ON GAUSSIAN SAMPLING, SMOOTHING PARAMETER AND APPLICATION TO 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 applications, 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 basis, or even switching the ambient space, even when the target lattice is not represented as full-rank in the ambiant space. We show how to correct the induced distortion with the "convolution-like" technique of Peikert ([29], 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* $\mathsf{A}_n, \mathsf{D}_n, \mathsf{E}_n$ as base cases, extending the state-of-the-art for root lattice sampling, which was limited to $\mathbf{Z}^n$.

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.

Lastly, we sprinkle our exposition with several new estimates for the smoothing parameter of lattices, stemming from our algorithmic constructions.

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 [10, 33, 15], identity-based encryption [19], functional encryption [2], (non-interactive) zero-knowledge proofs [31] and last but not least fully homomorphic encryption [18, 6]. 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 [24], *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) [19] and the *Peikert* sampler [29], 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 sampling 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 samplable 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 [19, 11].

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 [33] and DILITHIUM [10], epitomizing two known classes[1] 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 to reach all NIST security levels. MITAKA relies on a so-called hybrid sampler [32], 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.

Contributions. In this work, we aim at trespassing 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 where 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 an unifying thread in all our exposition. To complete our modular framework, we also provide ad-hoc samplers of essentially optimal widths for root lattices, in order to use them as fundamental blocks to instantiate more involved samplers. As an application, we obtain novel, optimal and efficient samplers over cyclotomic rings of prime and smooth conductors. We also study their impact on hash-and-sign type signatures. 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:

$$0 \longrightarrow \Lambda' \longrightarrow \Lambda \longrightarrow {}^{\Lambda}\!/_{\Lambda'} \longrightarrow 0.$$

[1] Recently, the scheme `Hawk` [12] presented a new intriguing signature paradigm tied to lattices.

Note that in this sequence, the quotient Λ/Λ' is not necessarily a lattice[2] 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

$$\eta_{5\varepsilon}(\Lambda) \leqslant \max\Big(\eta_\varepsilon(\Lambda'), \eta_\varepsilon\Big({}^{\Lambda}\!\big/_{\Lambda'}\Big)\Big),$$

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 $0 \subset \Lambda_1 \subset \cdots \subset \Lambda_k = \Lambda$. Iterating the previous construction gives us a generic sampler for Λ. Namely, we have a first short exact sequence stemming from the filtration as:

$$0 \longrightarrow \Lambda_1 \longrightarrow \Lambda \longrightarrow {}^{\Lambda}\!\big/_{\Lambda_1} \longrightarrow 0,$$

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 $0 \subset \Lambda_2/\Lambda_1 \subset \cdots \subset \Lambda_k = \Lambda/\Lambda_1$. 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 [19]), which corresponds to the particular case where $\mathrm{rk}(\Lambda_i) = i$ for all $1 \leqslant i \leqslant \mathrm{rk}(\Lambda)$, 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 [19]:

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

[2] Generally, 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.

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 well[3]. 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 normal[4] perturbation. This algorithm encompasses the sampler of Peikert [29], which reduces sampling in a lattice Λ to sampling *spherically* in $\mathbf{Z}^{\operatorname{rk}\Lambda}$ — 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 $\mathbf{TC}$:

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

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.

Sampling in root 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

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

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

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 and their duals, 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 [7, 8], 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 essentially *at the smoothing* of these lattices.

Cyclotomic sampling, 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 [21, 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 [33] 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 regime) and slightly shorter signatures for free compared to [15]. 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 behaviour of the ratio security-level over cyclotomic-conductors compared to [15]. These new data are gathered in Table 3.

Organization of the paper. After recalling some material about lattices and Gaussian measures in Section 2, we start with the first, central piece of our framework in Section 3: the sampling procedure over *short exact sequences* (Algorithm 1), and its natural recursive extension, the *filtered Sampler* (Algorithm 2). In Section 4, we present our *linear Sampler*; because of space constraints, its use for tensor sampling is postponed to Supplementary Material F. Section 5 is devoted to our samplers for (sometimes low dimensional) root lattices. Last, in Section 6, we instantiate many of our contributions into a hash-then-sign signature scheme with concrete parameters and analysis.

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 concatenatuion 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 $\mathbf{T}^\star = (\mathbf{T}^t\mathbf{T})^{-1}\mathbf{T}^t$.

2.1. Euclidean lattices.

A (real) *lattice* Λ is a finitely generated free $\mathbf{Z}$-module, endowed with a Euclidean norm $\|.\|$ 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 $\mathrm{rk}(\Lambda)$, 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 $\langle , \rangle$, 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 $\Lambda^\vee$; if it is not full rank, $\mathbf{B}(\mathbf{B}^t\mathbf{B})^{-1}$ is a basis of $\Lambda^\vee$.

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 $\pi_{V^{\perp}}$ 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 $\pi_{V^{\perp}}$, we have $\mathbf{P}^2 = \mathbf{P}$ and $\mathbf{P}^t = \mathbf{P}$. Given a basis $\mathbf{B} = (\mathbf{b}_1, \ldots, \mathbf{b}_n)$ of a lattice Λ, we denote its Gram-Schmidt orthogonalization by $\mathbf{B}^* = (\mathbf{b}_1^*, \ldots, \mathbf{b}_n^*)$, where $\mathbf{b}_i^* = \pi_{(b_1,\ldots,b_{i-1})^{\perp}}(\mathbf{b}_i)$.

