a unifying thread in all our exposition. To complete our modular framework, we also provide ad-hoc samplers of essentially optimal widths for root lattices, to use them as fundamental blocks to instantiate more involved samplers. Optimality is deduced from a new theoretical bound on the smoothing parameter obtained from the relation between the Gaussian function over a lattice and its *theta* series. When the kissing number and minimum of the (dual) lattice are known, it gives a tighter bound on

the smoothing, unconditionally on the value of ε – unfortunately, it is unlikely to get this information for an arbitrary lattice. As an application, we obtain novel, optimal, and efficient samplers over cyclotomic rings of prime and smooth conductors and optimized trapdoors in the spirit of [32]. The technical details of our contributions are as follows.

Exploiting the decomposition over short exact sequences. Given a lattice Λ and one of its sublattices Λ' , we can associate the short exact sequence of \mathbf{Z} -modules:

$$0 \longrightarrow \Lambda' \longrightarrow \Lambda \longrightarrow \Lambda/\Lambda' \longrightarrow 0.$$

Note that in this sequence, the quotient Λ/Λ' is not necessarily a lattice¹ itself, and as such, Λ cannot be identified as a lattice to $\Lambda' \oplus \Lambda/\Lambda'$. We show how to deal with this extension of groups to extend samplers for Λ' and Λ/Λ' into a sampler for Λ , for standard deviations above the smoothing parameters of the Λ' component. In particular, we identify precisely the projection of the Gaussian measure onto the quotient, recovering the known situation where Λ' is either full-rank or primitive. This construction translates into a simple bound on the smoothing parameter, namely

$$\eta_{3\varepsilon}(\Lambda) \leq \max\left(\eta_\varepsilon(\Lambda'), \eta_\varepsilon\left(\Lambda/\Lambda'\right)\right),$$

where the notion of smoothing is generalized to accommodate non-lattice quotients. Note that the choice of the sublattice is arbitrary here. This suffices, for instance, to deal with the problem of domain extension and restriction of samplers: given a sampler over Λ , how can one extend it to an overlattice or restrict it to a sublattice?

A filtered sampler. A filtration of a lattice is an increasing sequence of lattices $0 \subset \Lambda_1 \subset \dots \subset \Lambda_k = \Lambda$. Iterating the previous construction gives us a generic sampler for Λ . Namely, we have a first short exact sequence stemming from the filtration:

$$0 \longrightarrow \Lambda_1 \longrightarrow \Lambda \longrightarrow \Lambda/\Lambda_1 \longrightarrow 0,$$

and by our sampler over sequences, we can efficiently sample in Λ if we know how to sample in both Λ_1 and Λ/Λ_1 . However, we can remark that quotienting by Λ_1 induces a filtration

¹Generally, the quotient is a product of the *torsion part*, which is a finite abelian group and its *free part*, which corresponds to a lattice too. Even when the quotient is torsion-free, Λ does not identify to $\Lambda' \oplus \Lambda/\Lambda'$ as lattices in general.

$0 \subset \Lambda_2/\Lambda_1 \subset \cdots \subset \Lambda_k = \Lambda/\Lambda_1$. Hence, we can recursively apply this technique and devise a sampler for Λ from samplers over $(\Lambda_{i+1}/\Lambda_i)_i$. This approach yields a natural generalization of Klein's sampler (as presented in [18]), which corresponds to the particular case where $\text{rk}(\Lambda_i) = i$ for all $1 \leq i \leq \text{rk}(\Lambda)$, and the successive quotients correspond to the Gram-Schmidt orthogonalization. Expectedly, we obtain a bound on the smoothing parameter of Λ in terms of the smoothing parameter of these quotients, generalizing that of [18]:

$$\eta_\varepsilon(\Lambda) \leq \max_{1 \leq i \leq k} \eta_{\frac{\varepsilon}{k+1}}(\Lambda_i/\Lambda_{i-1}).$$

In a later section, we show how this abstract sampler and its designated bound can lead to significant improvements over the Klein-Peikert dichotomy on a concrete example.

A linear sampler. *Change of basis* is a natural technique in linear algebra allowing to re-express sets of linear equations in more congenial forms, by looking at the coordinates of a linear space under a different basis. It is a deep principle undertaking numerous aspects of numerical algorithm, whether by making incremental changes (like in Gaussian elimination or lattice reduction), or in one take (e.g., computing the Discrete Fourier transform representation). Unsurprisingly, we can apply it to discrete Gaussian sampling as well². Hence, from a high-level point of view, one can design a Gaussian sampler in a given lattice Λ as long as one can sample discrete Gaussians in the lattice spanned by a (fixed) congenial basis \mathbf{C} , which can even live in a different space. This process amounts to controlling the *distortion* on the Gaussian distribution induced by the change-of-basis procedure, and to smooth it out with a carefully chosen normal³ perturbation. This algorithm encompasses the sampler of Peikert [26], which reduces sampling in a lattice Λ to sampling *spherically* in $\mathbf{Z}^{\text{rk } \Lambda}$ — this can be done coordinate-wise. This construction yields a natural bound on the smoothing parameter, writing a basis \mathbf{B} of Λ as the product \mathbf{TC} :

$$\eta_\varepsilon(\Lambda) \leq s_1(\mathbf{T}) \cdot \eta_\varepsilon(\mathcal{L}(\mathbf{C}))$$

for $s_1(\mathbf{T})$ being the largest singular value of \mathbf{T} . Again, note that the choice of the decomposition is arbitrary (as long as \mathbf{C} is invertible). A generic sampler in tensor lattices $\Lambda_1 \otimes \Lambda_2$ follows almost immediately.

²Of course, change of basis works very well for continuous Gaussians: it simply amounts to matrix-vector multiplication.

³What matters for proofs is that the perturbation distribution has good convolution properties with Gaussian kernels.

Sampling in remarkable lattices. The previous contributions aim at building a framework for efficient Gaussian sampling, by joining existing samplers through *extensions* (namely module extension for the exact sequence sampler, linear extension for the linear sampler and tensor extension for the tensor sampler). It means that we need to be able to sample in some base cases to fully instantiate these higher-order constructions. We thus introduce a set of ad-hoc samplers for some of the so-called root lattices (A_n lattices, the face-centered lattices D_n , the Gosset lattice E_8) emerging in many contexts. They are, for example, well-known for their outstanding geometric properties, e.g., enjoying quasi-linear decoding [5, 6], or their appearance in more mathematical topics such as the classification of Lie algebras. In particular, our samplers rely on their well-understood structures and exceptional isomorphisms between them, coming from the latter topic, and we reach standard deviations quite close to the smoothing of these lattices. Generally, we add another tribute to the deep connexions of these remarkable lattices with coding theory, as we combine our algebraic framework with the *k*ing construction [8] to devise samplers in the Leech and low-dimensional Barnes-Wall lattices. The technique allows as a byproduct to construct samplers in parity-check-like lattices.

Cryptographic impact. To showcase our framework in a cryptographic context, we demonstrate how to instantiate various samplers over some structured lattices. There are well-known identifications between certain ideals in prime cyclotomic rings and A_{p-1} lattices (or their duals), already subject to algorithmic works [20, 14]. Cyclotomic rings of smooth conductors can also identify as (direct sums of) prime cyclotomic rings. We exploit our ad-hoc samplers to devise novel samplers in cyclotomic rings: our result combines quasi-linear efficiency and optimal Gaussian width. To our knowledge, all previous approaches reached worse Gaussian widths, and at best equivalent efficiency.

We also detail the implication for the design of hash-and-sign signatures, where the ability to sample efficiently and precisely is crucial for the security and bandwidth of the scheme. We compare our variations with the recently proposed and state-of-the-art Falcon [30] and Mitaka [15] signatures. In particular, we show how to design hash-and-sign signatures more tightly on smooth cyclotomic fields, giving more security (around 20 bits in both classical and quantum regimes) and slightly shorter signatures for free compared to [15] (although bitsize is not the focus of this work). More interestingly, we show how to implement them on *prime*

cyclotomics, allowing a very tight choice of parameter selections. At a high-level, our results are also satisfying in the sense that they not only increase the security level for prime cyclotomics compared to [15], they also show a more regular growth and behavior of the ratio security level over cyclotomic-conductors compared to [15]. Then, we also give new instantiations of the recent framework of [32] for compact gadget-based sampling with a target lattice constructed as a tensor of the root E_8 and \mathbf{Z}^n . We again get slightly shorter signatures for free (110 bytes shorter) for higher security, both in the classical and quantum settings.



2. ALGEBRAIC AND COMPUTATIONAL BACKGROUND

General notation. The bold capitals \mathbf{Z} , \mathbf{Q} , and \mathbf{R} refer respectively to the ring of integers, the field of rational and real numbers. Given a real number x , the integral roundings *floor*, *ceil* and *round to the nearest integer* are denoted respectively by $\lfloor x \rfloor$, $\lceil x \rceil$, $\lfloor x \rceil$. Let \ln denote the natural logarithm. For a real-valued function f and a countable set S , we write generically $f(S) = \sum_{x \in S} f(x)$ assuming that this sum is absolutely convergent. Vectors and matrices are understood column-wise. For \mathbf{A}, \mathbf{B} two matrices, we write $[\mathbf{A}, \mathbf{B}]$ for the concatenation of the columns from \mathbf{A} and \mathbf{B} . The transpose of a matrix \mathbf{T} is \mathbf{T}^t and if \mathbf{T} is non singular, its pseudo-inverse is $\mathbf{T}^* = (\mathbf{T}^t \mathbf{T})^{-1} \mathbf{T}^t$.

2.1. Euclidean lattices. A (real) *lattice* Λ is a finitely generated free \mathbf{Z} -module, endowed with a Euclidean norm $\|\cdot\|$ on the real vector space $\Lambda_{\mathbf{R}} := \Lambda \otimes_{\mathbf{Z}} \mathbf{R}$. By definition, there exists a finite family $(\mathbf{b}_1, \dots, \mathbf{b}_n) \in \Lambda^n$ of linearly independent elements such that $\Lambda = \bigoplus_{i=1}^n \mathbf{b}_i \mathbf{Z}$, and we write $\Lambda = \mathcal{L}(\mathbf{B})$, with the matrix $\mathbf{B} = [\mathbf{b}_1, \dots, \mathbf{b}_n]$. It is called a *basis* of Λ . Every basis has the same number of elements $\text{rk}(\Lambda)$, called the *rank* of the lattice. We let $\lambda_1(\Lambda)$ be the Euclidean norm of a shortest non-zero vector in Λ . The volume is $\det \Lambda = \sqrt{\det \mathbf{B}^t \mathbf{B}}$, for any basis \mathbf{B} of Λ .

In this work, when dealing with lattices embedded in \mathbf{R}^n , we only consider the standard Euclidean norm, corresponding to the canonical inner product $\langle \cdot, \cdot \rangle$, but we stress that most of our algorithms are agnostic to the choice of the norm. The dual of a lattice Λ is the lattice $\Lambda^\vee = \{\mathbf{x} \in \Lambda_{\mathbf{R}} \mid \langle \mathbf{x}, \mathbf{v} \rangle \in \mathbf{Z}, \forall \mathbf{v} \in \Lambda\}$, and we always endow it with the same norm as Λ . If Λ is a full-rank lattice of basis \mathbf{B} , then \mathbf{B}^{-t} is a basis of Λ^\vee ; if it is not full rank, $\mathbf{B}(\mathbf{B}^t \mathbf{B})^{-1}$ is a basis of Λ^\vee .

2.1.1. Orthogonality. For a subspace $V \subset \Lambda_{\mathbf{R}}$, let $V^\perp = \{\mathbf{y} \in \Lambda_{\mathbf{R}} \mid \langle \mathbf{y}, \mathbf{v} \rangle = 0, \forall \mathbf{v} \in V\}$ be the orthogonal. Let π_{V^\perp} denote the orthogonal projection onto V^\perp equipped with the restriction of the norm to that space. If \mathbf{P} is a matrix representation of π_{V^\perp} , we have $\mathbf{P}^2 = \mathbf{P}$ and $\mathbf{P}^t = \mathbf{P}$. Given a basis $\mathbf{B} = (\mathbf{b}_1, \dots, \mathbf{b}_n)$ of a lattice Λ , we denote its Gram-Schmidt orthogonalization by $\mathbf{B}^* = (\mathbf{b}_1^*, \dots, \mathbf{b}_n^*)$, where $\mathbf{b}_i^* = \pi_{(b_1, \dots, b_{i-1})^\perp}(\mathbf{b}_i)$.

2.1.2. Sublattices, quotient lattices. Let $(\Lambda, \|\cdot\|)$ be a lattice, and let Λ' be a finitely generated submodule of Λ . The restriction of $\|\cdot\|$ to Λ' endows it with a lattice structure : $(\Lambda', \|\cdot\|)$ is called a *sublattice* of Λ . If any basis of Λ' extends into a basis of Λ , then Λ' is called *primitive*. In this case, the quotient Λ/Λ' is endowed with a canonical lattice structure by defining:

$\|\mathbf{v} + \Lambda'\|_{\Lambda/\Lambda'} = \inf_{\mathbf{v}' \in \Lambda'_R} \|\mathbf{v} - \mathbf{v}'\|$. Then, there is an isometry between $(\Lambda/\Lambda', \|\cdot\|_{\Lambda/\Lambda'})$ and $(\pi_{\Lambda'^{\perp}}(\Lambda), \|\cdot\|)$. Effectively, this means we represent quotient lattices by computing the projection of a given basis for Λ . We write $\Lambda = \Lambda' \perp \Lambda''$ to highlight that Λ is the *orthogonal* direct sum of two lattices. In this case, $\pi_{\Lambda'^{\perp}}(\Lambda) = \Lambda''$ and we have an *isometry* $\Lambda \cong \Lambda' \oplus \Lambda/\Lambda'$.

Whether Λ' is primitive or not, the quotient Λ/Λ' always decomposes as a product of its *torsion part* T (finite subgroup of torsion elements) and its *torsion-free* part. Torsion elements in the quotient represent $\mathbf{x} \in \Lambda$ such that $a\mathbf{x} \in \Lambda'$ for some $a \in \mathbf{Z}$, that is, the set $\Lambda \cap \Lambda'_R$. The torsion-free part is itself a lattice: if $\overline{\Lambda'}$ is the (primitive) lattice generated by Λ' and a system of representative for T , it identifies to $\Lambda/\overline{\Lambda'}$, with the quotient norm. It is thus equivalent for Λ' to be primitive and for Λ/Λ' to be torsion-free. When Λ' has full-rank, Λ/Λ' is just the torsion group T . Usual solutions to perform the lift from a coset representative to a lattice point use *Babai's rounding* or *Babai's nearest plane* algorithm.

2.1.3. *Effective lifting.* Given a coset $\mathbf{t} + \Lambda'$ of the quotient Λ/Λ' , we might need to find a representative of this class in Λ . While any element could be theoretically taken, from an algorithmic point of view, we shall take an element of norm somewhat small, so that its coefficients remain polynomial in the input representation of the lattice. Below is the pseudo-code for Nearest Plane based lifting.

Algorithm 1 \therefore Lift (by Babai's nearest plane)

Input: A lattice basis $\mathbf{B} = (\mathbf{b}_1, \dots, \mathbf{b}_n)$ of Λ' in Λ , a vector $\mathbf{t} \in \Lambda$.
Result: A vector $\mathbf{s} \in \Lambda$ in the coset $\mathbf{t} + \Lambda'$.

- 1 Compute the Gram-Schmidt orthogonalization $(\mathbf{b}_1^*, \dots, \mathbf{b}_k^*)$ of \mathbf{B}
- 2 $\mathbf{s} \leftarrow \mathbf{t}$
- 3 **for** $i = k$ **downto** 1 **do** $\mathbf{s} \leftarrow \mathbf{s} - \left\lfloor \frac{\langle \mathbf{s}, \mathbf{b}_i^* \rangle}{\|\mathbf{b}_i^*\|^2} \right\rfloor \mathbf{b}_i$
- 4 **return** \mathbf{s}

2.1.4. *Filtrations.* A *filtration* of a lattice Λ is an increasing sequence of sublattices $\{0\} = \Lambda_0 \subset \Lambda_1 \subset \Lambda_2 \subset \dots \subset \Lambda_k = \Lambda$ where each Λ_i is a primitive sublattice of Λ_{i+1} . Let $\text{rk}(\Lambda_i) = d_i$, then $0 = d_0 < d_1 < d_2 < \dots < d_k = \text{rk}(\Lambda)$. A filtration is called *complete* if $d_i = i$ for all i : for example, any basis of Λ gives a complete filtration. Filtrations are compatible with quotienting: a filtration $(\Lambda_i)_i$ of Λ yields a filtration $(\Lambda_{i+j}/\Lambda_j)_i$ of Λ/Λ_j .

2.2. Discrete Gaussian distributions. Let Σ be a positive definite matrix. We define $\rho_\Sigma(\mathbf{x}) = \exp(-\pi\mathbf{x}^t\Sigma^{-1}\mathbf{x})$ as the Gaussian kernel of covariance Σ . Equivalently, we could call it the standard Gaussian mass for the norm induced by Σ^{-1} . In that case, one sees that a Gaussian function is always *isotropic*, i.e., its value only depends on the designated norm of its input. When $\Sigma = s^2\mathbf{I}_n$, the subscript Σ is shortened⁴ in s^2 and s is called the *width*.

Let now $\Lambda \subset \mathbf{R}^m$ of rank $n \leq m$. The discrete Gaussian distribution over Λ with center $\mathbf{c} \in \Lambda_{\mathbf{R}}$ and covariance $\Sigma \in \mathbf{R}^{m \times m}$ is defined by the density

$$\mathcal{D}_{\Lambda, \mathbf{c}, \Sigma}(\mathbf{x}) = \frac{\rho_\Sigma(\mathbf{x} - \mathbf{c})}{\rho_\Sigma(\Lambda - \mathbf{c})}, \quad \forall \mathbf{x} \in \Lambda.$$

When $\mathbf{c} = \mathbf{0}$, we omit the script \mathbf{c} .

2.2.1. Smoothing parameter. For a lattice Λ and real parameter $\varepsilon > 0$, the *smoothing parameter* $\eta_\varepsilon(\Lambda)$ is the smallest $s > 0$ such that $\rho_{\frac{1}{s^2}}(\Lambda^\vee) \leq 1 + \varepsilon$. When the Gaussian width s exceeds the smoothing parameter, all the lattice cosets have roughly the same mass.

Lemma 2.1 ([26, Lemma 2.4]). *Given a lattice Λ , $\varepsilon \in (0, 1)$ and $\Sigma \succ \eta_\varepsilon(\Lambda)^2\mathbf{I}_n$, then, for any $\mathbf{c} \in \Lambda_{\mathbf{R}}$, $\rho_\Sigma(\Lambda + \mathbf{c}) \in [\frac{1-\varepsilon}{1+\varepsilon}, 1] \rho_\Sigma(\Lambda)$.*

The following result recalls that cosets' mass has exponential decay from the origin. A useful consequence is to express the Gaussian mass by means of a sublattice and its corresponding projection.

Lemma 2.2. *Let $\Lambda \subset \mathbf{R}^m$ be a lattice and $\mathbf{x} \in \mathbf{R}^m$. For $\Sigma \succ 0$, let P be the orthogonal projection onto $\Lambda_{\mathbf{R}}^\perp$, where orthogonality is taken with respect to the inner product $\mathbf{x}^t\Sigma^{-1}\mathbf{y}$. Then we have $\rho_\Sigma(\mathbf{x} + \Lambda) \leq \rho_\Sigma(P(\mathbf{x})) \cdot \rho_\Sigma(\Lambda)$. If moreover Λ is primitive in Λ' , we have $\rho_\Sigma(\Lambda') \leq \rho_\Sigma(\Lambda)\rho_\Sigma(P(\Lambda'))$. The equality case occurs when $\Lambda' = \Lambda \perp P(\Lambda')$.*

2.3. On root lattices. So-called *root lattices* are families of special lattices with nice geometry deriving from root systems. They enjoy, for instance, good decoding properties (see [5, 6], or more recently and closely related to this work, see [14, 31]). Most of their fundamental quantities are well-understood, and general exposition can be found in [22, Chapter 4] or [7]. We only recall here the definitions of three types of root lattices (A_n , E_n and D_n), and highlight some properties of the A_n family.

⁴Most of the prior literature uses s or $\sqrt{\Sigma}$, that is, an analog of standard deviation instead of the covariance.