2.1.2. *Sublattices, quotient lattices.* Let $(\Lambda, \|\cdot\|)$ be a lattice, and let Λ' be a submodule of Λ. Then the restriction of $\|\cdot\|$ to Λ' endows Λ with a lattice structure. The pair $(\Lambda', \|\cdot\|)$ is called a *sublattice* of Λ. If any bases 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: $\|\mathbf{v} + \Lambda'\|_{\Lambda/\Lambda'} = \inf_{\mathbf{v}' \in \Lambda'_{\mathbf{R}}} \|\mathbf{v} - \mathbf{v}'\|$. Then, there is an isometry between $(\Lambda/\Lambda', \|\cdot\|_{\Lambda/\Lambda'})$ and $(\pi_{V^{\perp}}(\Lambda), \|\cdot\|)$. 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, $\pi_{V^{\perp}}(\Lambda) = \Lambda''$ 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 $\Lambda \cap \Lambda'_{\mathbf{R}}$. 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.

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. An effective solution to do so consists in using, for instance, *Babai's rounding* or *Babai's nearest plane* algorithms — a pseudo-code is given in Supplementary Material A.

2.1.4. *Filtrations.* A *filtration* of a lattice Λ is an increasing sequence of sublattices $\{0\} = \Lambda_0 \subset \Lambda_1 \subset \Lambda_2 \subset \cdots \subset \Lambda_k = \Lambda$ where each Λ_i is a primitive sublattice of Λ_{i+1}. Let $\mathrm{rk}(\Lambda_i) = d_i$, then $0 = d_0 < d_1 < d_2 < \cdots < d_k = \mathrm{rk}(\Lambda)$. 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: if $(\Lambda_i)_i$ is a filtration of Λ, then $(\Lambda_{i+j}/\Lambda_j)_i$ is a natural filtration of the lattice Λ/Λ_i.

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 shortened in s^2 and s is called the *width.*

Let now $\Lambda \subset \mathbf{R}^m$ of rank $n \leqslant m$. 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

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

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 $\rho_{\frac{1}{s^2}}(\Lambda^\vee) \leqslant 1+\varepsilon$. When the Gaussian width s exceeds the smoothing parameter, all the lattice cosets have roughly the same mass. We propose a slight generalization of [26] for general covariances encompassing this claim.

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

The following result recalls that cosets' mass has exponential decay from the origin. A useful consequence is to express the Gaussian mass by mean of a sublattice and its corresponding projection (a proof is provided in Supplementary Material C.1).

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 $\mathbf{x} \mapsto \mathbf{x}^t\Sigma^{-1}\mathbf{x}$. Then we have $\rho_\Sigma(\mathbf{x}+\Lambda) \leqslant \rho_\Sigma(P(\mathbf{x})) \cdot \rho_\Sigma(\Lambda)$. If moreover Λ is primitive in Λ', we have $\rho_\Sigma(\Lambda') \leqslant \rho_\Sigma(\Lambda)\rho_\Sigma(P(\Lambda'))$. The equality case occurs when $\Lambda' = \Lambda \perp P(\Lambda')$.

2.3. **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 [7, 8], or more recently and closely related to this work, see [14, 34]). Most of their fundamental quantities are well-understood, and general exposition can be found in [23, Chapter 4] or [9]. 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.

Definition 2.1 (Root lattices)**.** *For integer $n > 0$, the root lattices $\mathsf{A}_n, \mathsf{D}_n, \mathsf{E}_n$ of rank n are respectively defined as*

$$\mathsf{A}_n = \{\mathbf{v} \in \mathbf{Z}^{n+1} \mid v_1 + \cdots + v_{n+1} = 0\},$$

$$\mathsf{D}_n = \{\mathbf{v} \in \mathbf{Z}^{n} \mid v_1 + \cdots + v_{n} \in 2\mathbf{Z}\},$$

$$\mathsf{E}_n = \left\{\mathbf{v} \in \mathbf{Z}^{n} \cup \left(\mathbf{Z} + \frac{1}{2}\right)^n \mid v_1 + \cdots + v_{n} \in 2\mathbf{Z}\right\}.$$

We will particularly focus on A_n lattices and their dual. If $(\mathbf{e}_i)_{i\leqslant n+1}$ denotes the canonical basis of $\mathbf{R}^{n+1}$, they are generated by $(\mathbf{e}_i - \mathbf{e}_{i+1})_{1\leqslant i \leqslant n}$ and span the hyperplane $\mathbf{1}^{\perp}$, where $\mathbf{1} = (1, \ldots, 1)$. Their volume is $\sqrt{n+1}$, and $\lambda_1(\mathsf{A}_n) = \sqrt{2}$. Their dual are $\mathsf{A}_n^{\vee} = \pi_{\mathbf{1}^{\perp}}(\mathbf{Z}^{n+1})$, with $\lambda_1(\mathsf{A}_n^{\vee}) = \sqrt{n/(n+1)}$, and A_n has index $n+1$ in $\mathsf{A}_n^{\vee}$. Noticeably, A_2 identifies with the famous hexagonal lattice.