Definition 2.1 (Root lattices). *For integer $n > 0$, the root lattices A_n, D_n of rank n are respectively defined as*

$$A_n = \{\mathbf{v} \in \mathbf{Z}^{n+1} \mid v_1 + \cdots + v_{n+1} = 0\}, \quad D_n = \{\mathbf{v} \in \mathbf{Z}^n \mid v_1 + \cdots + v_n \in 2\mathbf{Z}\}.$$

We also let

$$E_n = D_n \cup \left(\left(\frac{1}{2}, \dots, \frac{1}{2} \right) + D_n \right).$$

When n is even, E_n is a lattice and an important special case is the Gosset lattice E_8 . If we let $(\mathbf{e}_i)_i$ denotes the canonical basis of \mathbf{R}^{n+1} , the A_n lattices admit the basis $(\mathbf{e}_i - \mathbf{e}_{i+1})_i$. This basis has Gram matrix

$$(1) \quad G_n = \begin{pmatrix} 2 & -1 & 0 & 0 & \cdots & 0 & 0 \\ -1 & 2 & -1 & 0 & \cdots & 0 & 0 \\ \vdots & \ddots & \ddots & \ddots & \ddots & \vdots & \vdots \\ 0 & 0 & 0 & \cdots & -1 & 2 & -1 \\ 0 & 0 & 0 & \cdots & 0 & -1 & 2 \end{pmatrix} \in \mathbf{Z}^{n \times n}.$$

The dual lattice A_n^\vee admits the basis $(\frac{n}{n+1}\mathbf{e}_i - \frac{1}{n+1}\sum_{1 \leq i \neq j \leq n} \mathbf{e}_j)_i$, with associated Gram matrix

$$(2) \quad G_n^\vee = \frac{1}{n+1} \begin{pmatrix} n & -1 & \cdots & -1 & -1 \\ -1 & n & \cdots & -1 & -1 \\ \vdots & \vdots & \vdots & \vdots & \vdots \\ -1 & -1 & \cdots & -1 & n \end{pmatrix}.$$

As we will particularly focus on the A_n lattices, we collect also some of their geometric properties. Note that they span the hyperplane $\mathbf{1}^\perp$, where $\mathbf{1} = (1, \dots, 1)$. Their volume is $\sqrt{n+1}$, and $\lambda_1(A_n) = \sqrt{2}$. Their dual is $A_n^\vee = \pi_{\mathbf{1}^\perp}(\mathbf{Z}^{n+1})$, with $\lambda_1(A_n^\vee) = \sqrt{n/(n+1)}$, and A_n has index $n+1$ in A_n^\vee . Noticeably, A_2 identifies with the famous hexagonal lattice.

2.3.1. On cyclotomic rings. In Section 9, we need some background on cyclotomic rings and their geometry. Most of the used material is relatively standard.

We let \mathbf{Z}_m^* be the multiplicative group of the $\phi(m)$ integers invertible modulo m , where ϕ is Euler's totient function. Let ζ_m be a primitive m -th root of unity. The m -th cyclotomic polynomial $\Phi_m(X) = \prod_{i \in \mathbf{Z}_m^*} (X - \zeta_m^i) \in \mathbf{Z}[X]$. The degree of $\Phi_m(X)$ is $d = \phi(m)$. We call $\mathcal{R}_m = \mathbf{Z}[\zeta_m] \simeq \mathbf{Z}[X]/(\Phi_m(X))$ the m -th cyclotomic ring and $K_m = \mathbf{Q}[\zeta_m]$ the m -th

cyclotomic field. Any $f \in K_m$ can be uniquely written as $f = \sum_{i=0}^{n-1} f_i \zeta_m^i$ with $f_i \in \mathbf{Q}$. The coefficient embedding identifies f to its vector of coefficients (f_0, \dots, f_{d-1}) .

The cyclotomic field K_m has exactly d embeddings fixing elements of \mathbf{Q} . Concretely, the embedding ψ_i for $i \in \mathbf{Z}_m^*$ is defined by $\psi_i(\zeta_m) = \zeta_m^i$. Let $\psi(f) = (\psi_i(f))_{i \in \mathbf{Z}_m^*} \in K_m^d$ be the canonical embedding of $f \in K_m$. We can use it to define the conjugate of f as $f^* = (\overline{\psi_1(f)}, \dots, \overline{\psi_d(f)})$. The trace is the \mathbf{Q} -linear map defined as $\text{Tr}(f) = \sum_{i \in \mathbf{Z}_m^*} \psi_i(f)$. It gives an inner product as $\text{Tr}(fg^*) = \langle \psi(f), \psi(g) \rangle$ and an euclidean norm $\|\psi(f)\|^2 = \text{Tr}(ff^*)$. This product can be extended to vectors in a way that gives a positive definite hermitian form over K_m^k , compatible with the geometry. We only need the case of $k = 2$. The form is defined as $\langle (f, g), (F, G) \rangle = f^*F + g^*G$. In particular, elements $\langle (f, g), (f, g) \rangle = ff^* + gg^*$ have all their embeddings real and positive. Thus for all pairs $(f, g), (F, G)$ of vectors in K_m^2 ,

$$(F, G)^* := (F, G) - \frac{\langle (f, g), (F, G) \rangle}{\langle (f, g), (f, g) \rangle} (f, g)$$

is well-defined. One checks that $\langle (f, g), (F, G)^* \rangle = 0$, which gives a notion of orthogonality over K_m^2 and accordingly, $(F, G)^*$ is the Gram-Schmidt orthogonalization of (F, G) with respect to (f, g) for the trace product.

The Vandermonde matrix associated to the m -th primitive roots $\zeta_{m,1}, \dots, \zeta_{m,\phi(m)}$ of 1 is

$$V_m = \begin{bmatrix} 1 & \zeta_{m,1} & \zeta_{m,1}^2 & \dots & \zeta_{m,1}^{\phi(m)-1} \\ 1 & \zeta_{m,2} & \zeta_{m,2}^2 & \dots & \zeta_{m,2}^{\phi(m)-1} \\ \vdots & & & \dots & \\ 1 & \zeta_{m,\phi(m)} & \zeta_{m,\phi(m)}^2 & \dots & \zeta_{m,\phi(m)}^{\phi(m)-1} \end{bmatrix}.$$

We have $\psi(x) = V_m(x_1, \dots, x_{\phi(m)-1})^t$. When m is a prime p , it is known that the largest and smallest singular values are $s_1(V_p) = \sqrt{p}$ and $s_{p-1}(V_p) = 1$. The corresponding Gram matrix $V_p \overline{V_p}^t$ has determinant $\Delta_K = p^{p-2}$, which is the discriminant (or the squared discriminant sometimes) of the field K_p .

2.4. Identifications between cyclotomic ideals and root lattices. Most of the next facts are well-known, and can be found e.g. in [14, 20, 31]. A first interesting decomposition arises when $m = p^\ell q^k$ for primes $p \neq q$: we have $\mathcal{R}_m = \mathcal{R}_{p^\ell} \otimes \mathcal{R}_{q^k}$. Additionally, we have $\mathcal{R}_{p^k} = \bigoplus_{i=1}^{p^k-1} \mathcal{R}_p$, which is an orthogonal sum under the canonical embedding. Recall the basis $\mathbf{b}_i := \frac{p-1}{p} \mathbf{e}_i - \frac{1}{p} \sum_{1 \leq i \neq j \leq p} \mathbf{e}_j$ of \mathbf{A}_{p-1}^\vee , seen in Section 2.3. The key fact for what follows is the

isometry between \mathcal{R}_p and \mathbf{A}_{p-1}^\vee

$$\tau : \mathcal{R}_p \longrightarrow \mathbf{A}_{p-1}^\vee, \text{ defined by } \tau(\zeta_p^i) = \mathbf{b}_{i+1} \text{ for } 0 \leq i \leq p-2,$$

with normalization for the isometry seen here as $\|\psi(x)\|^2 = p\|\tau(x)\|^2$.

A second identification involves the co-different ideal \mathcal{R}_p^\vee of \mathcal{R}_p : it is the principal ideal $\mathcal{R}_p^\vee = \langle \frac{1-\zeta_p}{p} \rangle$, and acts the dual lattice to \mathcal{R}_p . From what precedes, we get a map $\tau^\vee : \mathcal{R}_p^\vee \rightarrow \mathbf{A}_{p-1}$ sending the dual basis of $1, \zeta_p, \dots, \zeta_p^{p-1}$ for the trace product, also called *decoding basis* by [20], to the basis $(\mathbf{e}_1 - \mathbf{e}_i)_{2 \leq i \leq p}$ of \mathbf{A}_{p-1} . We now have $p\|\psi(x)\|^2 = \|\tau^\vee(x)\|^2$. This map turns out to be mildly less interesting for our purpose. Nevertheless, the ideal $\langle 1 - \zeta_p \rangle$ admits the \mathbf{Z} -basis $(1 - \zeta_p), \dots, \zeta_p^{p-2} - \zeta_p^{p-1}$. Letting $u_i = \zeta_p^i - \zeta_p^{i+1}$, we have:

$$\text{tr}(u_i u_i^*) = \|\psi(u_i)\|^2 = 2p \quad \text{and} \quad \text{tr}(u_i u_j^*) = \langle \psi(u_i), \psi(u_j) \rangle = \begin{cases} -p & \text{if } |i - j| = 1 \\ 0 & \text{else.} \end{cases}$$

We recognize a scaling by p of the Gram matrix (1) of \mathbf{A}_{p-1} when described by the basis $(\mathbf{e}_i - \mathbf{e}_{i+1})_i$. Hence, if we define

$$\phi : \langle 1 - \zeta_p \rangle \longrightarrow \mathbf{A}_{p-1} \text{ by } \phi(u_{i-1}) = \mathbf{e}_i - \mathbf{e}_{i+1} \text{ for } 1 \leq i \leq p-1,$$

it identifies both lattices and we have $\|\psi(x)\|^2 = p\|\phi(x)\|^2$. For modularity, we gather this observation in the following result.

Proposition 1 (Adapted from [31, Chap.1]). *Let p be a prime, ζ_p a primitive p -th root of 1, and ψ the canonical embedding of \mathcal{R}_p . There exists a linear map $\phi : \langle 1 - \zeta_p \rangle \longrightarrow \mathbf{A}_{p-1}$ such that we have $\|\psi(x)\|^2 = p\|\phi(x)\|^2$, for all $x \in \langle 1 - \zeta_p \rangle$.*

The map ψ can be computed using the Vandermonde matrix V_p associated with the p -th primitive roots of 1. We have $s_1(V_p) = \sqrt{p}$ and $s_{p-1}(V_p) = 1$, where s_{p-1} is the smallest singular value. This implies $\frac{1}{p}\|x\|^2 \leq \|\phi(x)\|^2 \leq \|x\|^2$ for all $x \in \langle 1 - \zeta_p \rangle$.

3. ALGEBRAIC EXTENSIONS AND SAMPLING

3.1. Gaussian measures over short sequences of groups. For a lattice Λ , we want to study the behavior of the Gaussian measure $\mathcal{D} = \mathcal{D}_{\Lambda, \mathbf{c}, \Sigma}$ with regards to exact sequences of \mathbf{Z} -modules⁵:

$$(3) \quad 0 \longrightarrow \Lambda' \longrightarrow \Lambda \longrightarrow \Lambda/\Lambda' \longrightarrow 0.$$

⁵We highlight that this is a short sequence of *groups* and not necessarily of *lattices*.

Exactness means that the kernel of each arrow is exactly the image of the arrow preceding it. It implies that Λ' identifies to a submodule of Λ and that the map $\Lambda \rightarrow \Lambda/\Lambda'$ is surjective. We *do not assume* that Λ, Λ' have the same rank, nor that we have an exact sequence of *lattices*, nor that it splits (which would mean that $\Lambda \cong \Lambda' \times \Lambda/\Lambda'$ as \mathbf{Z} -modules).

Recall from Section 2 that Λ/Λ' decomposes as the direct sum $\mathbb{T} \oplus \Lambda'_f$ of its *torsion part* \mathbb{T} and its free part. The free part can be seen as $\Lambda/\overline{\Lambda}'$, where $\overline{\Lambda}'$ is the lattice spanned by Λ' and a set of representative of \mathbb{T} . This denser lattice can be understood as a “primitivation” of Λ' . Hence, $\Lambda/\overline{\Lambda}'$ identifies⁶ to the lattice $\pi_{(\Lambda'_R)^\perp}(\Lambda)$. Let us give a concrete exemple.

3.1.1. *An example of torsion in quotient of lattices.* Let $\Lambda = \mathbf{Z}^2$ be the square lattice of rank 2.

- **Simple Torsion, or “full rank”.** Let $\Lambda' = (2\mathbf{Z})^2$ be its index-4 sublattice consisting of vectors with even coefficients. Then $\Lambda/\Lambda' = \mathbf{Z}^2/(2\mathbf{Z})^2 = (\mathbf{Z}/2\mathbf{Z})^2$ is the Klein group, a finite group of order 4, and its free part is $\{0\}$. Representative for the torsion are for instance $(0, 0), (1, 0), (0, 1)$ and $(1, 1)$.

- **Torsion-free.** Let $\Lambda' \cong \mathbf{Z}$ be its sublattice of rank 1 consisting of vectors with null first coefficient. Then $\Lambda/\Lambda' \cong \mathbf{Z}$ is a lattice of rank 1: the set of vectors with second coefficient equal to zero.

- **Mixed case.** Let $\Lambda' \cong 2\mathbf{Z}$ be its sublattice of rank 1 consisting of vectors with their second coefficient even. Then we have: $\Lambda/\Lambda' \cong \mathbf{Z} \times \mathbf{Z}/2\mathbf{Z}$, encoding the coset space as choosing the parity of the second coefficient (i.e., the torsion part $\mathbf{Z}/2\mathbf{Z}$) and then choosing the first coefficient without conditions (i.e., the free part \mathbf{Z}). A non trivial representative for the torsion is $(0, 1)$, and indeed the lattice spanned by Λ' and $(0, 1)$ is \mathbf{Z} .

The measure \mathcal{D} decomposes into two components, which can then be normalized to probability distributions:

- **the restriction:** over the sublattice Λ' , which identifies as $\mathcal{D}' = \rho/\rho(\Lambda')$.
- **the pushforward:** $\pi_*\mathcal{D}$ onto the quotient Λ/Λ' . By definition, for any witness \mathbf{x} of a Λ' -coset in Λ , we have $\pi_*\mathcal{D}(\mathbf{x}) = \mathcal{D}(\mathbf{x} + \Lambda')$.

⁶The important catch here is about *which* orthogonality we are considering: in our proof, the orthogonality *must* be with respect to the norm induced by the covariance matrix of the target Gaussian, that is, $\mathbf{x} \mapsto \mathbf{x}^t \Sigma^{-1} \mathbf{x}$. This allows us to use that $\mathbf{x}, \mathbf{y} \in \Lambda_{\mathbf{R}}$ such that $\mathbf{x}^t \Sigma^{-1} \mathbf{y} = 0$, we have $\rho_{\Sigma}(\mathbf{x} + \mathbf{y}) = \rho_{\Sigma}(\mathbf{x}) \cdot \rho_{\Sigma}(\mathbf{y})$.

Understanding the latter is the focus of the next lemma. In the lemma below, we distinguish the quotient map $\pi : \Lambda \rightarrow \Lambda/\Lambda'$ from the orthogonal projection $\bar{\pi} := \pi_{(\Lambda'_R)^\perp}$.

Lemma 3.1. *Let $\Lambda' \subset \Lambda$ be lattices and \mathbb{T} the torsion part of Λ/Λ' . If $\Sigma \succ \eta_\varepsilon(\Lambda')$ and $\mathcal{D} = \mathcal{D}_{\Lambda, \mathbf{c}, \Sigma}$, then the pushforward distribution proportional to $\pi_* \mathcal{D}$ is at total variational distance $\frac{\varepsilon}{1-\varepsilon}$ of the distribution defined by $|\mathbb{T}|^{-1} \cdot \mathcal{D}_{\pi(\Lambda), \pi(\mathbf{c}), \Sigma}$, where $|\mathbb{T}|$ is the cardinality of \mathbb{T} .*

For the sake of notational clarity, we restrict to the case of centered distributions but the proof readily adapts to the general case.

Proof. Our first goal is to describe the coset $\mathbf{x} + \Lambda'$, and recall that we denoted $\bar{\pi}$ the orthogonal projection from Λ to $\Lambda/\overline{\Lambda'}$. Consider a section⁷ $s : \Lambda/\overline{\Lambda'} \rightarrow \Lambda$, that is, a linear map such that $\bar{\pi} \circ s = \text{Id}$. From its properties, we see that $\mathbf{x} \in s(\bar{\pi}(\mathbf{x})) + \overline{\Lambda'}$. Let now $\mathbb{T} = \{\mathbf{t} + \Lambda'\}_{\mathbf{t}}$ be a system of representative of the torsion points. Since $\overline{\Lambda'}$ is the disjoint union of cosets $\mathbf{t} + \Lambda'$, there is a unique one such that $\mathbf{x} \in s(\bar{\pi}(\mathbf{x})) + \mathbf{t} + \Lambda'$, and it follows that $\mathbf{x} + \Lambda' = s(\bar{\pi}(\mathbf{x})) + \mathbf{t} + \Lambda'$. By definition and orthogonality, the pushforward of the discrete Gaussian under π therefore acts as

$$(4) \quad \mathcal{D}(\mathbf{x} + \Lambda') = \rho(s(\bar{\pi}(\mathbf{x})) - \bar{\pi}(\mathbf{x}) + \mathbf{t} + \Lambda') \cdot \mathcal{D}(\bar{\pi}(\mathbf{x})).$$

Similarly, the total measure of the quotient Λ/Λ' can be written

$$(5) \quad \mathcal{D}(\pi^{-1}(\Lambda/\Lambda')) = \sum_{(\mathbf{t}, \bar{\pi}(\mathbf{x}))} \rho(s(\bar{\pi}(\mathbf{x})) - \bar{\pi}(\mathbf{x}) + \mathbf{t} + \Lambda') \mathcal{D}(\bar{\pi}(\mathbf{x})).$$

By assumption on the covariance parameter, we are above the smoothing of Λ' , so all the Λ' -cosets have roughly the same Gaussian mass as Λ' . More precisely, taking ratio between Equalities (4) and (5) and using Lemma 2.1, we get our claim:

$$(6) \quad \frac{\pi_* \mathcal{D}(\mathbf{x} \bmod \Lambda')}{\pi_* \mathcal{D}(\Lambda/\Lambda')} \in \left[\frac{1-\varepsilon}{1+\varepsilon}, \frac{1+\varepsilon}{1-\varepsilon} \right] \cdot \frac{1}{|\mathbb{T}|} \cdot \frac{\rho(\bar{\pi}(\mathbf{x}))}{\rho(\bar{\pi}(\Lambda))}.$$

∴

Lemma 3.1 satisfyingly recovers the *extreme* cases which are frequently encountered in the literature:

⁷Such a map always exists: indeed, as $\overline{\Lambda'}$ is primitive, one can always find a sublattice Λ_0 such that $\Lambda = \Lambda_0 \oplus \overline{\Lambda'}$ and the section can be defined by identifying the vectors of a basis of Λ_0 with their projections by $\bar{\pi}$.

- If Λ' is full-rank, we have $|\mathbb{T}| = [\Lambda : \Lambda']$ and the projection sends all points to 0, so that $\pi_*\mathcal{D}$ is statistically close to the uniform distribution over the finite group of Λ' -cosets.
- If Λ' is primitive, the quotient is torsion-free, $\pi = \bar{\pi}$ and we recover that $\pi_*\mathcal{D}$ is essentially the orthogonal projection of the discrete distribution, that is, a discrete Gaussian distribution over $\pi(\Lambda)$.

An interesting subcase happens when an *orthogonal* decomposition $\Lambda = \Lambda' \perp \Lambda''$ is known. We then have a short exact sequence $0 \rightarrow \Lambda' \rightarrow \Lambda \rightarrow \Lambda'' \rightarrow 0$. But now, the Gaussian measure splits perfectly, so that the pushforward is *exactly* the projected distribution.

Lemma 3.2. *Let Λ', Λ'' be two lattices, $\Lambda = \Lambda' \perp \Lambda''$, and π the orthogonal projection onto $\Lambda'_{\mathbf{R}}^\perp$. If $\mathcal{D} = \mathcal{D}_{\Lambda, \mathbf{t}, s^2}$, then we have $\pi_*\mathcal{D} = \mathcal{D}_{\Lambda'', \pi(\mathbf{t}), s^2}$.*

Proof. The assumptions give $\pi(\Lambda) = \Lambda''$. Decompose $\mathbf{z} \in \Lambda$ as $\mathbf{z} = \mathbf{z}' + \pi(\mathbf{z})$ and similarly $\mathbf{t} = \mathbf{t}' + \pi(\mathbf{t})$. We use orthogonality twice: on the one hand, it gives $\rho_{s^2}(\pi(\mathbf{z}) - \mathbf{t} + \Lambda') = \rho_{s^2}(\pi(\mathbf{z}) - \pi(\mathbf{t}))\rho_{s^2}(\Lambda' - \mathbf{t}')$. On the other hand, it also gives $\rho_{s^2}(\Lambda - \mathbf{t}) = \rho_{s^2}(\Lambda' - \mathbf{t}')\rho_{s^2}(\Lambda'' - \pi(\mathbf{t}))$. Taking ratios gives the result. $\ddot{::}$

3.1.2. *Smoothing parameter and short sequences.* The decomposition induced by the quotient translates into a generic bound on the smoothing parameter:

Proposition 2 (Modularity of smoothing parameter). *Let Λ be a lattice and $0 < \varepsilon < \sqrt{17} - 4$, then*

$$\eta_{3\varepsilon}(\Lambda) \leq \min_{\Lambda' \subset \Lambda} \max\left(\eta_\varepsilon(\Lambda'), \eta_\varepsilon\left(\frac{\Lambda}{\Lambda'}\right)\right),$$

where the minimum ranges over all possible sublattices of Λ .

Proof. For completeness, we recall the *Poisson Summation Formula* for Gaussian functions. For any rank n lattice Λ , we have

$$(7) \quad \rho_s(\Lambda) = \frac{s^n}{\det \Lambda} \cdot \rho_{1/s}(\Lambda^\vee).$$

Let Λ' be any sublattice of Λ and $s \geq \max(\eta_\varepsilon(\Lambda'), \eta_\varepsilon(\Lambda/\Lambda'))$. Now, consider the orthogonal projection π of Λ onto the orthogonal space to $\Lambda'_{\mathbf{R}}$, for the standard inner product. Lemma 2.2 gives that $\rho_s(\Lambda) \leq \rho_s(\overline{\Lambda'}) \cdot \rho_s(\pi(\Lambda))$. Using Identity (7) and the fact that $\det(\Lambda) = \det(\pi(\Lambda)) \det(\overline{\Lambda'})$, this is equivalent to $\rho_{1/s}(\Lambda^\vee) \leq \rho_{1/s}(\overline{\Lambda'}^\vee) \rho_{1/s}(\pi(\Lambda)^\vee)$. Because they have the same rank, $\Lambda' \subset \overline{\Lambda'}$ is equivalent to $\overline{\Lambda'}^\vee \subset \Lambda'^\vee$, so we have $\rho_{1/s}(\overline{\Lambda'}^\vee) \leq \rho_{1/s}(\Lambda'^\vee)$. By

assumption on s , this implies $\rho_{1/s}(\Lambda^\vee) \leq (1 + \varepsilon)^2$, and we conclude by noting that our choice of sublattice was arbitrary. $\ddot{::}$

Note that the bound makes appear $\overline{\Lambda'}$ and not Λ' itself in the quotient. The intuition behind this, maybe surprising, detail stems from the fact that the pushforward measure is driven only by Λ' and the torsion-free part of the quotient. Indeed, we can geometrically interpret the smoothing parameter to be the minimal width to *smooth out* the lattice structure i.e., the pushforward over $\Lambda_{\mathbf{R}}/\Lambda$ is the uniform distribution. But then remark that $(\Lambda/\Lambda')_{\mathbf{R}} = (\Lambda/\overline{\Lambda'})_{\mathbf{R}}$ as real spaces, making all the torsion elements geometrically irrelevant w.r.t. the smoothing.

3.1.3. Towards a Gaussian sampler. The bound of Proposition 2 can be turned into a natural sampler built from given samplers over Λ/Λ' and Λ' , or oracles for them. First, sample in the quotient with the appropriate distribution, lift the result to the full lattice, and sample around this point in the sublattice Λ' . Remark that all $\mathbf{x} \in \Lambda$ write uniquely as $\mathbf{x} = \overline{\mathbf{x}'} + \pi(\mathbf{x})$ with $\overline{\mathbf{x}'} \in \overline{\Lambda'}$ and $\overline{\mathbf{x}'} = \mathbf{t} + \mathbf{x}'$, since it also belongs to a unique coset $\mathbf{t} + \Lambda'$ with $\mathbf{t} \in \mathbf{T}$. Above the smoothing of Λ' , sampling according to the pushforward selects such a coset and $\pi(\mathbf{x})$ with essentially the correct distribution. Similarly, above the smoothing of $\pi(\Lambda)$ we cannot really distinguish in which coset of $\pi(\Lambda)$ a Gaussian around $\pi(\mathbf{x})$ belongs.

All-in-all, this strategy leads to Algorithm 2, where we even allow sampling approximatively in the sets Λ' and Λ/Λ' —this will be proved useful to recursively chain calls of this sampler, as we do in Section 4.2.

The proof relies on Lemma 3.1 and the examination of the samples. Two smoothing arguments over Λ' are used: once to apply Lemma 3.1, and once to trade cosets for larger ε .

Theorem 1 (Correctness of the short exact sampler). *When $\Sigma \succ \eta_\varepsilon(\Lambda')$, Algorithm 2 is correct. Moreover, let \mathcal{D} be the distribution of its output. For $\varepsilon < \frac{1}{2}$, we have*

$$\sup_{\mathbf{v} \in \Lambda} \left| \frac{\mathcal{D}(\mathbf{v})}{\mathcal{D}_{\Lambda, \mathbf{t}, \Sigma}(\mathbf{v})} - 1 \right| \leq 6(\delta + \varepsilon).$$

In particular, \mathcal{D} is within statistical distance $3(\delta + \varepsilon)$ of $\mathcal{D}_{\Lambda, \mathbf{t}, \Sigma}$.

Proof. Let $\varepsilon > 0$ and $\mathbf{v} \in \Lambda$, and as usual denote by \mathbf{T} the torsion subgroup of Λ/Λ' . Independence of the sampling oracles of lines 3 and 4 and law of total probability yields:

$$\mathcal{D}(\mathbf{v}) = \Pr(\mathbf{u}' = \mathbf{v} - \mathbf{u}_q \mid \mathbf{q} = \pi(\mathbf{v})) \Pr(\mathbf{q} = \pi(\mathbf{v})).$$

By hypothesis on the oracle \mathcal{O}_q and Lemma 3.1, we have:

$$(8) \quad \Pr(\mathbf{q} = \pi(\mathbf{v})) \in \left[\frac{1 - \delta}{1 + \delta} \frac{1 - \varepsilon}{1 + \varepsilon}, \frac{1 + \delta}{1 - \delta} \frac{1 + \varepsilon}{1 - \varepsilon} \right] \cdot \frac{1}{|\mathbf{T}|} \cdot \mathcal{D}_{\pi(\Lambda), \pi(\mathbf{t}), \Sigma}(\pi(\mathbf{v})).$$

On the other hand we have $\pi(\mathbf{u}_q) = \mathbf{q}$ and thus:

$$(9) \quad \begin{aligned} \Pr(\mathbf{u}' = \mathbf{v} - \mathbf{u}_q | \mathbf{q} = \pi(\mathbf{v})) &= \frac{\rho_{\Sigma}(\mathbf{v} - \mathbf{u}_q - (\text{Id} - \pi)(\mathbf{t} - \mathbf{u}_q))}{\rho_{\Sigma}(\Lambda' - (\text{Id} - \pi)(\mathbf{t} - \mathbf{u}_q))} \\ &= \frac{\rho_{\Sigma}((\text{Id} - \pi)(\mathbf{v} - \mathbf{t}))}{\rho_{\Sigma}(\Lambda' - (\text{Id} - \pi)(\mathbf{t} - \mathbf{u}_q))}. \end{aligned}$$

Because $\Sigma \succ \eta_{\varepsilon}(\Lambda') \geq \eta_{\varepsilon}(\overline{\Lambda}')$, we obtain that

$$\begin{aligned} \rho_{\Sigma}(\Lambda' - (\text{Id} - \pi)(\mathbf{t} - \mathbf{u}_q)) &\in \left[1, \frac{1 + \varepsilon}{1 - \varepsilon} \right] \cdot \frac{1}{|\mathbf{T}|} \rho_{\Sigma}(\overline{\Lambda}' - (\text{Id} - \pi)(\mathbf{t} - \mathbf{u}_q)) \\ &\in \left[\frac{1 - \varepsilon}{1 + \varepsilon}, \left(\frac{1 + \varepsilon}{1 - \varepsilon} \right)^2 \right] \cdot \frac{1}{|\mathbf{T}|} \rho_{\Sigma}(\overline{\Lambda}' - (\text{Id} - \pi)(\mathbf{t})). \end{aligned}$$

We then find that:

$$\mathcal{D}(\mathbf{v}) \in \left[\frac{1 - \delta}{1 + \delta} \left(\frac{1 - \varepsilon}{1 + \varepsilon} \right)^3, \frac{1 + \delta}{1 - \delta} \left(\frac{1 + \varepsilon}{1 - \varepsilon} \right)^2 \right] \cdot \frac{\rho_{\Sigma}(\pi(\mathbf{v} - \mathbf{t}))}{\rho_{\Sigma}(\pi(\Lambda) - \pi(\mathbf{t}))} \cdot \frac{\rho_{\Sigma}((\text{Id} - \pi)(\mathbf{v} - \mathbf{t}))}{\rho_{\Sigma}(\overline{\Lambda}' + (\text{Id} - \pi)(\mathbf{t}))}$$

Routine calculations conclude the proof. $\ddot{::}$

Algorithm 2 $\ddot{::}$ Short exact sequence sampler

Input:

- A sublattice $\Lambda' \subset \Lambda$, a centre \mathbf{t}
- an oracle \mathcal{O}' for $\mathcal{D}_{\Lambda', *, \Sigma}$
- an oracle \mathcal{O}_q over Λ / Λ' , $\frac{1 + \delta}{1 - \delta}$ -close to the pushforward of $\mathcal{D}_{\Lambda, *, \Sigma}$

Output: $\mathbf{v} \in \Lambda$ following distribution statistically close to $\mathcal{D}_{\Lambda, \mathbf{t}, \Sigma}$

- 1 if $\Lambda = \{0\}$ then return 0
- 2 Compute $\pi : \Lambda_{\mathbf{R}} \rightarrow \Lambda_{\mathbf{R}} / \Lambda'_{\mathbf{R}}$, the orthogonal projection onto $\Lambda'_{\mathbf{R}}^{\perp}$ for the norm induced by Σ^{-1}
- 3 $\mathbf{q} \leftarrow \mathcal{O}_q(\Lambda / \Lambda', \pi(\mathbf{t}), \Sigma)$; $\mathbf{u}_q \leftarrow \mathbf{Lift}(\mathbf{q}, \Lambda)$
- 4 $\mathbf{u}' \leftarrow \mathcal{O}'(\Lambda', (\text{Id} - \pi)(\mathbf{t} - \mathbf{u}_q), \Sigma)$
- 5 return $\mathbf{u}_q + \mathbf{u}'$

4. GENERIC APPLICATIONS OF THE SHORT-SEQUENCE SAMPLER

We now present two generic applications of this abstract sampler: domain extensions and restrictions, and a broad generalization of the so-called Klein/GPV sampler [18]. The formers are simple, elementary but important illustrations of the use of the pushforward distribution; the latter extends the toolbox for Gaussian sampling in cryptography. In Section 8, we will present several concrete samplers in remarkable lattices using these generic constructions, reaching closer to the smoothing than the approaches used previously.

4.1. Domain extension and restriction.

4.1.1. *Extension to an overlattice.* Let Λ' be a full-rank sublattice of Λ , so that the quotient Λ/Λ' is of torsion (i.e. the free part of this quotient is reduced to $\{0\}$) and suppose that we have access to an oracle \mathcal{O} for $\mathcal{D}_{\Lambda',*,\Sigma}$ for a parameter $\Sigma \succ \eta_\varepsilon(\Lambda')$. By Lemma 3.1, the pushforward $\pi_*\mathcal{D}_{\Lambda',*,\Sigma}$ is at distance at most $\frac{\varepsilon}{1-\varepsilon}$ of the uniform distribution over Λ/Λ' . Hence specializing Algorithm 2 with \mathcal{O}' and a uniform sampler for \mathcal{O}_q yields the following:

Corollary 1 (Domain extension). *Let $\varepsilon > 0$ and Λ be a lattice, Λ' one of its sublattices of finite index. For any oracle \mathcal{O}' realizing a discrete Gaussian sampling in Λ' at any center and covariance $\Sigma \succ \eta_\varepsilon(\Lambda')$, there exists an algorithm sampling at distance at most 6ε of $\mathcal{D}_{\Lambda',*,\Sigma}$ using at most one oracle call to \mathcal{O}' .*

In a nutshell, the ability to sample in Λ' and Λ/Λ' enables to reconstruct samples in Λ : we do a *domain extension* of the discrete Gaussian over Λ' to the overlattice Λ .

4.1.2. *Restriction to a sublattice.* Conversely, it is easy to sample in a sublattice Λ' when we already know how to sample in Λ , and Λ' has finite index $[\Lambda : \Lambda']$: sample in Λ and reject all samples not landing in Λ' . The number of tries is of course driven by $[\Lambda : \Lambda']$, which can be proven when sampling above the smoothing of Λ' . In fact, it makes it a specific case of the *rejection sampling* technique, with trivial rejection probabilities. In the upcoming Section 8, we showcase some practical examples with root lattices.

Proposition 3 (Domain restriction). *Let $\varepsilon > 0$ and Λ be a lattice and Λ' one of its sublattices of finite index. For any oracle \mathcal{O} realizing a discrete Gaussian sampling in Λ with covariance $\Sigma \succ \eta_\varepsilon(\Lambda')$, there exists a Gaussian sampler (for the same covariance) in Λ' using on expectation $[\Lambda : \Lambda']$ calls to \mathcal{O} .*

Proof. The procedure is as follows: get a sample x from \mathcal{O} and return it if $x \in \Lambda'$, otherwise restart. The probability of $x \in \Lambda'$ is exactly $p' = \rho_\Sigma(\Lambda')/\rho_\Sigma(\Lambda)$ by definition of \mathcal{O} . As such the expected number of repetitions before a success is (as the expectation of a geometric distribution) $\frac{1}{p'}$. Since $\Sigma \succ \eta_\varepsilon(\Lambda')$, Lemma 3.1 implies that $p' \in \frac{1}{[\Lambda:\Lambda']} [1, \frac{1+\varepsilon}{1-\varepsilon}]$, bounding $1/p'$ by $[\Lambda:\Lambda']$. The correctness of the process follows from conditional probabilities. $\ddot{::}$

Remark. We point out the possible connection with the averaging recombination technique used in [1], where a domain restriction is performed, and samples are then combined with a domain extension from 2Λ to Λ (using exponentially many vectors).

4.2. A filtration sampler. We now show how our short exact sequence sampler naturally extends to filtrations and allows to retrieve and generalize samplers appearing in cryptography, such as those in [15, 18, 29]. For example, in the most natural case where one would sample “coordinate-by-coordinate”, our algorithm recovers Klein/GPV sampler. More generally it gives a family of new samplers for a given lattice, depending on how one decides to sort and “cut in subspaces” its input basis, offering larger freedom in the design of sampling algorithms⁸.

4.2.1. Smoothing parameter bound over a filtration. We first highlight a new smoothing parameter bound deduced from a given filtration $\{0\} = \Lambda_0 \subset \Lambda_1 \subset \dots \subset \Lambda_k = \Lambda$ of a lattice Λ . It relies on repeatedly applying the splits of the smoothing parameter over the short sequences (Proposition 2) stemming from the filtration. Starting from the penultimate term Λ_{k-1} , we bound (ignoring here the exact values of ε to ease the exposition) the smoothing parameter of Λ by $\max(\eta(\Lambda_{k-1}), \eta(\Lambda/\Lambda_{k-1}))$. Applying Proposition 2 to Λ_{k-1} , we have $\eta(\Lambda_{k-1}) \leq \max(\eta(\Lambda_{k-1}/\Lambda_{k-2}), \eta(\Lambda_{k-2}))$. We go down in the filtration inductively until we reach Λ_1 . All in all, the smoothing parameter is dominated by the biggest term appearing in the splitting.

Lemma 1. *Let $k \geq 1$ be an integer, Λ a lattice and $\varepsilon \leq \frac{2}{k}$. We have*

$$\eta_\varepsilon(\Lambda) \leq \min_{(\Lambda_i)_i} \max_i \eta_{\frac{\varepsilon}{k+1}} \left(\Lambda_i / \Lambda_{i-1} \right),$$

where the minimum is taken over all possible filtrations of length k of Λ .

⁸In the same way that Klein/GPV sampler is a randomized version of Babai’s nearest plane algorithm, our technique can be interpreted as a randomized version of the nearest-colattice algorithm of Espitau and Kirchner [16].

Proof. Let V_i be the real space spanned by Λ_i , and P_i be the orthogonal projection onto V_i^\perp , where the orthogonality is for the quadratic form $\mathbf{x} \mapsto \mathbf{x}^t \Sigma^{-1} \mathbf{x}$. Lemma 2.1 gives $\rho_\Sigma(\Lambda) \leq \rho_\Sigma(P_{k-1}(\Lambda)) \cdot \rho_\Sigma(\Lambda_{k-1})$. Letting P_0 be the identity, we obtain by induction

$$\rho_\Sigma(\Lambda) \leq \prod_{i=1}^k \rho_\Sigma(P_{i-1}(\Lambda_i)) = \prod_{i=1}^k \rho_\Sigma(\Lambda_i/\Lambda_{i-1}).$$

Using the Poisson Summation Formula (7) and taking $\Sigma \succ \max_i \eta_{\varepsilon/(k+1)}(\Lambda_i/\Lambda_{i-1})$ gives

$$\rho_{\Sigma^{-1}}(\Lambda^\vee) \leq \prod_{i=1}^k \rho_{\Sigma^{-1}}\left((\Lambda_i/\Lambda_{i-1})^\vee\right) \leq \left(1 + \frac{\varepsilon}{k+1}\right)^k.$$

Calculations concludes the proof, using the fact that $(1 + \varepsilon/(k+1))^k < 1 + \varepsilon$

∴

The term $k+1$ is chosen to obtain a compact, readable statement with an identical smoothing quality for each quotient lattice.⁹ The idea behind the above bound allows to mildly relax the smoothness condition over lattice cosets: instead of the whole lattice, it is only needed to smooth the “worst” of the successive quotients deduced from the filtration for the cosets of the whole lattice to have essentially the same mass.

For example, it was shown¹⁰ in [18], and subsequently used at the core of several practical constructions, that for any rank n lattice Λ ,

$$\eta_\varepsilon(\Lambda) \leq \min_{\substack{(\mathbf{b}_1, \dots, \mathbf{b}_n) \\ \text{basis of } \Lambda}} \max_{1 \leq i \leq n} \eta_n^\varepsilon(\mathbf{Z}\mathbf{b}_i^*),$$

where the \mathbf{b}_i^* 's are the Gram-Schmidt vector of the corresponding basis. This bound corresponds to restricting Lemma 1 to filtrations of length n stemming from the \mathbf{b}_i 's as $\Lambda_i = \mathcal{L}(\mathbf{b}_1, \dots, \mathbf{b}_i)$. Indeed, we have that for any $0 \leq i < n$, Λ_{i+1}/Λ_i is isometric to $\mathbf{Z}\mathbf{b}_i^*$ (see also Section 2.1.2). While it could seem more likely that such a fine-grained filtration would give in general better smoothing bounds, we actually show that there are practical cryptographic cases where one can improve the situation by carefully selecting a different and *a priori* coarser-grained filtration.

⁹What matters in the proof is that $\prod_i (1 + \varepsilon_i) \leq 1 + \varepsilon$, where ε_i is a given smoothing quality for Λ_i/Λ_{i-1} and ε is the target quality for Λ .

¹⁰In its usual form for a fixed basis, the bound is $\eta_\varepsilon(\Lambda) \leq \max_{1 \leq i \leq n} \|\mathbf{b}_i^*\| \cdot \eta_\varepsilon(\mathbf{Z}^n)$.

4.2.2. *The filtered sampler.* Following our motto — smoothing bounds and sampling are built on the same underlying principles — we can transform Lemma 1 into a Gaussian sampler. In essence, the process corresponds to k successive calls of Algorithm 2, recursively progressing along the filtration.

Assume that we are given approximate oracles to sample discrete Gaussians in the sequence of lattices $(\Lambda_{i+1}/\Lambda_i)_i$, and a deterministic lift the first call considers the short exact sequence

$$0 \rightarrow \Lambda_1 \rightarrow \Lambda \rightarrow \Lambda/\Lambda_1 \rightarrow 0.$$

Algorithm 2 requires a pushforward oracle on Λ/Λ_1 , so since we do not have *a priori* an explicit access to it, we instantiate it as a recursive call over the quotient filtration $\{0\} = \Lambda_1/\Lambda_1 \subset \Lambda_2/\Lambda_1 \subset \dots \subset \Lambda/\Lambda_1$. Hence the callee now deals with the sequence $0 \rightarrow \Lambda_2/\Lambda_1 \rightarrow \Lambda/\Lambda_1 \rightarrow \Lambda/\Lambda_2 \rightarrow 0$. This is done until we reach the trivial sequence. Then, the algorithm climbs its way back in the recursion tree, providing samples in the lattices Λ_{i+1}/Λ_i .

Algorithm 3 \therefore Filtered sampler

Input: A filtration $\{0\} \subset \Lambda_1 \subset \dots \subset \Lambda_k = \Lambda$, a parameter $\Sigma > \max_{0 \leq i < k} \eta_\varepsilon(\Lambda_{i+1}/\Lambda_i)$ and a center $\mathbf{t} \in \Lambda \otimes \mathbf{R}$.

Output: $\mathbf{v} \in \Lambda$ following distribution statistically close to $\mathcal{D}_{\Lambda, \mathbf{t}, \Sigma}$

- 1 if $\Lambda = \{0\}$ then return 0
- 2 Compute $\pi : \Lambda \rightarrow \Lambda/\Lambda_1$
- 3 $\mathbf{z} \leftarrow \text{FilteredSampler}\left(\left(\Lambda_i/\Lambda_1\right)_i, \pi(\mathbf{t}), \Sigma\right)$
- 4 $\mathbf{u} \leftarrow \text{Lift}(\mathbf{z}, V_1)$
- 5 $\mathbf{u}' \leftarrow \mathcal{D}_{\Lambda_1, (\text{Id}-\pi)(\mathbf{t}-\mathbf{u}), \Sigma}$
- 6 return $\mathbf{u} + \mathbf{u}'$

Theorem 2 (Correctness of the filtered sampler). *Algorithm 3 is correct. Moreover, let \mathcal{D} be the distribution of its output. For any $\varepsilon < 1/k^2$, we have*

$$\sup_{\mathbf{v} \in \Lambda} \left| \frac{\mathcal{D}(\mathbf{v})}{\mathcal{D}_{\Lambda, \mathbf{t}, \Sigma}(\mathbf{v})} - 1 \right| \leq (2k + 1)\varepsilon.$$

In particular, \mathcal{D} is within statistical distance $(k + 1)\varepsilon$ of $\mathcal{D}_{\Lambda, \mathbf{t}, \Sigma}$.