2.4. **On cyclotomic rings.** In the final Section 6, we need some background on cyclotomic rings and their geometry. Most of the material used is relatively standard, but some other aspects might be less known. For completeness purposes, we have put these recalls in Supplementary Material B as it is not our primary focus.

3. Discrete Gaussians and short exact sequences

3.1. **The short exact sequence sampler.**

3.1.1. *Split of the mass over short exact sequences.* In this section, we generically study the behavior of discrete Gaussians within *general* short exact sequences. For a given lattice Λ, we work with the following exact sequence of $\mathbf{Z}$-modules:

$$0 \longrightarrow \Lambda' \longrightarrow \Lambda \longrightarrow {}^{\Lambda}\!/_{\Lambda'} \longrightarrow 0, \tag{1}$$

meaning that the kernel of each arrow is exactly the image of the arrow preceding it. It implies in particular that the map $\Lambda' \to \Lambda$ is an injection (i.e., Λ' identifies to a submodule of Λ) and that the map $\Lambda \to \Lambda/\Lambda'$ 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 $\mathrm{T} \oplus \Lambda'_f$ of its *torsion part* 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 T. This denser lattice can be understood as a "primitivation" of Λ'. We detail an example in Supplementary Material A.

For the sake of notational simplicity, we now focus on centered discrete Gaussian distribution and omit the parameter s in the following discussion. Over this sequence, such a distribution $\mathcal{D} = \rho/\rho(\Lambda)$ over Λ decomposes into two components measures, which can then be normalized to probability distributions:

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

Understanding the latter is the focus of the next lemma. Recall that $\overline{\Lambda'}$ is primitive, so that $\Lambda/\overline{\Lambda'}$ identifies to the lattice $\pi_{(\Lambda'_{\mathbf{R}})^\perp}(\Lambda)$. The 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 $\rho_\Sigma(\mathbf{x} + \mathbf{y}) = \rho_\Sigma(\mathbf{x}) \cdot \rho_\Sigma(\mathbf{y})$. We shall however abuse notation and use π indifferently for the quotient map or for the orthogonal projection, as this will cause no harm.

Lemma 3.1. *Let $\Lambda' \subset \Lambda$ be lattices and* 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 $\pi_\star\mathcal{D}$ is at total variational distance $\frac{\varepsilon}{1-\varepsilon}$ of the distribution defined by $|\mathrm{T}|^{-1} \cdot \mathcal{D}_{\pi(\Lambda),\pi(\mathbf{c}),\Sigma}$, where $|\mathrm{T}|$ is the cardinality of* T.

For the sake of notational clarity, we also omit the parameter Σ in the proof, and restrict to the case of centered distributions (but the proof readily adapts).

Proof. Let $\mathrm{T} = \{\mathbf{t} + \Lambda'\}_{\mathbf{t}}$ be a system of representative of the torsion points, i.e., for any $\mathbf{x} \in \Lambda$, there is a unique $\mathbf{t}$ such that $\mathbf{x} \mod \Lambda' = \pi(\mathbf{x}) + \mathbf{t}$. By definition and orthogonality, the pushforward of the discrete Gaussian acts as

$$\mathcal{D}(\pi^{-1}(\mathbf{x} \bmod \Lambda')) = \rho(\mathbf{t} + \Lambda')\mathcal{D}(\pi(\mathbf{x})). \tag{2}$$

Therefore the total measure of the quotient Λ/Λ' can be written

$$\mathcal{D}\Big(\pi^{-1}\big({}^{\Lambda}\!/\!{}_{\Lambda'}\big)\Big) = \sum_{(\mathbf{t},\pi(\mathbf{x}))} \rho(\mathbf{t}+\Lambda')\mathcal{D}(\pi(\mathbf{x})) = \mathcal{D}(\pi(\Lambda)) \cdot \sum_{\mathbf{t}} \rho(\mathbf{t}+\Lambda'). \tag{3}$$

Since $\overline{\Lambda'}$ is the disjoint union of the $\mathbf{t}+\Lambda'$, so we have $\rho(\overline{\Lambda'}) = \sum_{\mathbf{t}} \rho(\mathbf{t}+\Lambda')$. By assumption on Σ and Lemma 2.1, we get $\rho(\mathbf{t}+\Lambda') \in [\frac{1-\varepsilon}{1+\varepsilon}, 1] \cdot \rho(\Lambda')$ for all $\mathbf{t}+\Lambda' \in \mathrm{T}$, which implies $\rho(\overline{\Lambda'}) \in \left[\frac{1-\varepsilon}{1+\varepsilon}, 1\right] \cdot |\mathrm{T}|\rho(\Lambda')$. Now we normalize the pushforward measure by combining the Identities (2) and (3) to obtain

$$\frac{\pi_\star\mathcal{D}(\mathbf{x} \bmod \Lambda')}{\pi_\star\mathcal{D}(\Lambda/\Lambda')} \in \left[\frac{1-\varepsilon}{1+\varepsilon}, \frac{1+\varepsilon}{1-\varepsilon}\right] \cdot \frac{1}{|\mathrm{T}|} \cdot \frac{\rho(\pi(\mathbf{x}))}{\rho(\pi(\Lambda))}. \tag{4}$$

■

Lemma 3.1 satisfyingly recovers the *extreme* cases which are frequently encountered in the literature. On the one hand, if Λ' is full-rank, we have $|\mathrm{T}| = [\Lambda : \Lambda']$ and the projection sends all points to 0, so that $\pi_\star\mathcal{D}$ is statistically close to the uniform distribution over the finite group of Λ'-cosets. On the other hand, if Λ' is primitive, the quotient is torsion-free, and we recover that $\pi_\star\mathcal{D}$ 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 $0 \to \Lambda' \to \Lambda \to \Lambda'' \to 0$, but 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 $\pi_\star\mathcal{D} = \mathcal{D}_{\Lambda'',\pi(\mathbf{t}),s^2}$.*

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. ■

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

[Modularity of smoothing parameter] Let Λ be a lattice and $0 < \varepsilon < \sqrt{17} - 4$, then

$$\eta_{3\varepsilon}(\Lambda) \leqslant \min_{\Lambda' \subset \Lambda} \max\Big(\eta_\varepsilon(\Lambda'), \eta_\varepsilon\Big({}^{\Lambda}\!/\!{}_{\Lambda'}\Big)\Big).$$

where $\eta_\varepsilon(\Lambda/\overline{\Lambda'})$ is set by convention to 0 if the quotient is of torsion, and where the minimum ranges over all possible sublattices of Λ.

The proof is detailed in Supplementary Material D.1. The intuition behind the — apparently arbitrary — choice of 0 for the torsion part stems from the following. Above the smoothing of the sublattice Λ', the distribution over the quotient is already (almost-)uniform. Hence, if we interpret the smoothing parameter to be the minimal width to *smooth out* the lattice structure (i.e., have a pushforward over $\Lambda_{\mathbf{R}}/\Lambda$ which is uniform), there is no additional condition. Indeed, scalar extension by $\mathbf{R}$ over the sequence (1) yields:

$$0 \longrightarrow \Lambda'_{\mathbf{R}} \longrightarrow \Lambda_{\mathbf{R}} \longrightarrow {}^{\Lambda}\!/\!_{\Lambda'} \otimes_{\mathbf{Z}} \mathbf{R} \longrightarrow 0.$$

Decompose now $\Lambda/\Lambda' \cong \mathrm{T} \oplus \mathbf{Z}^{r-r'}$ into its torsion and free part, where $r = \operatorname{rk}\Lambda$ resp. $r' = \operatorname{rk}\Lambda'$. As T is a finite abelian group, we have $\mathrm{T} \otimes_{\mathbf{Z}} \mathbf{R} = \{0\}$, and thus we have (as abelian groups) $\Lambda/\Lambda' \otimes_{\mathbf{Z}} \mathbf{R} \cong (\mathrm{T} \otimes_{\mathbf{Z}} \mathbf{R}) \oplus \mathbf{R}^{r-r'} \cong \mathbf{R}^{r-r'}$. In other words, the extension of scalar makes the torsion *vanish* so that it corresponds to the space spanned by $\Lambda/\overline{\Lambda'}$, the torsion-free part of Λ/Λ'. As such, the pushforward measure is driven only by Λ' and the torsion-free part of the quotient.

3.1.3. *Towards a Gaussian sampler.* The bound of Section 3.1.2 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 (or an approximation of this 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 $\overline{\mathbf{x}'} = \mathbf{t} + \mathbf{x}'$ uniquely too, since it also belongs to a unique coset $\mathbf{t} + \Lambda'$ with $\mathbf{t} \in \mathrm{T}$. Above the smoothing of Λ', 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.[5] All-in-all, this strategy leads to Algorithm 1, where we even allow sampling approximatively in the sets Λ' and Λ/Λ'—this will be proved usefull to recursively chain calls of this sampler, as we do in Section 3.3.

[5] Equivalently, it could fall into any of them.

```
Algorithm 1: 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 ${}^{\Lambda}\!/\!_{\Lambda'}$, $\frac{1+\delta}{1-\delta}$-close to the pushforward of $\mathcal{D}_{\Lambda,\star,\Sigma}$
Output: $\mathbf{v} \in \Lambda$ following distribution statistically close to $\mathcal{D}_{\Lambda,\mathbf{t},\Sigma}$
if $\Lambda = \{0\}$ then return $0$
Compute $\pi : \Lambda_{\mathbf{R}} \to {}^{\Lambda_{\mathbf{R}}}\!/\!_{\Lambda'_{\mathbf{R}}}$, the orthogonal projection onto $\Lambda'^{\perp}_{\mathbf{R}}$ for the norm
    induced by $\Sigma^{-1}$
$\mathbf{q} \leftarrow \mathcal{O}_q\left({}^{\Lambda}\!/\!_{\Lambda'}, \pi(\mathbf{t}), \Sigma\right)$; $\mathbf{u}_q \leftarrow$ Lift$(\mathbf{q}, \Lambda)$
$\mathbf{u}' \leftarrow \mathcal{O}'(\Lambda', (Id - \pi)(\mathbf{t} - \mathbf{u}_q), \Sigma)$
return $\mathbf{u}_q + \mathbf{u}'$
```

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

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

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

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 switch cosets "at the cost of some ε". The details are given in Supplementary Material D.1.