It suffices to proceed by induction along the filtration repeatedly calling Algorithm 3.

Proof. At Step 5, the algorithm recalls itself on the filtration obtained by quotienting by its first element. The first time it happens, the input filtration is then $\{0\} \subset \Lambda_2/\Lambda_1 \subset \dots \subset \Lambda/\Lambda_1$. Let $V_i = \Lambda_i \otimes \mathbf{R}$. By isomorphism theorems, the quotient is then (isometric to) Λ/Λ_2 , and by induction, at the i -th call, the input filtration is therefore $\{0\} \subset \Lambda_{i+1}/\Lambda_i \subset \dots \subset \Lambda/\Lambda_i$, with a corresponding flag of subspaces $(V_{i+1} \cap V_i^\perp) \subset \dots \subset (V \cap V_i^\perp)$. By assumptions, the Gaussian oracle is always able to sample in the first element Λ_{i+1}/Λ_i of its input filtration. The standard deviation is always above the smoothing parameter of such lattices, Algorithm 3 outputs an element in the lattice thanks to the properties of Algorithm 1.

We now analyze the distribution of the outputs. Let $\varepsilon > 0$ and set $\delta = \frac{1-\varepsilon}{1+\varepsilon}$. We use the following loop invariant: if the input filtration contains $k - 1$ elements and the target center is \mathbf{t}' , then the probability that Algorithm 3 outputs some vector \mathbf{v} belongs to $[\delta^{k-1}, \delta^{1-k}] \cdot \mathcal{D}_{\Lambda', \mathbf{t}', \Sigma}(\mathbf{v})$, where Λ' is the lattice spanned by the input filtration. This hypothesis is satisfied for any filtration with 1 element. Let us assume now that it is true up to some $k \geq 1$. By construction, we can write $\mathbf{u} = \mathbf{z} + \mathbf{v}$ for some $\mathbf{v} \in V_1$. In particular, we have $(\text{Id} - \pi)(\mathbf{t} - \mathbf{u}) = (\text{Id} - \pi)(\mathbf{t}) - \mathbf{v}$. Next, let $P(\mathbf{z})$ resp. $P(\mathbf{u}')$ the probability to obtain \mathbf{z} resp. \mathbf{u}' . Orthogonality gives us $\rho_\Sigma(\mathbf{u} + \mathbf{u}' - \mathbf{t}) = \rho_\Sigma(\mathbf{z} - \pi(\mathbf{t}))\rho_\Sigma(\mathbf{u}' - (\text{Id} - \pi)(\mathbf{t}) + \mathbf{v})$. The induction hypothesis then yields

$$\begin{aligned} P(\mathbf{z})P(\mathbf{u}') &\in \left[\delta^{k-1}, \delta^{1-k} \right] \cdot \frac{\rho_\Sigma(\mathbf{z} - \pi(\mathbf{t}))}{\rho_\Sigma(\Lambda/\Lambda_1 - \pi(\mathbf{t}))} \cdot \frac{\rho_\Sigma(\mathbf{u}' - (\text{Id} - \pi)(\mathbf{t}) + \mathbf{v})}{\rho_\Sigma(\Lambda_1 - (\text{Id} - \pi)(\mathbf{t}) + \mathbf{v})} \\ &= \left[\delta^{k-1}, \delta^{1-k} \right] \cdot \mathcal{D}_{\Lambda, \mathbf{t}, \Sigma}(\mathbf{u} + \mathbf{u}') \cdot \frac{\rho_\Sigma(\Lambda - \mathbf{t})}{\rho_\Sigma(\Lambda/\Lambda_1 - \pi(\mathbf{t}))\rho_\Sigma(\Lambda_1 - (\text{Id} - \pi)(\mathbf{t}) + \mathbf{v})}. \end{aligned}$$

Since $\Sigma \geq \eta_\varepsilon(\Lambda_1)$, we have $\rho_\Sigma(\Lambda_1 - (\text{Id} - \pi)(\mathbf{t}) + \mathbf{v}) \in [\delta, \delta^{-1}] \cdot \rho_\Sigma(\Lambda_1 - (\text{Id} - \pi)(\mathbf{t}))$. With a last orthogonality argument, we obtain our claim on the distribution, using that $\left(\frac{1+\varepsilon}{1-\varepsilon}\right)^{k-1} < \frac{1+k\varepsilon}{1-k\varepsilon}$ when $\varepsilon < \frac{1}{k^2}$. $\ddot{:}$

4.3. Recovering some known samplers.

4.3.1. *Klein/GPV sampler.* As we saw, this sampler corresponds to taking the full filtration associated to a lattice basis $(\mathbf{b}_1, \dots, \mathbf{b}_n)$ giving a lower bound on the width in $\max_i \eta_\varepsilon(\Lambda_{i+1}/\Lambda_i) = \max_i \eta_\varepsilon(\mathbf{b}_i^* \mathbf{Z}) = \eta_\varepsilon(\mathbf{Z}) \cdot \max_i(\|\mathbf{b}_i^*\|)$.

4.3.2. *Klein/GPV sampler over a ring.* This sampler works at the ring level of a module over some ring of integer $\mathcal{O}_{\mathbf{K}}$ in a number field \mathbf{K} (for example on NTRU lattices which are rank two module over a cyclotomic ring). More precisely given a module basis $(\mathbf{m}_1, \dots, \mathbf{m}_d)$ over

$\mathcal{O}_{\mathbf{K}}$, we make use of the full filtration¹¹ $\mathfrak{m}_1\mathcal{O}_{\mathbf{K}} \subset \mathfrak{m}_1\mathcal{O}_{\mathbf{K}} \oplus \mathfrak{m}_2\mathcal{O}_{\mathbf{K}} \subset \dots$. Each recursive call thus consists in calling the oracles over the quotients $\mathfrak{m}_{i+1}\mathcal{O}_{\mathbf{K}}/\mathfrak{m}_i\mathcal{O}_{\mathbf{K}}$, which are scaling by an algebraic number of $\mathcal{O}_{\mathbf{K}} \cong \mathbf{Z}^{\deg \mathbf{K}}$. When instantiating this oracle with subsequently described Algorithm 4 (or Peikert’s [26] for instance), it retrieves the so-called *hybrid sampler* used in [15].

4.3.3. *Fast Fourier Orthogonalization sampler.* Introduced in [12, 30], this sampler reaches the same quality as Klein’s sampler but run in quasi-linear time in the dimension, by exploiting the structure of tower of subfields in power-of-two cyclotomic fields. It is retrieved as the filtered sampler where the oracle over the ring is the sampler itself, called recursively. More precisely given a basis $\mathfrak{m}_1, \mathfrak{m}_2$ of a module Λ over the ring of integers $\mathcal{O}_{\mathbf{K}}$ of the cyclotomic field of conductor 2^k , we have the short exact sequence:

$$0 \longrightarrow \mathfrak{m}_1\mathcal{O}_{\mathbf{K}} \longrightarrow \Lambda \longrightarrow \Lambda/\mathfrak{m}_1\mathcal{O}_{\mathbf{K}} \longrightarrow 0$$

where once again the submodule $\mathfrak{m}_1\mathcal{O}_{\mathbf{K}}$ shall be understood as a sublattice through the canonical embedding map. Now remark that an oracle call is made on the modules of rank 1 $\mathfrak{m}_1\mathcal{O}_{\mathbf{K}}$ and $\Lambda/\mathfrak{m}_1\mathcal{O}_{\mathbf{K}}$. However, these modules are also modules of rank 2 over the cyclotomic field of conductor 2^{k-1} . As such, for each of them we can apply the same technique recursively, requiring samples in modules of rank 2 over smaller and smaller fields, until we eventually reach \mathbf{Q} , where we know how to sample.

4.3.4. *A filtered tensor sampler.* Let A be a lattice given with a complete filtration $\{0\} \subset A_1 \subset \dots \subset A_\ell \subset A$ and another lattice B . Suppose that we have a gaussian sampler \mathcal{O} in B . Then remark that the tensor filtration $\{0\} \subset A_1 \otimes B \subset \dots \subset A_\ell \otimes B \subset A \otimes B$ is a filtration of $A \otimes B$ and that since each quotient A_{i+1}/A_i is of dimension 1, the quotients $(A_{i+1} \otimes B)/(A_i \otimes B)$ are actually isometrics to scalings of B — the scaling factor being exactly the covolume of the line $A_{i+1} \otimes B/A_i \otimes B$. Hence, from the sampler \mathcal{O} , we can sample in $A \otimes B$ by the ℓ calls generated when applying the **Filtered Sampler** to our filtration.

¹¹We make a slight abuse of notations here by silently identifying a submodule with the corresponding sublattice of the lattice attached to the module. To be perfectly formal, we shall understand the elements of the filtration as viewed under the canonical embedding map recalled in Section 2.3.1.

5. THE LINEAR SAMPLER

5.1. Smoothing parameters and linear transformations. The algorithms presented in Section 3 sample without leaving the ambient space of the lattice. However, in certain cases, it is of interest to transfer the problem to another space – where the local geometry eases the sampling process – and transfer the result back to the original lattice. In a sense, as all lattices can be seen as a transformation of the integer lattice \mathbf{Z}^n , and as most practical Gaussian samplers rely on the ability to sample integral Gaussians, this observation is already implicit in previous works. As expected, such back and forth between different spaces will generate bias because of the *distortion* incurred by the underlying linear transformation. To enforce the correctness of the output distribution, it must be corrected. For example, the filtered sampler of Section 4.2.2 *iteratively* corrects the transformation to the space attached to the filtration one subspace at a time. A *global* approach to the problem consists in considering any lattice as a linear transformation of another lattice, but not always \mathbf{Z}^n . This gives the following bound on the smoothing parameter, possibly implicit in previous works.

Lemma 2. *Let Λ be a lattice of rank n in \mathbf{R}^m , then $\eta_\varepsilon(\Lambda) \leq \inf s_1(\mathbf{T}) \cdot \eta_\varepsilon(\mathcal{L}(\mathbf{C}))$, where the infimum is taken over all pair (\mathbf{T}, \mathbf{C}) such that $\Lambda = \mathcal{L}(\mathbf{T}\mathbf{C})$ and $\mathbf{C} \in \mathbf{R}^n$ is invertible.*

Proof. Let \mathbf{T}, \mathbf{C} be any such decomposition of a given basis \mathbf{B} of Λ . The basis of the dual lattice Λ^\vee is then $\mathbf{B}^\vee = \mathbf{B}(\mathbf{B}^t\mathbf{B})^{-1}$, and as such for any vector $\mathbf{z} \in \mathbf{R}^m$ we have $(\mathbf{B}^\vee\mathbf{z})^t(\mathbf{B}^\vee\mathbf{z}) = \mathbf{z}^t(\mathbf{B}^t\mathbf{B})^{-1}\mathbf{z} = (\mathbf{C}^{-t}\mathbf{z})^t(\mathbf{T}^t\mathbf{T})^{-1}(\mathbf{C}^{-t}\mathbf{z})$. This implies that for any $s > 0$, we have $\rho_{\frac{1}{s^2}}(\Lambda^\vee) = \rho_{\frac{\mathbf{T}^t\mathbf{T}}{s^2}}(\mathcal{L}(\mathbf{C})^\vee)$. Asking $s \geq \eta_\varepsilon(\Lambda)$ is thus equivalent to asking that $s^2(\mathbf{T}^t\mathbf{T})^{-1} \succ \eta_\varepsilon(\mathcal{L}(\mathbf{C}))^2 \cdot \mathbf{I}_n$, as stated. ∴

5.2. Sampling by linear transformation. The global approach is an algorithmic formulation of the proposition of Peikert [26, Theorem 3.1]. For the sake of simplicity, we will restrict ourselves to the case of *continuous* perturbations in our presentation. As explained, on a high level the transformation of a *fixed* lattice distorts the geometry in the initial space and consequently any ellipsoid in that space. The bias can be corrected to any target ellipsoidal shape by adding a large enough perturbation, and up to rescaling.

Going formal, one can prove the correctness of the approach thanks to the nice properties of Gaussian distributions, and the scaling factor appears implicitly as a condition of positive-definiteness involving the smoothing parameter of the initial lattice.

Theorem 3 (Correctness of the linear sampler). *Let $r \geq \eta_\varepsilon(\Lambda(\mathbf{C}))$. If $s_n(\Delta) > r^2 \cdot s_1(\mathbf{T})^2$, then Algorithm 4 is correct. Moreover, let \mathcal{D} be the distribution of its output. For $\varepsilon < 1/2$, we have*

$$\sup_{\mathbf{v} \in \Lambda} \left| \frac{\mathcal{D}(\mathbf{v})}{\mathcal{D}_{\Lambda, \mathbf{t}, \Delta}(\mathbf{v})} - 1 \right| \leq 4\varepsilon.$$

In particular, \mathcal{D} is within statistical distance 2ε of $\mathcal{D}_{\Lambda, \mathbf{t}, \Delta}$.

Proof. The support of the output distribution is correct by construction. By construction, the probability of sampling some $\mathbf{p} \in \mathbf{C}_{\mathbf{R}} := \text{span}_{\mathbf{R}}(\mathbf{C})$, $\mathbf{x} \in \mathcal{L}(\mathbf{C})$ and outputting \mathbf{y} is the marginalized distribution

$$P(\mathbf{y}) = \det(\Sigma_{\mathbf{p}})^{-1/2} \cdot \int_{\mathbf{C}_{\mathbf{R}}} \frac{\rho_{\Sigma_{\mathbf{p}}}(\mathbf{p}) \cdot \rho_{r^2}(\mathbf{x} - \mathbf{T}^* \mathbf{t} - \mathbf{p})}{\rho_{r^2}(\mathcal{L}(\mathbf{C}) - \mathbf{T}^* \mathbf{t} - \mathbf{p})} d\mathbf{p}.$$

Gaussian functions have also good multiplicative properties (see for instance [26, Fact 2.10]): there exists¹² a positive definite matrix Σ' and a vector \mathbf{t}' , both over $\mathbf{C}_{\mathbf{R}}$, such that:

$$\rho_{\Sigma_{\mathbf{p}}}(\mathbf{p}) \rho_{r^2}(\mathbf{x} - \mathbf{T}^* \mathbf{t} - \mathbf{p}) = \rho_{\Sigma}(\mathbf{x} - \mathbf{T}^* \mathbf{t}) \rho_{\Sigma'}(\mathbf{p} - \mathbf{t}'),$$

where we also use that $\Sigma_{\mathbf{p}} + r^2 = \Sigma$. Combining these two equalities, we rewrite the distribution of the output as

$$P(\mathbf{y}) = \det(\Sigma_{\mathbf{p}})^{-1/2} \cdot \rho_{\Sigma}(\mathbf{x} - \mathbf{T}^* \mathbf{t}) \cdot \int_{\mathbf{C}_{\mathbf{R}}} \frac{\rho_{\Sigma'}(\mathbf{p} - \mathbf{t}')}{\rho_{r^2}(\mathcal{L}(\mathbf{C}) - \mathbf{T}^* \mathbf{t} - \mathbf{p})} d\mathbf{p}.$$

By definition of the pseudo-inverse, we have that $\mathbf{T}^* \mathbf{y} = \mathbf{x}$ and that $\mathbf{T} \mathbf{T}^*$ is the orthogonal projection \mathbf{P} onto $\Lambda_{\mathbf{R}}$. This gives $(\mathbf{x} - \mathbf{T}^* \mathbf{t})^t \Sigma^{-1} (\mathbf{x} - \mathbf{T}^* \mathbf{t}) = (\mathbf{y} - \mathbf{t})^t \mathbf{P}^t \Delta^{-1} \mathbf{P} (\mathbf{y} - \mathbf{t})$ and we obtain

$$\rho_{\Sigma}(\mathbf{x} - \mathbf{T}^* \mathbf{t}) = \rho_{\Delta}(\mathbf{y} - \mathbf{t}).$$

Observe that if $m = n$, we have $\mathbf{T}^* = \mathbf{T}^{-1}$ so that $\mathbf{P} = \mathbf{I}_n$ and the result holds as well. By assumptions on r and thanks to Lemma 2.1, we now have

$$P(\mathbf{y}) \in \left[1, \frac{1 + \varepsilon}{1 - \varepsilon} \right] \cdot \left(\frac{\det \Sigma'}{\det \Sigma_{\mathbf{p}}} \right)^{1/2} \cdot \frac{\rho_{\Delta}(\Lambda - \mathbf{t})}{\rho_{r^2}(\mathcal{L}(\mathbf{C}))} \cdot \mathcal{D}_{\Lambda, \mathbf{t}, \Delta}(\mathbf{y}),$$

and summing over all \mathbf{y} to handle the normalization constants, we deduce

$$P(\mathbf{y}) \in \left[\frac{1 - \varepsilon}{1 + \varepsilon}, \frac{1 + \varepsilon}{1 - \varepsilon} \right] \cdot \mathcal{D}_{\Lambda, \mathbf{t}, \Delta}(\mathbf{y}).$$

Our claims on the maximum relative error and the statistical distance follow. Lastly, we show that the condition on the singular values is sufficient. The algorithm is correct as soon as

¹²Their expressions can be made explicit but are not needed to understand the rest of the proof.

$\Sigma \succ r^2 \mathbf{I}_n$, or equivalently when $s_1(\mathbf{T}^t \Delta^{-1} \mathbf{T}) < r^{-2}$. Using standard properties of operator norms, we see that $s_1(\mathbf{T}^t \Delta^{-1} \mathbf{T}) \leq s_1(\mathbf{T})^2 \cdot s_n(\Delta)^{-1}$, and the results follow. $\ddot{::}$

Algorithm 4 $\ddot{::}$ Linear sampler

Input:

- Two matrices $\mathbf{T} \in \mathbf{R}^{m \times n}$, $\mathbf{C} \in \mathbf{R}^{n \times n}$ with $m \geq n$ and \mathbf{C} invertible such that $\mathbf{TC} = \mathbf{B}$ is a basis of a lattice Λ ;
- a center $\mathbf{t} \in \Lambda \otimes \mathbf{R}$;
- a parameter $r \geq 0$ and a positive definite matrix $\Delta \in \mathbf{R}^{m \times m}$ such that $\Sigma := (\mathbf{T}^t \Delta^{-1} \mathbf{T})^{-1} \succ r^2 \mathbf{I}_n$;

Output: $\mathbf{y} \in \Lambda$ with distribution statistically close to $\mathcal{D}_{\Lambda, \mathbf{t}, \Delta}$.

```

1  $\Sigma_{\mathbf{p}} \leftarrow \Sigma - r^2 \mathbf{I}$ 
2  $\mathbf{p} \leftarrow \mathcal{N}_{\Sigma_{\mathbf{p}}}$ 
3  $\mathbf{x} \leftarrow \mathcal{D}_{\mathcal{L}(\mathbf{C}), \mathbf{T}^* \mathbf{t} + \mathbf{p}, r^2} / * \mathbf{T}^*$  is the pseudo-inverse */
4 ; return  $\mathbf{y} := \mathbf{T} \mathbf{x}$ 

```

Remark. This sampler also relies on a continuous Gaussian sampler, but fundamentally, the required property is that the product of the density functions of the perturbation and the lattice sampler can be understood: this is the core fact used to ensure the correctness of the output.

5.3. Example of elementary instantiations. This abstract framework allows to easily recover known samplers and extend them in few directions. We already highlighted that using $\mathbf{C} = \mathbf{I}_n$ and full-rank lattices retrieves Peikert's sampler [26].

- **SVD based sampler::** The singular value decomposition gives $\mathbf{B} = \mathbf{X} \Delta \mathbf{Y}^t$ with \mathbf{X}, \mathbf{Y} orthogonal matrices and Δ diagonal. It amounts to a diagonalisation of the Gram matrix $\mathbf{B}^t \mathbf{B}$, or in other words to work in a basis of eigenvectors. As \mathbf{X} and \mathbf{Y} are orthogonal, the sampler reaches standard deviations starting $s_1(\mathbf{B}) \cdot \eta_\varepsilon(\mathbf{Z}^n)$, and gives a quality equivalent to the one of Peikert.
- **QR-based sampler::** Using the QR decomposition $\mathbf{B} = \mathbf{QR}$, one can use e.g. the above procedure to sample in $\mathcal{L}(\mathbf{R})$ with coordinate domain \mathbf{Z}^d , and then obtains a discrete Gaussian in $\mathcal{L}(\mathbf{B})$. Since \mathbf{Q} is orthogonal, the sampler reaches standard deviations starting $s_1(\mathbf{R}) \cdot \eta_\varepsilon(\mathbf{Z}^n) = s_1(\mathbf{B}) \cdot \eta_\varepsilon(\mathbf{Z}^n)$. Alternatively, as $s_1(\mathbf{Q}) = 1$, if one knows

how to sample in $\mathcal{L}(\mathbf{R})$ then the linear sampler can reach standard deviations greater than $\eta_\varepsilon(\mathcal{L}(\mathbf{R}))$.

- **Gram-Schmidt based sampler:** Let $\mathbf{B} = \tilde{\mathbf{B}}\mathbf{U}$ be the Gram-Schmidt decomposition of \mathbf{B} . In particular, \mathbf{U} is upper triangular with 1's on its diagonal, and therefore its own Gram-Schmidt orthogonalization corresponds to the identity matrix. By Lemma 1, we thus have $\eta_\varepsilon(\mathcal{L}(\mathbf{U})) \leq \eta_{\varepsilon/(n+1)}(\mathbf{Z})$. Using the standard bound $\eta_\varepsilon(\mathbf{Z}) \leq \sqrt{\ln(2(1+1/\varepsilon)/\pi)}$, the sampler with that decomposition reaches standard deviations very close to $s_1(\tilde{\mathbf{B}}) \cdot \eta_\varepsilon(\mathbf{Z}^n)$. This sampler is used in [15, Algorithm 4, 5] in an approximate version.
- **Hybrid sampler:** More generally, one can decompose the ambient space of Λ in orthogonal subspaces, use a sampler in each subspaces and obtain a discrete Gaussian in Λ by multiplying by a block-orthogonal matrix. This approach is similar to the hybrid sampler of [15, 29], when decomposing the NTRU lattice.

6. APPLICATION: SAMPLING IN TENSOR LATTICES

As a direct application of Section 5, we present a novel (up to our knowledge) approach to sample in tensor products of lattices. They appear naturally with rings of cyclotomic integers of smooth conductors. In particular, one could use the algorithm of this section to sample in $\mathcal{R}_{p^\ell q^k}$ for $p \neq q$ prime, in several ways depending on how the ring is embedded in a euclidean space.

6.1. Tensors and related bounds. Let $\mathbf{A} \in \mathbf{R}^{m_1 \times n_1}$ and $\mathbf{B} \in \mathbf{R}^{m_2 \times n_2}$. Recall that $\mathcal{L}(\mathbf{A}) \otimes \mathcal{L}(\mathbf{B}) = \mathcal{L}(\mathbf{A} \otimes \mathbf{B})$, where $\mathbf{A} \otimes \mathbf{B}$ is the Kronecker product between these matrices:

$$\mathbf{A} \otimes \mathbf{B} = \begin{bmatrix} a_{1,1}\mathbf{B} & a_{1,2}\mathbf{B} & \cdots & a_{1,n_1}\mathbf{B} \\ \vdots & \vdots & \ddots & \vdots \\ a_{m_1,1}\mathbf{B} & a_{m_1,2}\mathbf{B} & \cdots & a_{m_1,n_1}\mathbf{B} \end{bmatrix}$$

The *mixed product property* states that

$$\mathbf{A} \otimes \mathbf{B} = (\mathbf{A} \otimes \mathbf{I}_{m_2})(\mathbf{I}_{n_1} \otimes \mathbf{B}) = (\mathbf{I}_{m_1} \otimes \mathbf{B})(\mathbf{A} \otimes \mathbf{I}_{n_2}).$$

Tensoring with identities preserves geometric properties.

Lemma 3. *Let $\mathbf{A} \in \mathbf{R}^{m \times n}$ of full column rank, $k \neq \ell$ positive integers, and $\varepsilon < 1/2$. We have:*

- (1) $s_1(\mathbf{A} \otimes \mathbf{I}_k) = s_1(\mathbf{I}_\ell \otimes \mathbf{A}) = s_1(\mathbf{A})$;
- (2) $\eta_\varepsilon(\mathcal{L}(\mathbf{I}_k \otimes \mathbf{A})) \geq \eta_{\varepsilon/k}(\mathcal{L}(\mathbf{A})) \geq \eta_{2\varepsilon}(\mathcal{L}(\mathbf{I}_k \otimes \mathbf{A}))$.

$$(3) \eta_\varepsilon(\mathcal{L}(\mathbf{I}_k \otimes \mathbf{A})) = \eta_\varepsilon(\mathcal{L}(\mathbf{A} \otimes \mathbf{I}_k));$$

Proof. Since $\mathbf{I}_k \otimes \mathbf{A} = \text{diag}(\mathbf{A}, \dots, \mathbf{A})$, we have $s_1(\mathbf{I}_k \otimes \mathbf{A}) = s_1(\mathbf{A})$. Next, the dual basis of $\mathbf{I}_k \otimes \mathbf{A}$ is $\mathbf{I}_k \otimes \mathbf{A}^\vee$. We then have

$$\rho_{1/s}(\mathcal{L}(\mathbf{I}_k \otimes \mathbf{A})^\vee) = \rho_{1/s}\left(\bigoplus_{i \leq k} \mathcal{L}(\mathbf{A})^\vee\right) = \rho_{1/s}(\mathcal{L}(\mathbf{A})^\vee)^k.$$

Taking $s = \eta_\varepsilon(\mathcal{L}(\mathbf{I}_k \otimes \mathbf{A}))$ gives $\rho_{1/s}(\mathcal{L}(\mathbf{A})) \leq (1 + \varepsilon)^{1/k} \leq 1 + \varepsilon/k$, so $s \geq \eta_{\varepsilon/k}(\mathcal{L}(\mathbf{A}))$. Similarly, taking $s = \eta_{\varepsilon/k}(\mathcal{L}(\mathbf{A}))$ gives $\rho_{1/s}(\mathcal{L}(\mathbf{I}_k \otimes \mathbf{A})) \leq (1 + \varepsilon/k)^k \leq 1 + 2\varepsilon$, and the second property follows. To conclude, one checks that for all k , there always exist permutations $\mathbf{P} \in \mathbf{R}^{km \times km}$, $\mathbf{Q} \in \mathbf{R}^{kn \times kn}$ such that $\mathbf{P}(\mathbf{A} \otimes \mathbf{I}_k)\mathbf{Q} = \mathbf{I}_k \otimes \mathbf{A}$. Since permutations matrices are isometries of the standard Euclidean norm, we have $\|\mathbf{P}(\mathbf{A} \otimes \mathbf{I}_k)\mathbf{Q}\mathbf{x}\| = \|(\mathbf{A} \otimes \mathbf{I}_k)\mathbf{Q}\mathbf{x}\|$ for all $\mathbf{x} \in \mathbf{R}^{kn}$. Adding that \mathbf{Q} is an invertible transformation of \mathbf{R}^{kn} that stabilizes \mathbf{Z}^{kn} , we obtain $s_1(\mathbf{A} \otimes \mathbf{I}_k) = s_1(\mathbf{I}_k \otimes \mathbf{A})$ and $\rho_{1/s}(\mathcal{L}(\mathbf{I}_k \otimes \mathbf{A})) = \rho_{1/s}(\mathcal{L}(\mathbf{A} \otimes \mathbf{I}_k))$. $\ddot{}$

6.1.1. *Smoothing bound for tensors.* Combining these properties with Lemma 2 gives a bound on the smoothing parameter of a tensor product of lattices. Our bound involves singular values of the left factor \mathbf{T} in the mixed-product decomposition of $\mathbf{A} \otimes \mathbf{B}$, whereas the bound given in [24, Corollary 2.7] is associated to the maximal length $\|\mathbf{B}\|_{GS}$ of the Gram-Schmidt vectors of that factor. It is known that $\|\mathbf{B}\|_{GS} \leq s_1(\mathbf{B})$, so that our bound is generally of worse quality, but its effectiveness does not rely on Gram-Schmidt orthogonalization. In essence, this is the same tradeoff as between Peikert's randomized round-off approach [26] and Klein's randomized nearest plane (e.g., [18]).

Lemma 6.1. *Let $\Lambda = \mathcal{L}(\mathbf{A}) \otimes \mathcal{L}(\mathbf{B})$ for matrices $\mathbf{A} \in \mathbf{R}^{m_1 \times n_1}$ and $\mathbf{B} \in \mathbf{R}^{m_2 \times n_2}$ of full rank, where $m_i \geq n_i$. We have*

$$\eta_\varepsilon(\Lambda) \leq \min\left(s_1(\mathbf{A}) \cdot \eta_{\varepsilon/(2n_1)}(\mathcal{L}(\mathbf{B})), s_1(\mathbf{B}) \cdot \eta_{\varepsilon/(2n_2)}(\mathcal{L}(\mathbf{A}))\right).$$

6.2. **A sampling algorithm for tensor lattices.** Using Algorithm 4, this bound directly translates into Algorithm 5, where it is assumed that oracles for discrete Gaussian sampling in $\mathcal{L}(\mathbf{A})$ and $\mathcal{L}(\mathbf{B})$ are given.

Algorithm 5 \therefore : Tensor sampler

Input: Two matrices $\mathbf{A} \in \mathbf{R}^{m_1 \times n_1}$, $\mathbf{B} \in \mathbf{R}^{m_2 \times n_2}$ with $m_i \geq n_i$, giving a basis of $\Lambda = \mathcal{L}(\mathbf{A} \otimes \mathbf{B})$; a center $\mathbf{t} \in \Lambda \otimes \mathbf{R}$; $r \geq 0$; a positive definite matrix $\Delta \in \mathbf{R}^{m_1 m_2 \times n_1 n_2}$

Output: $\mathbf{y} \in \Lambda$ with distribution statistically close to the discrete Gaussian $\mathcal{D}_{\Lambda, \mathbf{t}, \Delta}$.

- 1 Select $(\mathbf{T}, \mathbf{C}) \in \{(\mathbf{A} \otimes \mathbf{I}_{m_2}, \mathbf{I}_{n_1} \otimes \mathbf{B}), (\mathbf{I}_{m_1} \otimes \mathbf{B}, \mathbf{A} \otimes \mathbf{I}_{n_2})\}$ minimizing $s_1(\mathbf{T}) \cdot \eta_\varepsilon(\mathcal{L}(\mathbf{C}))$
- 2 $\Sigma \leftarrow (\mathbf{T}^t \Delta^{-1} \mathbf{T})^{-1}$
- 3 **return** *Algorithm 4*($\mathbf{T}, \mathbf{C}, \mathbf{t}, \Delta, \Sigma, r$)

Corollary 2 (of Theorem 3). *Let $\Lambda = \mathcal{L}(\mathbf{A}) \otimes \mathcal{L}(\mathbf{B})$ for matrices $\mathbf{A} \in \mathbf{R}^{m_1 \times n_1}$ and $\mathbf{B} \in \mathbf{R}^{m_2 \times n_2}$ of full rank, where $m_i \geq n_i$. Let $(\mathbf{T}, \mathbf{C}) \in \{(\mathbf{A} \otimes \mathbf{I}_{m_2}, \mathbf{I}_{n_1} \otimes \mathbf{B}), (\mathbf{I}_{m_1} \otimes \mathbf{B}, \mathbf{A} \otimes \mathbf{I}_{n_2})\}$ be the pair minimizing $s_1(\mathbf{T}) \cdot \eta_\varepsilon(\mathcal{L}(\mathbf{C}))$, and let also $r \geq \eta_\varepsilon(\mathbf{C})$. Algorithm 5 is correct when $s_n(\Delta) \geq r^2 \cdot s_1(\mathbf{T})^2$. Moreover, let \mathcal{D} be the distribution of its output. For $\varepsilon < 1/2$, we have*

$$\sup_{\mathbf{v} \in \Lambda} \left| \frac{\mathcal{D}(\mathbf{v})}{\mathcal{D}_{\Lambda, \mathbf{t}, \Delta}(\mathbf{v})} - 1 \right| \leq 4\varepsilon.$$

In particular, \mathcal{D} is within statistical distance 2ε of $\mathcal{D}_{\Lambda, \mathbf{t}, \Delta}$.

The above statement is formulated to achieve the “smallest possible” covariance matrix, which is often the goal in practice for security reasons. In other contexts, it might happen that the other decomposition is chosen or needed. The choice of ε in our formulation somewhat hides a dimensional factor coming from tensoring \mathbf{A} or \mathbf{B} with an identity matrix, as expressed in the second property in Lemma 3.

7. AN EXACT CALCULATION OF THE SMOOTHING PARAMETER.

As we will be dealing with remarkable lattices, it is often possible to have a very accurate understanding of their meaningful quantities. The smoothing parameter is no exception and is the topic of this section, of independent interest. Our main ingredient here is the general

identification of the Gaussian mass with the *theta series* of a lattice:

$$(10) \quad \begin{aligned} \rho_{1/s^2}(\Lambda^\vee) &= 1 + \kappa(\Lambda^\vee) \cdot \exp(-\pi s^2)^{\lambda_1(\Lambda^\vee)^2} + \kappa_2 \exp(-\pi s^2)^{n_2^2} + \dots \\ &= \theta_{\Lambda^\vee}(\exp(-\pi s^2)), \end{aligned}$$

where we have sorted the vectors of Λ^\vee by their increasing squared norm¹³, and $\kappa(\Lambda^\vee)$ is the *kissing number* of Λ^\vee . Let now $q = \exp(-\pi s^2)$, then determining the smoothing parameter of a lattice amounts to find q such that $\theta_{\Lambda^\vee}(q) - 1 = \varepsilon$. This can always be done by *series reversion*: there exists a series S such that $S(\theta_{\Lambda^\vee}(q) - 1) = q$. Routine calculations then show

$$s = \sqrt{\frac{1}{\pi} \ln \left(\frac{1}{S(\varepsilon)} \right)}.$$

Note that this is an *exact* expression, and formulae from formal series theory even give the coefficients of S . Of course, calculating the actual value for some lattice Λ requires knowing those of θ_{Λ^\vee} .

Let us give more details about series reversion. As mentioned, a standard fact is that if $h(z) = h_1 z + h_2 z^2 + \dots$ is a power series, then there always exists another power series $g(z) = g_1 z + g_2 z^2 + \dots$ such that $g(h(z)) = z$. In other words, the existence of the reversion is no concern. Plugging the expression of h and equating the coefficients gives an explicit formula for the coefficient of g from those of h . In the general case, the ones on the expansion of g are:

$$(11) \quad \begin{aligned} g_1 &= h_1^{-1} & g_3 &= -h_1^{-5}(2h_2^2 - h_1 h_3) \\ g_2 &= -h_1^{-3} h_2 & g_4 &= -h_1^{-7}(5h_1 h_2 h_3 - h_1^2 h_4 - 5h_2^3) \end{aligned}$$

The interested reader can find a general expression for the n -th coefficient with *Morse-Feshbach's formula*. In the general case, the Gaussian mass does not exactly correspond to a power series, even up to renormalization. This is not a concern as one can extend the reversion approach to rational exponents, see the lemma below.

Lemma 4. *Let $\Lambda \subset \mathbf{Q}^m$ be any lattice, and let $T_\Lambda(q) = \theta_\Lambda(q) - 1 = \kappa_1 q^{\lambda_1(\Lambda^\vee)^2} + \dots$. There exists a unique series $S(q) = (q/\kappa_1)^{1/\lambda_1(\Lambda^\vee)^2} + \dots$ such $S(T_\Lambda(q)) = q$.*

Proof. Since Λ is rational, there exists an integer $d > 0$ such that $d\Lambda^\vee \subset \mathbf{Z}^m$. By considering $T(q) = T_\Lambda(q^{d^2})$, we can assume without loss of generality that T_Λ is a *power series*, that is,

¹³The parameters κ_2 and n_2 are placeholders for the number κ_2 of vectors of norm n_2 , the smallest possible norm in the lattice that is larger than λ_1 .

all exponents are positive integers. We let $n = d^2 \lambda_1(\Lambda^\vee)^2$, and we consider now the formal series $T(X) = \kappa X^n + \dots = \kappa X^n(1 + \tilde{T}(X)) \in \mathbf{R}[[X]]$ defined by the coefficients of T_Λ . We can define formally the n -th root of $1 + \tilde{T}(X)$ using the coefficients of the Taylor expansion of $x \mapsto x^{1/n}$. Then, the formal series $U(X) = \kappa^{1/n} X(1 + \tilde{T}(X))^{1/n} \in \mathbf{R}[[X]]$ is such that $U(X)^n = T(X)$, and its first term is $\kappa^{1/n} X$. We can apply formal series reversion: there exists a formal series $U^{-1}(X)$ such that $U^{-1}(U(X)) = X$. Now, the series $S(X) = U^{-1}(X^{1/n}) \in \mathbf{R}[[X^{1/n}]]$ satisfies that $S(T(X)) = X$. $\ddot{::}$

Thankfully, for all exceptional lattices, the first terms of the theta series are well-known (the minima and kissing numbers of remarkable lattices are often found in [7], among others). We can use them with the previous discussion to obtain the following estimates.

Lemma 5. *For any lattice Λ , we have the following estimate, valid for $\varepsilon > 0$:*

$$\eta_\varepsilon(\Lambda) = \frac{1}{\lambda_1(\Lambda^\vee)} \cdot \sqrt{\frac{1}{\pi} \ln \left(\frac{\kappa(\Lambda^\vee)}{\varepsilon} (1 + o(1)) \right)}.$$

In particular, we have the following approximations:

- $\eta_\varepsilon(\mathbf{Z}^n) \approx \sqrt{\frac{1}{\pi} \ln \left(\frac{2n}{\varepsilon} \right)}$ and for $n \geq 5$, $\eta_\varepsilon(\mathbf{D}_n) \approx \sqrt{\frac{1}{\pi} \ln \left(\frac{2n}{\varepsilon} \right)}$;
- $\eta_\varepsilon(\mathbf{A}_n) \approx \frac{1}{\lambda_1(\mathbf{A}_n^\vee)} \cdot \sqrt{\frac{1}{\pi} \ln \left(\frac{2(n+1)}{\varepsilon} \right)} \approx \sqrt{\frac{n+1}{n}} \cdot \eta_\varepsilon(\mathbf{Z}^n)$;
- $\eta_\varepsilon(\mathbf{E}_8) \approx \frac{1}{\sqrt{2}} \cdot \sqrt{\frac{1}{\pi} \ln \left(\frac{240}{\varepsilon} \right)}$ and $\eta_\varepsilon(\mathcal{L}) \approx \frac{1}{2} \cdot \sqrt{\frac{1}{\pi} \ln \left(\frac{196560}{\varepsilon} \right)}$, where \mathcal{L} is the Leech lattice.

Proof. All these estimates are obtained following the same pattern: identify the theta series θ of the dual, find the reversion series S such that $S(\theta(z) - 1) = z$; then set $z = \exp(-\pi s^2)$ and the smoothing is reached when $\theta(z) - 1 = \varepsilon$. In our context where $z = \exp(-\pi s^2)$, we only really care about the first few terms in the expansion. As explained earlier in the section, if S is the reversion of $\rho_{1/s}(\Lambda^\vee) - 1 = \kappa z^{\lambda_1} + \dots$, where $\kappa = \kappa(\Lambda^\vee)$ is the kissing number of the dual and $\lambda_1 = \lambda_1(\Lambda^\vee)$ its minimum in the Euclidean norm, then the smoothing parameter is obtained as

$$\eta_\varepsilon(\Lambda) = \sqrt{\frac{1}{\pi} \ln \left(\frac{1}{S(\varepsilon)} \right)}.$$

From the expression of the coefficient in the reversion, one checks that

$$S(\varepsilon) = \kappa^{-1} \varepsilon^{\frac{1}{\lambda_1}} (1 + o(1)),$$

which gives our first claim:

$$S(\varepsilon)^{-1} = \kappa\varepsilon^{-\frac{1}{\lambda_1^2}}(1 + o(1))^{-1} = \kappa\varepsilon^{-\frac{1}{\lambda_1^2}}(1 + o(1)).$$

Checking the well-known theta series of these exceptional lattices in [7] is enough to obtain our estimates. $\ddot{::}$

In Lemma 4, the estimate of $\eta_\varepsilon(\mathbf{Z}^n)$ is not new, but the proof strategy we give is. The second result can be understood intuitively as follows: D_n^\vee is the disjoint union of \mathbf{Z}^n and $\mathbf{Z}^n + \frac{1}{2}\mathbf{1}$. It tells us that D_n and \mathbf{Z}^n have almost equivalent smoothing, as the first term in the theta series of their duals is the same. Additionally, $\lambda_1^\infty(D_n^\vee) = \frac{1}{2}$, so that the usual bound¹⁴ obtained from the shortest vector of the dual *in the ℓ_∞ norm* would give an overestimate by a factor ≈ 2 .

8. SAMPLING IN REMARKABLE LATTICES

This section collects various approaches to efficiently sample Gaussian in the remarkable lattices A_n , D_n and E_8 (i.e., root lattices) and the Barnes-Walls, Leech, and Nebe lattices. On the one hand, some of them will appear to be important building blocks for sampling cyclotomic integers, be incorporated in cryptographic gadget constructions, and can be seen as *base cases* or elementary functions to construct samplers on arbitrary lattices by combination (in the same way Klein's and Peikert's samplers are built around one-dimensional samplers). On the other hand, they are also a good way to illustrate practical use cases of our generic samplers from the previous sections.

8.1. Sampling in root lattices. Our ad-hoc samplers for the root lattices rely on exceptional orthogonal decompositions involving such lattices, and their close relationship in general. We will use properties of root lattices gathered in Section 2.3.

8.1.1. Sampling in roots. We start with samplers for the root lattices of small dimensions, as well as the D_n family. They are based on the ability to juggle between restrictions and extensions of lattices using Proposition 3 and Corollary 1, and exceptional isometries and geometric relations between them [22, Chap. 4.6].