We now present two applications of this abstract sampler: domain extensions and restrictions, and a broad generalization of the so-called Klein's sampler [19].

3.2. **Application I: full rank domain extension and restriction.**

3.2.1. *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 $\mathcal{D}_{\Lambda',\star,\Sigma}$ for a parameter $\Sigma \succ \eta_\varepsilon(\Lambda')$. By Lemma 3.1, the pushforward $\pi_\star(\mathcal{D}_{\Lambda,\star,\Sigma})$ is at distance at most $\frac{\varepsilon}{1-\varepsilon}$ of the uniform distribution over Λ/Λ'. Hence specializing Algorithm 1 with $\mathcal{O}'$ and a uniform sampler for $\mathcal{O}_q$ yields the following:

Corollary 3.1 (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 with covariance $\Sigma \succ \eta_\varepsilon(\Lambda')$, there exists an algorithm sampling at distance at most 6ε of $\mathcal{D}_{\Lambda,\star,\Sigma}$ using at most one oracle call to $\mathcal{O}'$.*

In a nutshell, the ability to sample in Λ' and from a pushfoward distribution over Λ/Λ' close to uniform enable to reconstruct samples in Λ: we do a *domain extension* of the discrete Gaussian over Λ' to the overlattice Λ. We point out the possible connection with the averaging recombination technique used in [1], where a domain extension from 2Λ to Λ is perfomed (using exponentially many vectors).

3.2.2. *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 5, we showcase some practical examples with root lattices.

Proposition 3.1 (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 repetition 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. ∎

3.3. **Application II: the filtered sampler.** We now show that our short exact sequence sampler naturally extends to filtrations and allows to retrieve and generalize samplers appearing in cryptography, such as those in [15, 19, 32]. For example, in the most natural case where one would sample "coordinate-by-coordinate", our algorithm recovers Klein's 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 algorithms [6].

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

3.3.1. *Smoothing parameter bound over a filtration.* We first highlight a new smoothing parameter bound deduced from a given filtration

$$\{0\} = \Lambda_0 \subset \Lambda_1 \subset \cdots \subset \Lambda_k = \Lambda$$

of a lattice Λ. It relies on repeatedly applying the splits of the smoothing parameter over a short sequence (Section 3.1.2) along this 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 Section 3.1.2 to Λ_{k-1}, we also have $\eta(\Lambda_{k-1}) \leqslant \max(\eta(\Lambda_{k-1}/\Lambda_{k-2}), \eta(\Lambda_{k-2}))$. We go down the filtration inductively until we reach Λ_1. All in all, the smoothing parameter is dominated by the biggest term appearing in the splitting. Keeping track of the growth of the ε and optimizing gives:

Let $k \geqslant 1$ be an integer, Λ a lattice and $\varepsilon \leqslant e^{-k/2}/\sqrt{2}$. We have

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

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

The proof is provided in Supplementary Material D.2. Note that the term $k+1$ is quite arbitrary and we can choose any real $k' \geqslant k$ instead, as long as the condition on ε is updated conjointly. The idea behind the bound above allows to mildly relax the smoothness condition over lattice cosets: instead of the whole lattice, it is only needed to smooth the "worst" successive quotient deduced from the filtration for the cosets of the whole lattice to have essentially the same mass.

[6] We will explicitly exploit these additional degrees of liberties Section 5.

For example, it was shown[7] in [19], and subsequently used at the core of several practical constructions, that for any rank n lattice Λ,

$$\eta_\varepsilon(\Lambda) \leqslant \min_{\substack{(\mathbf{b}_1,\ldots,\mathbf{b}_n) \\ \text{basis of } \Lambda}} \max_{1 \leqslant i \leqslant n} \eta_{\frac{\varepsilon}{n}}(\mathbf{Z}\widetilde{\mathbf{b}_i}).$$

This bound corresponds to restricting Section 3.3.1 to complete filtrations of length n, i.e. the filtration stemming from the $\mathbf{b}_i$'s as $\Lambda_i = \mathcal{L}(\mathbf{b}_1, \ldots, \mathbf{b}_i)$. Indeed, we have that for any $0 \leqslant i < n$, Λ_{i+1}/Λ_i is isometric to $\mathbf{Z}\mathbf{b}_i^*$, with $\mathbf{b}_i^*$ being the corresponding Gram-Schmidt vector of the basis, (see also Section 2.1.2). While it could seems 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.

3.3.2. *The filtered sampler.* Following our motto — smoothing bounds and sampling are built on the same underlying principles — we can transform Section 3.3.1 into a Gaussian sampler. In essence, the process corresponds to k successive calls of Algorithm 1, 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 (e.g. of Algorithm 5) The first call considers the short exact sequence

$$0 \to \Lambda_1 \to \Lambda \to {}^{\Lambda}\!/_{\Lambda_1} \to 0.$$

Algorithm 1 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 $\{0\} = {}^{\Lambda_1}\!/_{\Lambda_1} \subset {}^{\Lambda_2}\!/_{\Lambda_1} \subset \cdots \subset {}^{\Lambda}\!/_{\Lambda_1}$. Hence the callee now deals with the sequence $0 \to \Lambda_2/\Lambda_1 \to \Lambda/\Lambda_1 \to \Lambda/\Lambda_2 \to 0$. 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.