¹⁴From e.g. [25, Lemma 3.5], we have $\eta_\varepsilon(\Lambda) \leq \lambda_1^\infty(\Lambda^\vee)^{-1} \cdot \eta_\varepsilon(\mathbf{Z}^n)$ for all rank n lattices. While out of the scope of the present paper, it is possible to give a bound depending on $\lambda_1(\Lambda^\vee)$ in the ℓ_2 -norm instead, *without* a \sqrt{n} loss as in [25, Lemma 3.5], *unconditionally* on ε contrary to [27, Lemma 2.6], but involving the kissing number of the dual.

Theorem 4. *We can sample efficiently and at a standard deviation right above the smoothing parameter in the following root lattices: D_n for all $n > 1$, $A_2, A_3, A_4, A_5, A_6, A_8, E_6, E_7, E_8$.*

We now detail these samplers.

8.1.2. *Sampling in the face-centered lattice D_n .* The D_n lattice can be described as the vectors of \mathbf{Z}^n with coordinates in the canonical basis $(\mathbf{e}_i)_{i \leq n}$ summing to an even number. Its index in \mathbf{Z}^n is 2. This congenial definition leads to a very natural rejection-based approach from samples over \mathbf{Z}^n : a sample either belongs to D_n or either to its non-zero coset (with almost equiprobability above the smoothing parameter). In other words, it is an instantiation of the domain restriction approach of Proposition 3, from which Proposition 8.1 is obtained.

Algorithm 6 ∴ Face centered cubic sampler

Input: A parameter $\sigma \geq \eta_\varepsilon(D_n)$ and a center $\mathbf{t} = \sum t_i \mathbf{e}_i \in \mathbf{R}^n$.

Output: $\mathbf{v} \in D_n$ following distribution statistically close to $\mathcal{D}_{D_n, \sigma, \mathbf{t}}$

1 **repeat**

2 $z_1 \leftarrow \mathcal{D}_{\mathbf{Z}, \sigma, t_1}, \dots, z_n \leftarrow \mathcal{D}_{\mathbf{Z}, \sigma, t_n}$

3 **until** $\sum z_i \in 2\mathbf{Z}$

4 **return** $\mathbf{z} = \sum z_i \mathbf{e}_i$

Proposition 8.1. *Let \mathcal{D} the distribution of outputs of Algorithm 6. Then \mathcal{D} is at statistical distance at most ε of $\mathcal{D}_{D_n, \sigma, \mathbf{t}}$. Moreover, we have*

$$\sup_{\mathbf{z} \in D_n} \left| \frac{\mathcal{D}(\mathbf{z})}{\mathcal{D}_{D_n, \sigma, \mathbf{t}}(\mathbf{z})} - 1 \right| \leq 2\varepsilon,$$

and it requires two tries in average.

8.1.3. *Sampling in the A_3 lattice.* There exists an isometry between D_3 and A_3 (coming from the exceptional Lie isomorphisms in small dimensions), which allows us to transfer the previous sampler on D_3 into a sampler for A_3 . One possible isometry consists in sending $(1, 1, 0)$ to $(1, -1, 0, 0)$, $(1, 0, 1)$ to $(1, 0, -1, 0)$ and $(0, 1, 1)$ to $(1, 0, 0, -1)$. Indeed, one checks that both these bases have the same Gram matrix by direct computation. Hence we can sample in A_3 with standard deviation above $\sigma \geq \eta_\varepsilon(A_3)$ using (an expected number of) three samples in \mathbf{Z}^4 .

Remark. *A similar algorithm would enable to sample in the dual lattice $D_3^\vee = D_3 \cup ((1/2, 1/2, 1/2) + D_3)$ (disjoint union).*

8.1.4. *Sampling in the E_8 lattice.* The E_8 lattice (also known as the Gosset lattice) is an unimodular lattice in \mathbf{R}^8 . It belongs to the class of Barnes-Wall lattices, as one of its first non-trivial member. It is not isometric to \mathbf{Z}^8 , and realizes the densest lattice sphere packing in dimension 8. As seen e.g. in [5, 22], it can be written as a disjoint union of cosets as $E_8 = D_8 \cup (1/2, \dots, 1/2) + D_8$, where D_8 is the lattice of integer vectors with coordinates summing to an even number. This leads to the following sampling algorithm, where we let $\mathbf{h} = (1/2, \dots, 1/2)$. This time, it is an instantiation of the domain extension sampler of Corollary 1, from which Proposition 8.2 is obtained.

Algorithm 7 \therefore E_8 sampler

Input: A parameter $\sigma \geq \eta_\varepsilon(D_8)$ and a center $\mathbf{t} \in \mathbf{R}^8$.
Output: \mathbf{v} following distribution statistically close to $\mathcal{D}_{D_8, \sigma, \mathbf{t}}$

- 1 $b \leftarrow \text{Bernoulli}$
- 2 $\mathbf{z} \leftarrow \text{Algorithm 6}(\sigma, \mathbf{t} - b \cdot \mathbf{h})$
- 3 **return** $b \cdot \mathbf{h} + \mathbf{z}$

Proposition 8.2. *let \mathcal{D} the distribution of outputs of Algorithm 7. Then \mathcal{D} is at statistical distance at most 2ε of $\mathcal{D}_{E_8, \sigma, \mathbf{t}}$. Moreover, we have*

$$\sup_{\mathbf{z} \in E_8} \left| \frac{\mathcal{D}(\mathbf{z})}{\mathcal{D}_{E_8, \sigma, \mathbf{t}}(\mathbf{z})} - 1 \right| \leq 4\varepsilon$$

Sampling smaller by sampling bigger: a rejection story. It is noted that $\eta_\varepsilon(D_8)$ is larger than $\eta_\varepsilon(E_8)$ by $> 30\%$. Fortunately this gap can be closed by rejection sampling, which gives a size-speed tradeoff. Concretely, to sample $\mathcal{D}_{D_8, \sigma', \mathbf{t}}$ with $\sigma' \approx \eta_\varepsilon(E_8)$, one first samples $\mathbf{v} \leftarrow \mathcal{D}_{D_8, \sigma, \mathbf{t}}$ with $\sigma = \eta_\varepsilon(D_8)$ using *Algorithm 7*, then accepts \mathbf{v} with probability $\frac{\rho_{\sigma'}(\mathbf{v} - \mathbf{t})}{\rho_\sigma(\mathbf{v} - \mathbf{t})}$ otherwise restarts. It can be verified that the output distribution is close to $\sigma' \approx \eta_\varepsilon(E_8)$ and the average number of repetitions is $N = \frac{\rho_\sigma(E_8 - \mathbf{t})}{\rho_{\sigma'}(E_8 - \mathbf{t})} \approx (\frac{\sigma}{\sigma'})^8$. For practical parameters ($\sigma = \eta_\varepsilon(D_8), \sigma' = \eta_\varepsilon(E_8)$) with $\varepsilon = 2^{-36}$, the expected repetition times is $N \approx 11$ and the gain is $\frac{\sigma'}{\sigma} \approx 0.74$.

8.1.5. *Sampling in the A_8 lattice.* There exists a special isometry between A_8 and a sublattice of E_8 . Indeed, if $(\mathbf{e}_i)_i$ denotes the canonical basis of \mathbf{R}^8 , the vectors $\mathbf{h} = (\frac{1}{2}, \dots, \frac{1}{2})$ and $\mathbf{b}_i = \mathbf{e}_i + \mathbf{e}_{i+1}$ for $1 \leq i \leq 7$ all belong to E_8 , and their Gram matrix is identical to the Gram matrix of the basis $(\mathbf{f}_1 - \mathbf{f}_i)_{1 \leq i \leq 8}$ of A_8 , where $(\mathbf{f}_i)_i$ is the canonical basis of \mathbf{R}^9 . As $\det A_8 = 3$ and

E_8 is unimodular, we expect only 3 tries in average for a sample in E_8 to belong in the lattice isometric to A_8 , which leads to the following algorithm. The acceptance criteria is obtained by identifying the sublattice isometric to A_8 , and the rest is the domain restriction approach from Section 4.1.

Algorithm 8 ∴ A8 sampler

Input: A parameter $\sigma \geq \max(\eta_\varepsilon(D_8), \eta_\varepsilon(A_8))$ and a center $\mathbf{t} \in \mathbf{R}^3$.

Output: \mathbf{v} following distribution statistically close to $D_{A_8, \sigma, \mathbf{t}}$

```

1 repeat
2   |  $\mathbf{z} \leftarrow \text{E8sampler}(\sigma, \mathbf{t})$ 
3 until  $\mathbf{z}_1 - \mathbf{z}_2 - \dots - \mathbf{z}_8 \in 3\mathbf{Z}$ 
4 return  $\mathbf{z}$ 

```

Proposition 4. *let $\mathcal{D}(\mathbf{z})$ the distribution of outputs of Algorithm 8. Then $\mathcal{D}(\mathbf{z})$ is at statistical distance at most 10ε of $D_{A_8, \sigma, \mathbf{t}}$. Moreover, we have*

$$\sup_{\mathbf{z} \in A_8} \left| \frac{\mathcal{D}(\mathbf{z})}{D_{A_8, \sigma, \mathbf{t}}(\mathbf{z})} - 1 \right| \leq 20\varepsilon$$

Proof. Let Λ be the sublattice of E_8 generated by $(\mathbf{h}, \mathbf{b}_1, \dots, \mathbf{b}_7)$ and $\mathbf{u} = \mathbf{e}_1 - \mathbf{e}_2 - \dots - \mathbf{e}_8$. We show that $\Lambda = \{\mathbf{x} \in E_8 : \langle \mathbf{x}, \mathbf{u} \rangle \in 3\mathbf{Z}\}$, which means that Algorithm 8 outputs vectors in the correct lattice. That Λ is included in this set is clear from its basis vectors. For the reverse inclusion, let \mathbf{x} be in the set, and let $a_1 = \langle \mathbf{x}, \mathbf{u} \rangle / 3 \in \mathbf{Z}$. As $\mathbf{x} \in E_8$ in particular, we also know that \mathbf{x} is either in D_8 either in $\mathbf{h} + D_8$. In the first case, we can write $\mathbf{x} = \sum_i x_i \mathbf{e}_i$ with $x_i \in \mathbf{Z}$ for all i , and such that $\sum x_i = 2k$ for some $k \in \mathbf{Z}$. and then we have $x_1 - k = \frac{3}{2}a_1 \in \mathbf{Z}$, so that a_1 is even. Therefore, for $2 \leq i \leq 8$, $a_i = x_i - a_1/2$ is an integer. In the second case, $\mathbf{x} = \sum_i (x_i + \frac{1}{2}) \mathbf{e}_i$ with $x_i \in \mathbf{Z}$ such that $\sum_i x_i = 2k$ for some $k \in \mathbf{Z}$. Then we have $x_1 - k = \frac{3}{2}(a_1 + 1) \in \mathbf{Z}$, so that a_i is odd, which means that $a_i = x_i - (a_1 + 1)/2$ is an integer. In each cases, we find that $\mathbf{x} = a_1 \mathbf{h} + a_2 \mathbf{b}_1 + \dots + a_8 \mathbf{b}_7$, or equivalently, $\mathbf{x} \in \Lambda$. The rest of the proof comes from from Proposition 3. ∴

8.1.6. *Sampling in the A_2 lattice.* The Gosset lattice E_8 contains copies of A_2 . As explained in [7, Chap. 4, Sec. 8.6], this is the case for example of the lattice generated by the vectors $u_1 = e_1 + e_8$ and $u_2 = \frac{1}{2}\mathbf{1}$, where $(e_i)_i$ is the canonical basis for \mathbf{R}^8 — it shares indeed the same Gram matrix as A_2 — but there are many others. This choice yields a linear isometry defined by $\phi_1(u_i) = b_i$,

where $b_1 = (-1, 1, 0)$ and $b_2 = (-1, 0, 1)$ is known as the Korkine-Zolotarev basis of A_2 . This serves as a definition for the root lattice E_6 as the set of vectors in E_8 orthogonal to Λ .

In other words, we have a sublattice $E_6 \perp A_2$ in E_8 , where the containment has index 3. Combining the restriction sampler from Proposition 3 with Lemma 3.2, we can obtain a sample in each lattice E_6 and A_2 every three tries in average. First use the map ϕ_1 to sent the target center in $\text{span}(u_1, u_2)$, then sample from Algorithm 7 until the output is in $E_6 \perp A_2$, and map the result back.

Although simple and efficient, this approach rather inefficiently uses the available randomness. In our applications, we need large batches of samples of A_2 . We show how to amortize the randomness expenses by using an additional exceptional decomposition. From [22, Theorem 4.6.10], there is another index 3 containment as $A_2 \perp A_2 \perp A_2 \subset E_6$. More precisely, there are isometries ϕ_2, ϕ_3, ϕ_4 from A_2 to some pairwise-orthogonal sublattices of E_6 . Then the procedure is as follow: map the target center t to $t' = \sum_i \phi_i(t)$, and sample around t' with Algorithm 7. In the orthogonal decomposition of \mathbf{R}^8 given by the four copies of A_2 , the resulting sample can be written $x = \sum_i x_i$, and we may keep any x_i belonging to (a copy of) A_2 . The correctness proof is again obtained by combining Proposition 3 with Lemma 3.2. Every 9 tries in average, we now obtain 4 samples. Expressed in terms of the underlying samples in \mathbf{Z} , this improves the randomness loss from 1 every 48 to 1 every 36, a bit better than a 30% improvement. Below we give the version that better amortizes its randomness usage.

Algorithm 9 ∴ A_2 sampler

Input: A parameter $\sigma \geq \max(\eta_\varepsilon(D_8), \eta_\varepsilon(A_8))$ and a center $\mathbf{t} \in \mathbf{R}^3$.

Output: (x_1, \dots, x_4) , each following a distribution statistically close to $D_{A_2, \sigma^2, \mathbf{t}}$

- 1 Compute the isometries ϕ_1, \dots, ϕ_4
- 2 $c \leftarrow \phi_1(\mathbf{t}) + \dots + \phi_4(\mathbf{t})$
- 3 **repeat**
- 4 $\mathbf{z} \leftarrow \mathbf{E8sampler}(\sigma^2, c)$
- 5 **until** $\mathbf{z} = z_1 + \dots + z_4 \in A_2 \perp A_2 \perp A_2 \perp A_2$
- 6 **return** $(\phi_1^{-1}(z_1), \dots, \phi_4^{-1}(z_4))$

Lastly, we deal with the question of the actual width reach by the sampler. There are no smoothing condition coming from Lemma 3.2, so we can focus our attention on the sublattice

$E_6 \perp A_2$. Using Lemma 2.2 and the definition of the smoothing parameter, we obtain $\eta_\varepsilon(E_6 \perp A_2) \leq \max(\eta_{3\varepsilon}(E_6), \eta_{3\varepsilon}(A_2))$. For small enough ε , we have $\eta_\varepsilon(E_6) \leq \eta_\varepsilon(A_2)$, which means that we can sample in A_2 at optimal width with Algorithm 9.

8.1.7. *Sampling in the E_7 lattice.* We make use of the (dual covering) $L = (E_7^\vee \perp \mathbf{Z}) = E_8 \cup (\frac{1}{2}f + E_8)$ where $f = (0, \dots, 0, 1, 1)^T \in \mathbf{Z}^8$ (see [22, Chap.4, p.118]). Hence using Algorithm 7 we can use our E_8 sampler to sample into L and project orthogonally onto E_7^\vee . We now use the fact that since E_7 is integral it is contained in its dual and use the Proposition 3. As $|E_7| = 2$, the quotient E_7^\vee/E_7 is of cardinality 4, meaning that we expect 4 repetitions of this process on average.

8.1.8. *Sampling in A_n^\vee .* The following algorithm is described for the sake of masochist pleasure, as a hurtful reminder that proofs are never checked enough. It simply uses the identification $A_n^\vee = \pi_{\mathbf{1}^\perp}(\mathbf{Z}^{n+1})$, where $\mathbf{1} = (1, \dots, 1)$. More precisely, the identification above leads to the short exact sequence $0 \rightarrow \mathbf{Z} \cdot \mathbf{1} \rightarrow \mathbf{Z}^n \rightarrow A_n^\vee \rightarrow 0$. We readily see that $\eta_\varepsilon(\mathbf{Z} \cdot \mathbf{1}) = \sqrt{n} \cdot \eta_\varepsilon(\mathbf{Z})$, giving the reachable Gaussian width. The correctness and quality follow from Lemma 3.1.

Algorithm 10 $\therefore A_n^\vee$ sampler

Input: A parameter $\sigma > 0$ and a center $\mathbf{t} \in \mathbf{1}^\perp$.

Output: $\mathbf{v} \in A_n^\vee$ following distribution statistically close to $\mathcal{D}_{A_n^\vee, \mathbf{t}, \sigma^2}$

- 1 $\mathbf{z} \leftarrow \mathcal{D}_{\mathbf{Z}^{n+1}, \mathbf{t}', \sigma^2}$
- 2 $\mathbf{y} \leftarrow \mathbf{z} - (\frac{1}{n+1} \sum_j \mathbf{z}_j)_{i=1}^{n+1}$
- 3 **return** \mathbf{y}

For completeness, we succinctly describe an algorithm in the special case $n = 2$, where we actually reach optimal width. As any rank 2 lattice, A_2 is similar to its dual A_2^\vee (we also say that they are isodual). In this case, the similarity has scaling factor $\frac{1}{\sqrt{3}}$ and the rotation fixes $\mathbf{R} \cdot \mathbf{1}$, turning the hyperplane $\mathbf{1}^\perp$ by $\pi/2$. The algorithm obtains an A_2 sample from Algorithm 7 at $\eta_\varepsilon(A_2)$, scales it and rotates it. If the target distribution has a non-trivial center, one just needs to transform it accordingly before sampling in A_2 . The rotation preserves the norm of the scaled sample, and by standard argument on discrete Gaussians, the resulting width is $\frac{1}{\sqrt{3}}\eta_\varepsilon(A_2) = \eta_\varepsilon(A_2^\vee) \approx \frac{1}{\sqrt{2}}\eta_{2\varepsilon/3}(\mathbf{Z}^2)$.

8.1.9. *Sampling in A_n lattices.* We now study the Gaussian sampling problem for arbitrary A_n lattices. The generic case is trickier, as there is no known direct isomorphisms or decompositions involving other exceptional lattices. A possible approach consists in instantiating our framework of Section 4.2.2 and Section 5.2 using the base cases we just constructed. As a point of comparison, we first briefly give the results given by the generic use of standard Klein and Peikert samplers.

8.1.10. *Trivial instantiations: Peikert and Klein samplers.* Unrolling the Cholesky algorithm on the Gram-matrix $G_n(1)$ of the standard basis $(\mathbf{e}_i - \mathbf{e}_{i+1})_{1 \leq i \leq n}$ of A_n reveals that the maximal value of its diagonal coefficients is achieved on its first element, which value is $\sqrt{2}$. Hence, the Klein sampler allows performing Gaussian sampling at standard deviation above $\sqrt{2}\eta_\varepsilon(\mathbf{Z})$. As G_n is a tridiagonal Toeplitz matrix with pattern $(-1, 2, -1)$, its eigenvalues are of the form $2 + 2 \cos(k\pi/(n+1))$ for $1 \leq k \leq n$ [19]. Consequently, the largest singular value of this basis is $(2 + 2 \cos(\pi/(n+1)))^{1/2} \geq 2\sqrt{1 - \frac{\pi^2}{2n^2}}$, a worse reachable standard deviation. The other classic basis $(\mathbf{e}_1 - \mathbf{e}_i)_{2 \leq i \leq n+1}$ of A_n has a largest singular value of $\sqrt{n+1}$, which has an even worse geometry.

8.1.11. *Constructing a better filtration.* To showcase possible trade-offs using Algorithm 3, we now describe different filtrations for A_n lattices. Our approach here is to rely on samplers in larger exceptional lattices from the previous section—such as the A_8 sampler (Algorithm 8)—as subroutines. This new family of algorithms allows sampling very close to the smoothing parameter of the A_8 . These improvements also stem from an additional ingredient: the filtrations we highlight are close to being block-orthogonal. A practical benefit yielded by such filtrations is the more parallelizable nature of the resulting processes. While the next result is straightforward, we highlight it for the sake of reusability.

Proposition 5. *Let $n > k$ be integers and $n = (k+1)q + r$ the euclidean division of n by $(k+1)$. Then A_n admits a filtration as $0 = \Lambda_0 \subset \Lambda_1 \subset \Lambda_2 \subset \dots \subset \Lambda_q \subset A_n$, where for all $1 \leq i \leq q$, Λ_i is isometric to an orthogonal direct sum of i copies of A_k .*

The proof amounts to identifying several copies of smaller A_k lattices, *orthogonal to each other*, in the standard basis of A_n , by appropriately permuting the columns (for instance the first two vectors in the usual basis of A_3 generate a copy of A_2) and packing the remaining vectors all together in the final part of the filtration.