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

Algorithm 2: Filtered sampler

```
Input: A filtration $\{0\} \subset \Lambda_1 \subset \cdots \subset \Lambda_k = \Lambda$, a parameter
       $\Sigma > \max_{0 \leqslant i < k} \eta_\varepsilon\left({}^{\Lambda_{i+1}}\!/_{\Lambda_i}\right)$ 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 $\pi : \Lambda \to {}^{\Lambda}\!/_{\Lambda_1}$
3 $\mathbf{z} \leftarrow$ FilteredSampler$\left(\left({}^{\Lambda_i}\!/_{\Lambda_1}\right)_i, \pi(\mathbf{t}), \Sigma\right)$
4 $\mathbf{u} \leftarrow$ Lift$(\mathbf{z}, V_1)$
5 $\mathbf{u}' \leftarrow \mathcal{D}_{\Lambda_1,(\mathrm{Id}-\pi)(\mathbf{t}-\mathbf{u}),\Sigma}$
6 return $\mathbf{u} + \mathbf{u}'$
```

[Correctness of the filtered sampler] Algorithm 2 is correct. Moreover, let $\mathcal{D}$ be the distribution of its output. For small enough ε, we have

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

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

Proof. It suffices to proceed by induction along the filtration repeatedly calling Algorithm 2. The detailed proof can be found in Supplementary Material D.2. ■

3.4. **Recovering some known samplers.** The filtered sampler readily recovers some well-known samplers.

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

3.4.2. *Klein's 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 $(\mathbf{m}_1, \ldots, \mathbf{m}_d)$

over $\mathcal{O}_{\mathbf{K}}$, we make use of the full filtration[8] $\Lambda_1 = \mathbf{m}_1\mathcal{O}_{\mathbf{K}} \subset \Lambda_2 = \mathbf{m}_1\mathcal{O}_{\mathbf{K}} \oplus \mathbf{m}_2\mathcal{O}_{\mathbf{K}} \subset \cdots$. Each recursive call thus consists in calling the oracles over the quotients Λ_{i+1}/Λ_i. When instantiating this oracle with subsequently described Algorithm 3 (or Peikert's [29] for instance), it retrieves the so-called *hybrid sampler* used in [15].

3.4.3. *Fast Fourier Orthogonalization sampler.* Introduced in [13, 33], 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 $\mathbf{m}_1, \mathbf{m}_2$ 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:

$$0 \longrightarrow \mathbf{m}_1\mathcal{O}_{\mathbf{K}} \longrightarrow \Lambda \longrightarrow {}^{\Lambda}\!\big/_{\mathbf{m}_1\mathcal{O}_{\mathbf{K}}} \longrightarrow 0$$

where once again the submodule $\mathbf{m}_1\mathcal{O}_{\mathbf{K}}$ 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 $\mathbf{m}_1\mathcal{O}_{\mathbf{K}}$ and $\Lambda/\mathbf{m}_1\mathcal{O}_{\mathbf{K}}$. 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 reapply 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. The linear sampler

4.1. **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 of sampling 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

[8] We 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 Supplementary Material B.1.

be corrected. For example, the filtered sampler of Section 3.3.2 *iteratively* corrects the transformation to the space attached to the filtration, acting "subspace-by-subspace". 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, which can be useful when few quantitative informations about Λ are known.

Lemma 4.1. *Let Λ be a lattice of rank n in $\mathbf{R}^m$, then $\eta_\varepsilon(\Lambda) \leqslant \inf s_1(\mathbf{T}) \cdot \eta_\varepsilon(\mathcal{L}(\mathbf{C}))$, where the infimum is taken over all pair $(\mathbf{T}, \mathbf{C})$ such that $\Lambda = \mathcal{L}(\mathbf{TC})$ 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 $\Lambda^\vee$ 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 $s \geqslant \eta_\varepsilon(\Lambda)$ 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. ■

4.2. **Sampling by linear transformation.** While the previous section dealt with local corrections for Gaussian sampling, we are now interested in a *global* approach to Gaussian sampling. Following as always our motto that a bound on the smoothing parameter corresponds to a sampling algorithm, we consider a lattice $\mathcal{L}(\mathbf{B})$ as a transformation $\mathbf{T}$ of some initial lattice $\mathcal{L}(\mathbf{C})$. Our approach follows and generalizes the proposition of Peikert [29], where $\mathcal{L}(\mathbf{C}) = \mathbf{Z}^n$, and we now give an informal description.

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: this rescaling corresponds to the fact that "there must be enough available space in our target ellipsoid" so that we can "inflate" the starting one into it by adding perturbations. 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.