Proof. Recall that for any $n > 1$, the A_n lattice admits the basis

$$A_n = \begin{pmatrix} 1 & 0 & 0 & \cdots & 0 \\ -1 & 1 & 0 & \cdots & 0 \\ \vdots & -1 & \ddots & \ddots & \vdots \\ 0 & 0 & \ddots & 1 & 0 \\ 0 & 0 & \cdots & -1 & 1 \\ 0 & 0 & \cdots & 0 & -1 \end{pmatrix} \in \mathbb{Z}^{(n+1) \times n}.$$

By assumption on q and k , we have in particular $n = (k+1)q + r = kq + q + r$. There exists a permutation on the columns of A_n that gives a reordered basis

$$B_n = \left(\begin{array}{c|c|c|c|c|c} A_k & & & & C_1 & \\ \hline & A_k & & & C_2 & \\ \hline & & \ddots & & \vdots & \\ \hline & & & A_k & C_q & \\ \hline & & \cdots & & C_{q+1} & A_r \end{array} \right) \in \mathbb{Z}^{(n+1) \times n},$$

where

$$C_1 = (\varepsilon_{i,j}^{(1)}) \in \mathbb{Z}^{(k+1) \times q} : \varepsilon_{i,j}^{(1)} = \begin{cases} 1 & (i,j) = (k+1, 1) \\ 0 & \text{otherwise,} \end{cases}$$

$$C_t = (\varepsilon_{i,j}^{(t)}) \in \mathbb{Z}^{(k+1) \times q} \text{ for } 1 < t \leq q : \varepsilon_{i,j}^{(t)} = \begin{cases} -1 & (i,j) = (1, t-1) \\ 1 & (i,j) = (k+1, t) \\ 0 & \text{otherwise,} \end{cases}$$

$$C_{q+1} = (\varepsilon_{i,j}^{(q)}) \in \mathbb{Z}^{(r+1) \times q} : \varepsilon_{i,j}^{(q)} = \begin{cases} -1 & (i,j) = (1, q) \\ 0 & \text{otherwise.} \end{cases}$$

Immediately, the result follows. ∴

The remaining vectors have to be dealt with, but it turns out not to impact what follows. This allows to sample over the A_n lattice using [Algorithm 3](#). Let us call B_n the basis corresponding to the filtration of [Proposition 5](#). At the deepest level of recursion, we sample in the lattice Λ/Λ_q , using for example Klein sampler, or equivalently, [Algorithm 3](#) with the filtration corresponding to the projection of the last $q+r$ columns of B_n orthogonally to $V_q^\perp = \text{Span}(\Lambda_q)^\perp$. Then, all subsequent samplings happen in (a copy of) A_k , and for example, when $k = 8$, one calls

Algorithm 8 for these last q steps. For the sake of clarity, we restrict ourselves to $k = 8$ and give an equivalent *iterative* algorithm.

Theorem 5. *Let $n > 8$ be an integer and $n = 9q + r$ the Euclidean division of n by 9. Let $t \in \mathbf{R}^n$ and \mathcal{D} be the distribution of the A_n **sampler**, for $\sigma \geq \max\{\sqrt{9/8} \cdot \eta_\varepsilon(\mathbf{Z}^8), \eta_\varepsilon(A_8)\}$. Then for a small enough ε , the statistical distance between \mathcal{D} and $\mathcal{D}_{A_n, t, \sigma^2}$ is at most $(q + 1)\varepsilon$.*

Proof of Theorem 5. Let $0 = \Lambda_0 \subset \Lambda_1 \subset \Lambda_2 \subset \dots \subset \Lambda_q \subset A_n$ be the filtration given by Proposition 5. Since we use the Klein sampler as a base sampler in the quotient Λ_q , the result holds using Theorem 2 when

$$\sigma \geq \max\left\{\sigma', \eta_\varepsilon(A_k), \max_{1 \leq i \leq q+r} \|\mathbf{b}_{kq+i}^*\| \eta_\varepsilon(\mathbf{Z})\right\}$$

where σ' is required by the sampling over A_k and \mathbf{b}_i^* is the i -th vector in the Gram-Schmidt orthogonalization of \mathbf{B}_n . We implement the sampling over A_k with the tailored A_8 sampler Algorithm 8 that just needs $\sigma' \geq \eta_\varepsilon(A_8)$. Therefore it suffices to estimate $\|\mathbf{b}_i^*\|$. A routine computation verifies that the Gram-Schmidt orthogonalization of \mathbf{A}_n is

$$\tilde{\mathbf{A}}_n = \begin{pmatrix} 1 & \frac{1}{2} & \cdots & \frac{1}{n} \\ -1 & \frac{1}{2} & \cdots & \frac{1}{n} \\ 0 & -1 & \cdots & \frac{1}{n} \\ \vdots & \vdots & \ddots & \frac{1}{n} \\ 0 & 0 & \cdots & -1 \end{pmatrix} \in \mathbf{Q}^{(n+1) \times n},$$

thus the last Gram-Schmidt norm is $\sqrt{(n+1)/n}$. Since orthogonal projections shrink norms, we have $\|\mathbf{b}_{kq+i}^*\| \leq \sqrt{(ki+1)/(ki)}$ for $1 \leq i \leq q$ and $\|\mathbf{b}_{kq+i}^*\| \leq \sqrt{(kq+i)/(kq)}$ for $q+1 \leq i \leq q+r$. Immediately the minimal achieved quality factor for the final block is bounded by $\sqrt{(k+1)/k} \cdot \eta_\varepsilon(\mathbb{Z}^{q+r})$. $\ddot{:}$

Algorithm 11 \therefore A_n sampler

Input: $\sigma \geq \max\left\{\sqrt{\frac{9}{8}}\eta_\varepsilon(\mathbf{Z}^8), \eta_\varepsilon(A_8)\right\}$, a center $\mathbf{t} \in \text{Span}_{\mathbf{R}}(A_n)$, a filtration $(\Lambda_i)_i$ of A_n in the form of \mathbf{B}_n , as in Proposition 5.

Output: \mathbf{v} following distribution statistically close to $\mathcal{D}_{A_n, \mathbf{t}, \sigma^2}$

- 1 Compute $\mathbf{c}_i = \pi_{V_q^\perp}(\mathbf{b}_{i+kq})$ for $1 \leq i \leq q+r$
- 2 $\mathbf{t}_{q+1} \leftarrow \pi_{V_q^\perp}(\mathbf{t})$
- 3 $\mathbf{x}_{q+1} \leftarrow \text{Algorithm 3}(\{\mathbf{c}_1, \dots, \mathbf{c}_{q+r}\}, \sigma, \mathbf{t}_{q+1})$
- 4 $\mathbf{u} \leftarrow \text{Lift}(\mathbf{x}_{q+1}, V_q)$
- 5 $\mathbf{t}' \leftarrow \mathbf{t} - \mathbf{x}_{q+1}$
- 6 Compute the orthogonal projections \mathbf{t}'_j of \mathbf{t}' on $\text{Span}(\mathbf{b}_{jk+1}, \dots, \mathbf{b}_{jk})$ for $0 \leq j < q$
- 7 $\mathbf{x}_1, \dots, \mathbf{x}_q \leftarrow \text{Algorithm 8}(\sigma, \mathbf{t}'_1), \dots, \text{Algorithm 8}(\sigma, \mathbf{t}'_q)$ /* can be done in parallel */
- 8 **return** $\mathbf{x}_1 + \dots + \mathbf{x}_q + \mathbf{x}_{q+1}$

8.2. The king sampler.

8.2.1. *The iterated parity-check construction.* In [8], a unifying construction is introduced to generalize the parity-check lattices. Notably, this so-called *king-construction* recovers some constructions of famous remarkable lattices, including the Barnes-Walls, Leech and Nebe lattices. By leveraging a recursive combination of the conventional parity-check construction, this technique provides a convenient coset decomposition, which the authors turn into novel decoding algorithms. In line with the philosophy that a Gaussian sampler serves as a randomized decoding algorithm, we further demonstrate how our domain extension techniques blend into the creation of new samplers for the *king*-based lattices.

The *king* construction starts from a double inclusion $L \subset M \subset N$ of full-rank lattices. Fixing coset representatives to identify the quotients $A = N/M$ and $B = M/L$, define the *king* construction $\Gamma(L, A, B, k)$ as follows, starting from its parity check sublattice is:

$$(12) \quad \Gamma(L, B, k)_P = \left\{ (t_1, \dots, t_k) \mid t_i \in L + B, \sum_i t_i \in L \right\}.$$

and its definition as the coset decomposition:

$$(13) \quad \Gamma(L, A, B, k) = \bigcup_{m \in A} \left\{ \Gamma(L, B, k)_P + (m, \dots, m) \right\}$$

8.2.2. Towards a sampler.

Reduction to sublattice sampling. This decomposition then translates to a sampler, by plugging the bricks described in Section 4.1: if we denote by Γ_0 the quotient of the k ing construction by its parity-check sublattice, the short sequence $0 \rightarrow \Gamma(L, B, k)_P \rightarrow \Gamma(L, A, B, k) \rightarrow \Gamma_0 \rightarrow 0$ underlying the decomposition of Equation (13) allows with the results of Section 3 to reduce the problem to sampling in the parity-check $\Gamma(L, B, k)_P$.

Sampling in parity check lattices. To leverage the extension results presented in Section 3 once again, we will now focus on describing $M^k/\Gamma(L, B, k)_P$ in a concise manner. By considering the map $\text{tr} : (x_1, \dots, x_k) \mapsto \sum_i x_i$ over B , we can express this quotient as $B^k/\ker(\text{tr})$. The proof is a matter of elementary commutative algebra. Just as we have the chain of abelian groups $L \subset M \subset N$, we have a chain $L^k \subsetneq \Gamma \subsetneq M^k$ where Γ denotes the parity check lattice, for ease of notations. Recall that $B = M/L$, so by standard isomorphism theorems, we readily see $M^k/L^k \simeq (M^k/L^k)/(\Gamma/L^k) \simeq B^k/(\Gamma/L^k)$. Also recall that we considered the group homomorphism $\text{tr} : (x_1, \dots, x_k) \mapsto \sum_i x_i$ over B . As we have $\Gamma/L^k = \{(t_1, \dots, t_k) \mid \sum_i t_i = 0\} = \ker \text{tr}$, we conclude by identification.

Consequently, by utilizing the short exact sampler (Algorithm 2), we can effectively reduce the task of sampling within the parity check lattice $\Gamma(L, B, k)_P$ to two main steps: sampling a uniform element in $B/\ker(\text{tr})$ and sampling an element in L . It is worth noting that sampling within the aforementioned group is equivalent to sampling a uniform k -tuple of elements with a zero-sum. This particular distribution can be accurately simulated by independently and uniformly sampling the first $k - 1$ elements, followed by setting the final element to be the negation of the sum of these previously sampled elements.

Algorithm 12 \therefore king sampler**Input:**

- A chain of (full-rank) lattices $L \subsetneq M \subsetneq N$
- Sets of representants for $A = N/L$ and $B = M/L$
- A sampler $\mathcal{D}_{L,\cdot,\Sigma}$ at covariance $\Sigma > \eta_\varepsilon(N)$.
- a center $\mathbf{c} \in N^k \otimes \mathbf{R}$

Output: \mathbf{v} following a distribution close to $\mathcal{D}_{\Gamma(L,A,B,k),\mathbf{c},\Sigma}$

```

1  $\alpha \leftarrow \mathcal{U}(A)$ 
2  $\mathbf{m} \leftarrow (\alpha, \dots, \alpha)$ 
3 for  $i = 1$  to  $k - 1$  do  $t_i \leftarrow \mathcal{U}(B)$ 
4  $t_k \leftarrow -\sum_{i=1}^{k-1} t_i$ 
5  $\mathbf{t} \leftarrow (t_1, \dots, t_k)$ 
6  $u \leftarrow \mathcal{D}_{L^k, \mathbf{c}-\mathbf{t}-\mathbf{m}, \Sigma}$ 
7 return  $(\mathbf{u} + \mathbf{t} + \mathbf{m})$ 

```

Putting all together. By these two steps, we reduced the sampling in the *king* construction to samples in the small sublattice L . The complete pseudo-code is given in Algorithm 12 when expanding all steps of the short exact sequence technique and using direct set of representants to avoid lifting:

The correctness of the *king sampler* can be derived from the proven correctness of the short exact sampler (Theorem 1). It is important to note that in this context, the selection of a suitable covariance Σ greater than the smoothing of the first lattice in the chain is mandatory (L) However, an interesting observation is that if we possess a direct sampler for the parity check lattice $\Gamma(L, B, k)_P$, we have the opportunity to optimize lines 3-4-5-6 by directly sampling from this lattice, centered at the vector $\mathbf{c} - \mathbf{m}$. By employing this optimization, the condition can be weakened, requiring only that Σ exceeds the smoothing of the parity check lattice. This modification not only simplifies the algorithm but also enhances its precision.

8.2.3. *Recovering remarkable lattices.* Following [8] we have recursive descriptions when $k = 2, 3$, with a slight abuse on notation where we allow the lattices in the chain to be rotated or scaled. The corresponding decoding/sampling algorithms are adapted *mutatis mutandis*¹⁵.

- *Barnes-Walls:* This family can be defined as $BW_{2n} = \Gamma(\phi \cdot BW_n, B, 2)$ for $B = BW_n/(1+i)BW_n$ and the bootstrap $BW_2 = \mathbf{Z}[i] \cong \mathbf{Z}^2$, where ϕ represents the multiplication by $(1+i)$ (which translates to a rotation and scaling, when looking at the underlying \mathbf{Z} -lattice).
- *Leech:* The celebrated Leech lattice can be constructed as a 3-construction (also called Turyn construction) from the E_8 root lattice. For this, we select the chain $2E_8 \subseteq T_\theta \subset T$ where $T_\theta \cong T \cong E_8$, for a phase such that $\sqrt{2}e^{i\theta} = \frac{1}{2}(1+i\sqrt{7})$ (see [8] for a complete description). Other variants of the same approach exist.

9. APPLICATION I: IMPROVED SAMPLERS FOR MITAKA

MITAKA [15] is a variant of FALCON offering simpler implementations and flexible parameters. It can be theoretically instantiated over *arbitrary* cyclotomic fields. While concrete parameters and security estimates are provided, the preliminary implementation of MITAKA only covers the case of power-of-2 cyclotomics. The instantiation over other cyclotomic rings \mathcal{R}_m relies on how the Gaussian sampling over \mathcal{R}_m is performed. This is non-trivial as the canonical basis of these rings of integers fails to be orthogonal when the conductor m is not a power-of-2.

In this section, we present two novel approaches relying on our ad-hoc, explicit samplers for root lattices: one for cyclotomic rings with *prime* conductor, one for *smooth* conductor $m = 2^\ell 3^k$. We believe that the techniques introduced in this section could find further use in designs, providing more flexible parameters, more efficient samplers, and tighter security.

9.1. Hybrid sampling and representation of cyclotomic numbers. MITAKA is an NTRU-based instantiation of the GPV framework [18]. Its secret key is a short basis $\mathbf{b}_0 = (f, g)^t$, $\mathbf{b}_1 = (F, G)^t \in \mathcal{R}_m^2$ of the NTRU module $\mathcal{M}_{\text{ntru}} = (f, g)^t \mathcal{R}_m \oplus (F, G)^t \mathcal{R}_m$. The signing amounts to sampling a discrete Gaussian in $\mathcal{M}_{\text{ntru}}$ close to an arbitrary target (a hashed message), which is accomplished by the *hybrid sampler* [12, 29]. Let σ_{sig} be the standard deviation of the sampled lattice Gaussian. For better sizes and security against forgery, one wants to minimize σ_{sig} .

¹⁵We don't describe the construction of the Nebe here as it will not be used in the following practical applications. However, we point out that it is based also on 3-ing construction upon the Leech lattice, and as such our framework readily applies.

As seen in Section 4.3, the hybrid sampler leverages the filtration $\{0\} \subset \psi(\mathbf{b}_0\mathcal{R}_m) \subset \psi(\mathcal{M}_{\text{ntru}})$, where ψ denotes the canonical embedding extended to vectors. The calls to Algorithm 4 consider $\mathbf{b}_0\mathcal{R}_m$ and $\mathcal{M}_{\text{ntru}}/\mathbf{b}_0\mathcal{R}_m$ as linear transformations¹⁶ of \mathcal{R}_m . Under this identification, Lemma 1 and Lemma 2 show that the sampler of MITAKA reaches standard deviation as

$$\sigma_{sig} \geq \max(s_1(\psi(\mathbf{b}_0)), s_1(\psi(\mathbf{b}_1^*))) \cdot \alpha \cdot \eta_\varepsilon(\psi(\mathcal{R}_m)),$$

where $\alpha > 1$ encodes how close we are able to sample from the smoothing parameter of the base ring \mathcal{R}_m . For Algorithm 3 to reach the stated covariance, it requires two *elliptic* samples in \mathcal{R}_m . In [15], this is handled by Peikert’s sampler in $\psi(\mathcal{R}_m)$, or equivalently, Algorithm 4 with \mathbf{C} being the (canonical embedding of) the power basis — equivalently, the Fourier domain of this basis. This choice comes from the use of a well-chosen continuous perturbation, which has diagonal covariance in this representation so that its square root can be computed in quasi-linear time, avoiding costly Cholesky decompositions.

The next requirement of Algorithm 4 is a *spherical, discrete* sample in \mathcal{R}_m . For *power-of-two* cyclotomics, the canonical embedding $\psi(\mathcal{R}_m)$ is essentially a scaling of $\mathbf{Z}^{m/2}$, and so $\alpha = 1$. The situation is less favorable for more general cyclotomic rings. For example in prime cyclotomic, sampling directly the coefficients of $x = \sum_j x_j \zeta^j$ as spherical Gaussians means that $\psi(x)$ has covariance (proportional to) $V_p \overline{V_p}^t$, a matrix far from being diagonal. In other words, going back and forth the canonical embedding distorts severely the resulting sample in $\mathcal{M}_{\text{ntru}}$. Another approach is to sample directly in the Fourier domain; for prime or smooth conductors, the current best approaches yield $\alpha = \sqrt{p-1}$ and $\alpha = \sqrt{2}$ losses, respectively [15].

Changing the construction of the basis of $\mathcal{M}_{\text{ntru}}$ is not the topic of this paper. We focus instead on decreasing the contribution of α . Our goal is to show that a *different representation* of \mathcal{R}_m can significantly reduce this parameter. The hero of the story is the principal ideal¹⁷ $\langle 1 - \zeta_p \rangle$. Using Algorithm 3 over the filtration induced by the so-called *decoding basis* [20] $\zeta_p^i - \zeta_p^{i+1}$ of the ideal $\langle 1 - \zeta_p \rangle$, one can achieve generally, $\alpha = \sqrt{2}$, which is the length of the largest Gram-Schmidt vector of this basis. We will now push further thanks to Proposition 1 and other identifications.

¹⁶In practice, this second call is encoded by the orthogonalization \mathbf{b}_1^* of \mathbf{b}_1 in the cyclotomic field

¹⁷It is also known as (a scaling of) the co-different ideal.

9.2. **Sampling over cyclotomic fields of conductor $2^\ell \cdot 3^k$.** Here, we work in $\mathcal{R}_m = \mathbf{Z}[\zeta_m]$ with $m = 2^\ell \cdot 3^k$ and $\ell, k > 0$, as suggested in [15]. To our knowledge, very few works focus¹⁸ on such conductors. From e.g. [20, 31], the tensor decomposition $\mathcal{R}_m = \mathcal{R}_{2^\ell} \otimes \mathcal{R}_{3^k}$ leads to an *orthogonal* decomposition (tied to the *powerful basis* [20]) $\mathcal{R}_m \cong \mathbf{Z}^{\frac{m}{2}} \otimes (\bigoplus_{i=1}^{3^k-1} \mathcal{R}_3) \cong \bigoplus_{i=1}^{\frac{m}{6}} \mathcal{R}_3$. Alternatively, we have an orthogonal decomposition

$$(14) \quad \langle 1 - \zeta_m^{m/3} \rangle = \frac{m}{2} \mathcal{R}_m^\vee \cong \bigoplus_{i=1}^{m/6} \langle 1 - \zeta_3 \rangle.$$

(see also [20, Cor. 2.18]). We can use Algorithm 9 with this decomposition: sampling in \mathcal{R}_3 is done by $m/6$ independent sampling in A_2 , by orthogonality.

9.2.1. *Efficiency and signature quality.* From Algorithm 9, we obtain samples in A_2 of width at least $\eta_\varepsilon(D_8)$, for some chosen ε . We observe using the estimates of Lemma 4 that $\eta_\varepsilon(D_8) \leq \eta_\varepsilon(A_2)$, so we can sample each component in the decomposition (14) *at* the smoothing of A_2 . Taking into account the entire filtration, the resulting sampler reaches standard deviation starting

$$(15) \quad \sigma' = \eta_{6\varepsilon/m}(A_2) \approx \cdot \sqrt{\frac{3}{2}} \cdot \eta_{2\varepsilon/3}(\mathbf{Z}^{\frac{m}{3}}).$$

The running time is linear in the conductor m . As we need large batches of samples, the alternate approach of Algorithm 9 that amortizes randomness is a good choice here. While still a bit randomness-hungry, the resulting sampler is completely parallelizable and also memory-efficient: we only need to store a table for integer Gaussians of small width. Moreover, thanks to the small Gaussian parameter, the constant-time implementation is easy and efficient.

9.2.2. *Comparisons with other methods.* On the one hand, the basis $\mathbf{b}_0, \mathbf{b}_1$ is not changed between our methods and the previous ones. On the other hand, previous approaches such as [15] could only sample representants of \mathcal{R}_m to a standard deviation of $\sigma' \geq \sqrt{2} \cdot \eta_\varepsilon(\mathbf{Z}^{m/3})$. This translates quantitatively into a NIST security¹⁹ level-up for each 3-smooth conductors parameter sets proposed in [15], as reported in Table 1.

Another generic method for low-dimensional lattices with a small width is *tabulated sampling*. Concretely, one precomputes a CDT-like table for all short vectors of A_2 and then outputs

¹⁸FALCON showcased an FFO-style sampler over cyclotomic rings of conductor $3 \cdot 2^\ell$ in the round 1 of the NIST call. It was abandoned because of its high technicality. Such rings are also the focus of the implementation in [21].

¹⁹We use here the same security estimates as in [15], in the so-called Core-SVP model for a fair comparison.

TABLE 1. Concrete values for forgery compared to Mitaka base sampler.

	MITAKA			This work		
	Classical	Quantum	NIST Level	Classical	Quantum	NIST Level
$d = 648$	117	103	I ⁻	137	121	II
$d = 768$	147	129	II	170	150	III
$d = 864$	168	148	III	195	171	IV
$d = 972$	194	170	IV	224	197	V

the sample through table look-up. However, the size of the table for $\mathcal{D}_{A_2, \sigma'^2}$ is much larger than the one for $\mathcal{D}_{\mathbf{Z}, \sigma'^2/3}$ in our algorithm, which significantly lowers the speed of the constant-time implementation.

9.3. Sampling over prime cyclotomic fields. The sampler results from a combination of Proposition 1 with our efficient Algorithm 11 instantiated over A_{p-1} .

9.3.1. *Efficiency and signature quality.* Both approaches are linear in p , with Algorithm 11’s main cost coming from the sampling in A_8 (Algorithm 8), achieving a width $\max(\eta_\varepsilon(D_8), \eta_\varepsilon(A_8))$. Using the approximation of Lemma 4, we see that $\eta_\varepsilon(D_8) \leq \eta_\varepsilon(A_8)$: we can sample the components at the smoothing of A_8 . The resulting standard deviation in Algorithm 11 is thus

$$(16) \quad \sigma' = \eta_{\varepsilon/q}(A_8) \approx \sqrt{\frac{9}{8}} \cdot \eta_{2\varepsilon/9}(\mathbf{Z}^{8q}) \text{ with } q = \lfloor p/9 \rfloor.$$

The isochronous implementation for both approaches is easy and efficient, as the involved algorithms only rely on an integer sampler of a fixed width and simple rejection samplings. They are both highly parallelizable, thanks to the filtration shown in Proposition 5; and memory-efficient, as the base sampling has small width $\sqrt{9/8} \cdot \eta_\varepsilon(\mathbf{Z})$ and it does not need to store many intermediate values due to the parallelism.

9.3.2. *Comparisons with other methods.* In [20, Sec. 6.3], the ideal $\langle 1 - \zeta_p \rangle$ and the identification of prime-power cyclotomic rings were used to sample *continuous* Gaussians, by mean of the so-called “decoding basis”, which is the \mathbf{Z} -basis of the ideal. From the map ϕ , we can directly identify the Gram matrix of \mathcal{R}_p as a scaling by p of that of A_{p-1}^\vee to be the circulant matrix G_p of first line $(p-1, -1, \dots, -1)$. The largest element in the diagonal of the Cholesky of G_p is

TABLE 2. Comparisons with other samplers over prime cyclotomics.

	Quality	Running time
Peikert, canonical basis	$\sqrt{p} \cdot \eta_\varepsilon(\mathbf{Z}^{p-1})$	$O(p^2)$
Klein, canonical basis	$\sqrt{p-1} \cdot \eta_\varepsilon(\mathbf{Z}^{p-1})$	$O(p^2)$
Peikert, decoding basis	$\approx 2\sqrt{1 - \frac{\pi}{2p^2}} \cdot \eta_\varepsilon(\mathbf{Z}^{p-1})$	$O(p)$
Klein, decoding basis	$\sqrt{2} \cdot \eta_\varepsilon(\mathbf{Z}^{p-1})$	$O(p)$
Coefficient embedding	$\eta_\varepsilon(\mathbf{Z}^{p-1})$	$O(p)$
Ours (ϕ)	$\eta_{\varepsilon/q}(\mathbf{A}_8) \approx \sqrt{\frac{9}{8}} \cdot \eta_{2\varepsilon/9}(\mathbf{Z}^{8q})$	$O(p)$

$\sqrt{p-1}$, which drives the quality of a Klein/GPV approach (as done in [15]). An approach *à la Peikert* with Algorithm 4 is driven by the Vandermonde matrix V_p , and we have $s_1(V_p) = \sqrt{p}$. Considering now the decoding basis as a matrix A_p with Gram matrix $G_p(1)$ (equivalently, using the map ϕ), we have identified the meaningful quantities in Section 8.1.9. Comparisons between all approaches are displayed in Table 2, showing that our filtration choice improves on the state-of-the-art.

9.3.3. *Practical impact.* The improvement over MITAKA is significant, as seen in Table 3. On the one hand, we can use finely tailored conductors to match the requirements of the NIST level, which allows working in smaller dimensions. But we can also use NTT-friendly moduli q that are smaller than the traditional $q = 12289$ used for power-of-two cyclotomics, which also allows for reducing both the public key and signature sizes – however the improvement is mild, therefore we focus on the security.

10. APPLICATION II: NEW COMPACT LATTICE GADGETS

Lattice gadgets are an important ingredient to build efficient lattice trapdoors. Very recently, Yu, Jia, and Wang developed a new gadget framework [32] and proposed practical signature schemes based on it. In their scheme, the gadget can be in principle any square matrix supporting efficient decoding and Gaussian sampling. However, the concrete instantiations only use the simplest (scaled) integer lattice \mathbb{Z}^n as the gadget. To design more efficient gadgets, one is not only required to explore new lattice structures but also to develop specialized decoding and/or sampling algorithms.

TABLE 3. Intermediate parameters and security levels for prime-Mitaka.

	Conductor $m : \varphi(m)$	Modulus q	Quality α	Security (C/Q/NIST level)
MITAKA	2304 : 768	18433	2.20	167/151/NIST-II
This work	683 : 682	1367	2.125	157/138/NIST-II
MITAKA	2592 : 864	10369	2.25	192/174/NIST-III
This work	857 : 856	6857	2.215	207/182/NIST-III
MITAKA	2916 : 972	17497	2.30	220/199/NIST-IV
This work	919 : 918	3677	2.247	223/196/NIST-IV
FALCON	2048 : 1024	12289	1.17	285/258/NIST-V
MITAKA	2048 : 1024	12289	2.33	233/211/NIST-V
This work	1009 : 1008	10091	2.30	250/219/NIST-V

This section showcases a new practical construction based on the E_8 lattice. As shown in Section 8.1.4, the E_8 lattice has an efficient sampler achieving the Gaussian width $\eta_\varepsilon(E_8) < \eta_\varepsilon(\mathbb{Z}^8)$, which allows better size and security than the \mathbb{Z}^n -based instantiation.

10.1. The Yu-Jia-Wang compact gadget framework. In a general gadget-based trapdoor scheme, the preimage sampling is converted to the gadget sampling by using the trapdoor, following the idea of the Micciancio-Peikert trapdoor [23]. In the Yu-Jia-Wang framework, the gadget is a square matrix $\mathbf{G} \in \mathbb{Z}^{n \times n}$ along with $\mathbf{H} \in \mathbb{Z}^{n \times n}$ such that $\mathbf{GH} = q\mathbf{I}_n$, and the gadget sampling is implemented by the so-called *semi-random* sampler. Given a target \mathbf{u} , the semi-random sampler outputs a Gaussian preimage \mathbf{v} such that $\mathbf{Gv} = \mathbf{u} - \mathbf{e} \pmod q$ for a small error \mathbf{e} . The sampler proceeds in two steps:

- (1) *Deterministic decoding over $\mathcal{L}(\mathbf{G})$.* The sampler first computes an error \mathbf{e} such that $\mathbf{u} - \mathbf{e} \in \mathcal{L}(\mathbf{G})$. Let $\mathbf{u} - \mathbf{e} = \mathbf{Gc}$ for some $\mathbf{c} \in \mathbb{Z}^n$.
- (2) *Gaussian sampling over $\mathcal{L}(\mathbf{H})$.* The sampler then samples $\mathbf{v} \leftarrow D_{\mathcal{L}(\mathbf{H})+\mathbf{c},r}$ with $r \geq \eta_\varepsilon(\mathcal{L}(\mathbf{H}))$. It holds that $\mathbf{Gv} = \mathbf{u} - \mathbf{e} \pmod q$.

Note that either shorter \mathbf{e} or shorter \mathbf{v} can be beneficial to the security. The concrete instantiations in [32] choose the simplest gadget $(\mathbf{G}, \mathbf{H}) = (p\mathbf{I}_n, (q/p)\mathbf{I}_n)$. In this setting, $\|\mathbf{e}\|$ is around $p\sqrt{n/12}$ and $\|\mathbf{v}\|$ is around $(q/p) \cdot \eta_\varepsilon(\mathbb{Z}^n)$.

10.2. Compact gadget from the E_8 lattice. The E_8 lattice has a good basis with well-structured inverse as follows

$$\mathbf{B} = \begin{pmatrix} 1 & 1 & 0 & 0 & 0 & 0 & 0 & 1/2 \\ -1 & 1 & 1 & 0 & 0 & 0 & 0 & 1/2 \\ 0 & 0 & 1 & 1 & 0 & 0 & 0 & 1/2 \\ 0 & 0 & 0 & 1 & 1 & 0 & 0 & 1/2 \\ 0 & 0 & 0 & 0 & 1 & 1 & 0 & 1/2 \\ 0 & 0 & 0 & 0 & 0 & 1 & 1 & 1/2 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 1/2 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1/2 \end{pmatrix} \mathbf{B}^{-1} = \frac{1}{2} \cdot \begin{pmatrix} 1 & -1 & 1 & -1 & 1 & -1 & 1 & -1 \\ 1 & 1 & -1 & 1 & -1 & 1 & -1 & -1 \\ 0 & 0 & 2 & -2 & 2 & -2 & 2 & -2 \\ 0 & 0 & 0 & 2 & -2 & 2 & -2 & 0 \\ 0 & 0 & 0 & 0 & 2 & -2 & 2 & -2 \\ 0 & 0 & 0 & 0 & 0 & 2 & -2 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 2 & -2 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 4 \end{pmatrix}.$$

In our E_8 -based instantiation, the gadget $(\mathbf{G}, \mathbf{H}) = (p \cdot \mathbf{B}^{-1} \otimes \mathbf{I}_{n/8}, (q/p) \cdot \mathbf{B} \otimes \mathbf{I}_{n/8})$.

The lattice $\mathcal{L}(\mathbf{B}^{-1})$ has an orthogonal basis

$$\mathbf{D} = \frac{1}{2} \cdot \begin{pmatrix} 1 & -1 & & & & & & \\ & 1 & 1 & & & & & \\ & & & 2 \cdot \mathbf{I}_6 & & & & \\ & & & & & & & 4 \end{pmatrix}$$

that offers fast decoding. The error size is around $\sqrt{\frac{n}{8}} \cdot \sqrt{\frac{1}{24} + \frac{6}{12} + \frac{1}{3}} \approx 0.33\sqrt{n}$ slightly larger than $\sqrt{n/12} \approx 0.29\sqrt{n}$ the one in the \mathbb{Z} -based instantiation. The sampling on $\mathcal{L}(\mathbf{H})$ is tighter than that in the \mathbb{Z} -base gadget: by using the ad-hoc sampler shown in Section 8.1.4, one can sample with a Gaussian width $\eta_\varepsilon(E_8) \approx \frac{1}{\sqrt{2}} \cdot \eta_\varepsilon(\mathbb{Z}^8)$ (scaled by q/p). As shown in Table 4, the E_8 gadget gains 56 (resp. 113) bytes of the signature size and 3 (resp. 9) bits of security against forgery for EAGLE (the Ring-LWE-based scheme). The comparisons follow the suggested parameters in [32] where the secret standard deviation is $\sqrt{1/2}$. We note that the presented parameters may not be optimal: there seems a few bits of gap between the key recovery and the forgery security. Finding optimal parameters and closing this gap could be interesting future work.

Towards improving LIP based schemes. The Lattice Isomorphism Problem has been introduced in [13] and is a promising class of problems to construct efficient cryptography. More precisely the LIP-based hash-and-sign signature of [13] instantiates with any lattice for which

TABLE 4. Parameters for the compact gadget-based signatures.

	$(p, q/p)$	Forgery security (C/Q)	Sig. (bytes)
\mathbb{Z} -based EAGLE-512	(2000, 8)	83 / 75	1406
E_8 -based EAGLE-512	(2000, 8)	86 / 78	1350
\mathbb{Z} -based EAGLE-1024	(2700, 12)	189 / 172	3052
E_8 -based EAGLE-1024	(2700, 12)	198 / 179	2939

we can sample efficiently at small Gaussian widths, and for which lattice attacks are as hard as they can be. The preliminary analysis of the problem by the authors of [13] has revealed that the Barnes-Walls lattices could be good candidates if a sharp sampler is available. Using the 2-ing sampler Section 8.2 recursively allows us to instantiate the LIP scheme with any of the BW_{1024} or BW_{2048} lattice. The precise optimization and benchmark of the resulting scheme would be totally out of the scope of this paper and we let it as very promising future work.

Acknowledgments: We thank the anonymous reviewers for their valuable suggestions and pointing out a flawed proof in an earlier version of this work. Alexandre Wallet was supported by PEPR quantique France 2030 programme (ANR-22-PETQ-0008) and by the ANR ASTRID project AMIRAL (ANR-21-ASTR-0016).

REFERENCES

- [1] Divesh Aggarwal and Noah Stephens-Davidowitz. Just take the average! an embarrassingly simple 2^n -time algorithm for svp (and cvp). In *1st Symposium on Simplicity in Algorithms (SOSA 2018)*. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, 2018.
- [2] Shweta Agrawal. Stronger security for reusable garbled circuits, general definitions and attacks. In Jonathan Katz and Hovav Shacham, editors, *CRYPTO 2017, Part I*, volume 10401 of LNCS, pages 3–35. Springer, Heidelberg, August 2017.
- [3] L. Babai. On lovasz’ lattice reduction and the nearest lattice point problem. *Combinatorica*, 6(1):1–13, 1986.
- [4] Ilaria Chillotti, Nicolas Gama, Mariya Georgieva, and Malika Izabachène. TFHE: Fast fully homomorphic encryption over the torus. *Journal of Cryptology*, 33(1):34–91, January 2020.
- [5] J. Conway and N. Sloane. Fast quantizing and decoding and algorithms for lattice quantizers and codes. *IEEE Transactions on Information Theory*, 28(2):227–232, 1982.
- [6] J. Conway and N. Sloane. A fast encoding method for lattice codes and quantizers. *IEEE Transactions on Information Theory*, 29(6):820–824, 1983.

- [7] J. H. Conway and N. J. A. Sloane. *Sphere Packings, Lattices and Groups*. Grundlehren der Mathematischen Wissenschaften 290. Springer-Verlag, New York, 1988.
- [8] Vincent Corlay, Joseph J. Boutros, Philippe Ciblat, and Loïc Brunel. On the decoding of lattices constructed via a single parity check. *IEEE Transactions on Information Theory*, 2022.
- [9] Léo Ducas, Eike Kiltz, Tancreède Lepoint, Vadim Lyubashevsky, Peter Schwabe, Gregor Seiler, and Damien Stehlé. CRYSTALS-Dilithium: A lattice-based digital signature scheme. *IACR TCHES*, 2018(1):238–268, 2018. <https://tches.iacr.org/index.php/TCHES/article/view/839>.
- [10] Léo Ducas, Vadim Lyubashevsky, and Thomas Prest. Efficient identity-based encryption over NTRU lattices. In Palash Sarkar and Tetsu Iwata, editors, *ASIACRYPT 2014, Part II*, volume 8874 of *LNCS*, pages 22–41. Springer, Heidelberg, December 2014.
- [11] Léo Ducas, Eamonn W. Postlethwaite, Ludo N. Pulles, and Wessel P. J. van Woerden. Hawk: Module LIP makes lattice signatures fast, compact and simple. *IACR Cryptol. ePrint Arch.*, page 1155, 2022.
- [12] Léo Ducas and Thomas Prest. Fast Fourier Orthogonalization. In *ISSAC 2016*, pages 191–198, 2016.
- [13] Léo Ducas and Wessel P. J. van Woerden. On the lattice isomorphism problem, quadratic forms, remarkable lattices, and cryptography. In Orr Dunkelman and Stefan Dziembowski, editors, *EUROCRYPT 2022, Part III*, volume 13277 of *LNCS*, pages 643–673. Springer, Heidelberg, May / June 2022.
- [14] Léo Ducas and Wessel PJ van Woerden. The closest vector problem in tensored root lattices of type a and in their duals. *Designs, Codes and Cryptography*, 86(1):137–150, 2018.
- [15] Thomas Espitau, Pierre-Alain Fouque, François Gérard, Mélissa Rossi, Akira Takahashi, Mehdi Tibouchi, Alexandre Wallet, and Yang Yu. MITAKA: A Simpler, Parallelizable, Maskable Variant of FALCON. In *Eurocrypt 2022*, 2022.
- [16] Thomas Espitau and Paul Kirchner. The nearest-colattice algorithm. *ANTS 2020*, 2020.
- [17] Craig Gentry. Fully homomorphic encryption using ideal lattices. In Michael Mitzenmacher, editor, *41st ACM STOC*, pages 169–178. ACM Press, May / June 2009.
- [18] Craig Gentry, Chris Peikert, and Vinod Vaikuntanathan. Trapdoors for hard lattices and new cryptographic constructions. In Richard E. Ladner and Cynthia Dwork, editors, *40th ACM STOC*, pages 197–206. ACM Press, May 2008.
- [19] M.J.C. Gover. *The eigenproblem of a tridiagonal 2-Toeplitz matrix*, volume 198 of *Linear algebra and its applications*. Elsevier, 1994.
- [20] Vadim Lyubashevsky, Chris Peikert, and Oded Regev. A toolkit for ring-LWE cryptography. In Thomas Johansson and Phong Q. Nguyen, editors, *EUROCRYPT 2013*, volume 7881 of *LNCS*, pages 35–54. Springer, Heidelberg, May 2013.
- [21] Vadim Lyubashevsky and Gregor Seiler. NTTRU: Truly fast NTRU using NTT. *IACR TCHES*, 2019(3):180–201, 2019. <https://tches.iacr.org/index.php/TCHES/article/view/8293>.
- [22] Jacques Martinet. *Perfection and Eutaxy*, pages 67–108. 2003.
- [23] Daniele Micciancio and Chris Peikert. Trapdoors for lattices: Simpler, tighter, faster, smaller. In David Pointcheval and Thomas Johansson, editors, *EUROCRYPT 2012*, volume 7237 of *LNCS*, pages 700–718. Springer, Heidelberg, April 2012.

- [24] Daniele Micciancio and Chris Peikert. Hardness of SIS and LWE with small parameters. In Ran Canetti and Juan A. Garay, editors, *CRYPTO 2013, Part I*, volume 8042 of *LNCS*, pages 21–39. Springer, Heidelberg, August 2013.
- [25] Chris Peikert. Limits on the hardness of lattice problems in l_p norms, 2008.
- [26] Chris Peikert. An efficient and parallel Gaussian sampler for lattices. In Tal Rabin, editor, *CRYPTO 2010*, volume 6223 of *LNCS*, pages 80–97. Springer, Heidelberg, August 2010.
- [27] Chris Peikert, Oded Regev, and Noah Stephens-Davidowitz. Pseudorandomness of ring-LWE for any ring and modulus. In Hamed Hatami, Pierre McKenzie, and Valerie King, editors, *49th ACM STOC*, pages 461–473. ACM Press, June 2017.
- [28] Chris Peikert and Sina Shiehian. Noninteractive zero knowledge for NP from (plain) learning with errors. In Alexandra Boldyreva and Daniele Micciancio, editors, *CRYPTO 2019, Part I*, volume 11692 of *LNCS*, pages 89–114. Springer, Heidelberg, August 2019.
- [29] Thomas Prest. *Gaussian Sampling in Lattice-Based Cryptography*. PhD thesis, École Normale Supérieure, Paris, France, 2015.
- [30] Thomas Prest, Pierre-Alain Fouque, Jeffrey Hoffstein, Paul Kirchner, Vadim Lyubashevsky, Thomas Pornin, Thomas Ricosset, Gregor Seiler, William Whyte, and Zhenfei Zhang. Falcon: Submission to the NIST’s post-quantum cryptography standardization process. <https://csrc.nist.gov/Projects/post-quantum-cryptography/round-3-submissions>.
- [31] Wessel PJ van Woerden. *The closest vector problem in cyclotomic lattices*. PhD thesis, Leiden University, 2016.
- [32] Yang Yu, Huiwen Jia, and Xiaoyun Wang. Compact lattice gadget and its applications to hash-and-sign signatures. In *CRYPTO 2023*, pages 390–420, 2023.

* PQSHIELD, FRANCE, * UNIV RENNES, INRIA, IRISA, CNRS, FRANCE, † TSINGHUA UNIVERSITY, BEIJING, CHINA

Email address: t.espitau@gmail.com, alexandre.wallet@inria.fr, yang.yu0986@gmail.com