Algorithm 3: Linear sampler

```
Input:
    • Two matrices $\mathbf{T} \in \mathbf{R}^{m\times n}, \mathbf{C} \in \mathbf{R}^{n\times n}$ with $m \geqslant n$ and $\mathbf{C}$ invertible such that
      $\mathbf{TC} = \mathbf{B}$ is a basis of a lattice $\Lambda$;
    • a center $\mathbf{t} \in \Lambda \otimes \mathbf{R}$;
    • a parameter $r \geqslant 0$ 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 $\mathcal{D}_{\Lambda,\mathbf{t},\Delta}$.
$\Sigma_{\mathbf{p}} \leftarrow \Sigma - r^2\mathbf{I}$
$\mathbf{p} \leftarrow \mathcal{N}_{\Sigma_{\mathbf{p}}}$
$\mathbf{x} \leftarrow \mathcal{D}_{\mathcal{L}(\mathbf{C}),\mathbf{T}^\star\mathbf{t}+\mathbf{p},r^2}$ /* $\mathbf{T}^\star$ is the pseudo-inverse */
; return $\mathbf{y} := \mathbf{Tx}$
```

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

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

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

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.*

The proof is very similar in spirit of [15, 29], and amounts to a marginal distribution calculation combined with the nice properties of Gaussian functions with respect to multiplication (see e.g. [29, Fact 2.1]). It is therefore in Supplementary Material E. The main technical difference in the proof is that the considered lattices may not be full-rank in their ambient space. This is dealt with properties of the pseudo-inverse: $\mathbf{T}^\star\mathbf{T} = \mathbf{I}_n$ and $\mathbf{TT}^\star$ is the orthogonal projection onto $\Lambda\otimes\mathbf{R}$. With the notation of Algorithm 3, this allows to write $\rho_\Sigma(\mathbf{x}-\mathbf{T}^\star\mathbf{t}) = \rho_\Delta(\mathbf{y}-\mathbf{t})$ and to proceed. We note that if $m = n$, then $\mathbf{T}$ is invertible and the algorithm is correct as soon as $\Delta \succ r^2\mathbf{TT}^t$. Second, if a decomposition

$\Delta = \mathbf{L}^t\mathbf{L}$ is known, then a sufficient condition for correctness is $s_n(\mathbf{L}) > r \cdot s_1(\mathbf{T})$: this recovers the results of [29, 15].

4.3. **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 [29].

- **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 $\mathbf{B} = \widetilde{\mathbf{B}}\mathbf{U}$ 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 Section 3.3.1, we thus have $\eta_\varepsilon(\mathcal{L}(\mathbf{U})) \leqslant \eta_{\varepsilon/(n+1)}(\mathbf{Z})$. Using the standard bound $\eta_\varepsilon(\mathbf{Z}) \leqslant \sqrt{\ln(2(1+1/\varepsilon)/\pi}$, the sampler with that decomposition reaches standard deviations very close to $s_1(\widetilde{\mathbf{B}}) \cdot \eta_\varepsilon(\mathbf{Z}^n)$.
- **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, 32], when decomposing the NTRU lattice.

4.4. **Application: sampling in tensor lattices.** A lattice $\mathcal{L}(\mathbf{A}) \otimes \mathcal{L}(\mathbf{B})$ is generated by the matrix $\mathbf{A} \otimes \mathbf{B}$ which can be rewritten in as a matrix product involving $\mathbf{A}$ and $\mathbf{B}$. Therefore Algorithm 3 instantiates very well over such lattices. This gives (up to our knowledge) a novel way to sample in $\mathcal{L}(\mathbf{A}) \otimes \mathcal{L}(\mathbf{B})$, and a corresponding smoothing parameter bound. The details are moved to Supplementary Material F because of space constraints.

5. Sampling in root lattices

This section collects various approach to efficiently sample Gaussian in root lattices A_n, D_n and E_8. On the one hand, some of them will appear to be important building blocks for sampling cyclotomic integers 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 build around one dimensional samplers). On the other hand, they are also a good way to illustrate practical usecases of our generic samplers from the previous sections.

Our ad-hoc samplers in particular rely on exceptional orthogonal decompositions involving such lattices, and their close relationship in general; all the background material for this section is described in [23, Chapter 4].

5.1. **Sampling in low dimensional root lattices.** We start with samplers for the root lattices of small dimension, as well as the D_n family. They are based on the ability to juggle between restrictions and extensions of lattices using Proposition 3.1 and Corollary 3.1, and exceptional isometries between them [23, Chap. 4.6].

Theorem 5.1. *We can sample efficiently and at standard deviation right above the smoothing parameter in the following exceptional lattices:*

- D_n, $\mathsf{A}_n^\vee$ *for all* $n > 1$*;*
- $\mathsf{A}_2, \mathsf{A}_3, \mathsf{A}_4, \mathsf{A}_5, \mathsf{A}_6, \mathsf{A}_8$*;*
- $\mathsf{E}_6, \mathsf{E}_7, \mathsf{E}_8$.

Below, we only give the high-level ideas of the samplers used in Section 6. All remaining proofs and details[9] can be found in Supplementary Material G.

D_n **samplers:** The D_n lattice can be described as the vectors of $\mathbf{Z}^n$ which coordinates in the canonical basis $(\mathbf{e}_i)_{i \leqslant n}$ sum to an even number, so that $[\mathbf{Z}^n : \mathsf{D}_n] = 2$. This congenial

[9] One can also sample in $\mathsf{A}_n = \mathbf{Z}^{n+1} \cap \mathbf{1}^\perp$, checking when the sum of coordinates vanishes. This is clearly inefficient when n grows. In the next section we propose a far more efficient algorithm, when $n \geqslant 9$.

definition leads to a domain extension approach form as in Proposition 3.1: a sample either belongs to D_n or either to its non-zero coset (with almost equiprobability above $\eta_\varepsilon(\mathsf{D}_n)$).

$\mathsf{A}_n^\vee$ **samplers:** By definition, $\mathsf{A}_n^\vee$ is $\pi_{\mathbf{1}^\perp}(\mathbf{Z}^{n+1})$, where $\mathbf{1} = (1, \ldots, 1)$, and thus directly falls into the framework of Section 3: sample in $\mathbf{Z}^{n+1}$ and compute the projection of the sample. For a vector $\mathbf{z} = (z_1, \ldots, z_{n+1})$, we have $\pi_{\mathbf{1}^\perp}(\mathbf{z}) = (z_i - \frac{1}{n+1}\sum_j z_j)_{i \leqslant n+1}$. Note that this does not incur any smoothing condition for correctness (although the standard deviation parameter should be large enough to make the statistical property of the sample "Gaussian-like").

A_2 **sampler:** From above, we can sample easily in $\mathsf{A}_2^\vee$. The additional ingredient is that any rank 2 lattice is similar to its dual (see e.g. [23]). Here, the similarity given by $\mathsf{A}_2 = \sqrt{3}R_{-\pi/2}(\mathsf{A}_2^\vee)$, where $R_{-\pi/2}$ is the rotation by $-\pi/2$ of the hyperplane $\mathbf{1}^\perp$ around the origin. These observations lead to Algorithm 9 (detailed in Supplementary Material G).

E_8 **sampler:** the E_8 lattice is an unimodular lattice in $\mathbf{R}^8$ so in particular its determinant is 1. It is not, however, isometric to $\mathbf{Z}^8$. We have the exact sequence $0 \to \mathsf{D}_8 \to \mathsf{E}_8 \to \mathbf{Z}/2\mathbf{Z} \to 0$ by the covering of cosets $\mathsf{E}_8 = \mathsf{D}_8 \cup (1/2, \ldots, 1/2) + \mathsf{D}_8$. Combining our D_n sampler with Algorithm 1 gives Algorithm 10 to sample in E_8: flip a coin to decide the coset, sample a Gaussian in D_8, output the sum.

A_8 **sampler:** from [23, Theorem 4.6.7 and 4.6.12], A_8 is isometric to a lattice of index 3 in E_8. Combining the E_8 sampler and Proposition 3.1 in the natural way gives Algorithm 11.

We claim that all our approaches reach very close to the smoothing parameter of these lattices, and we now give estimates for the one we need in later sections. The main ingredient here is the identification of the Gaussian mass with the *theta series* of a lattice:

$$\begin{aligned}\rho_{1/s^2}(\Lambda^\vee) &= 1 + \kappa(\Lambda^\vee) \cdot \exp(-\pi s^2)^{\lambda_1(\Lambda^\vee)^2} + \kappa_2 \exp(-\pi s^2)^{n_2^2} + \cdots \\ &= \theta_{\Lambda^\vee}(\exp(-\pi s^2)),\end{aligned} \tag{5}$$

where we have sorted the vectors of $\Lambda^\vee$ by their increasing squared norm, and $\kappa(\Lambda^\vee)$ is the *kissing number* of $\Lambda^\vee$. 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

$$s = \sqrt{\frac{\ln(-S(\varepsilon))}{\pi}}.$$

Note that this is an *exact* expression, but that working out some terms of S requires to know those of $\theta_{\Lambda^\vee}$. Thankfully, for all exceptional lattices, the first terms of the theta series are well-known, and we obtain the next lemma (details are given in Supplementary Material G).

We have the following estimates for $\varepsilon > 0$:

- $\eta_\varepsilon(\mathbf{Z}^n) = \sqrt{\frac{1}{\pi}(\ln(\frac{2n}{\varepsilon}) + o(1/\varepsilon))}$;
- for $n \geqslant 5$, $\eta_\varepsilon(\mathbf{Z}^n) \leqslant \eta_\varepsilon(\mathsf{D}_n) = \sqrt{\frac{1}{\pi}(\ln(\frac{2n}{\varepsilon}) + o(1/\varepsilon))} \approx \eta_\varepsilon(\mathbf{Z}^n)$;
- $\eta_\varepsilon(\mathsf{A}_n) = \sqrt{\frac{n+1}{n}} \cdot \sqrt{\frac{1}{\pi}(\ln(\frac{2(n+1)}{\varepsilon}) + o(1/\varepsilon))} \approx \lambda_1(\mathsf{A}_n^\vee)^{-1} \cdot \eta_\varepsilon(\mathbf{Z}^n)$;
- $\eta_\varepsilon(\mathsf{A}_n^\vee) = \frac{1}{\sqrt{2}} \cdot \sqrt{\frac{1}{\pi}(\ln(\frac{n(n+1)}{\varepsilon}) + o(1/\varepsilon))} \approx \lambda_1(\mathsf{A}_n)^{-1} \cdot \eta_{2\varepsilon/(n+1)}(\mathbf{Z}^n)$

The second result can be understood intuitively as follows: $\mathsf{D}_n^\vee$ 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 are the same. On the other hand, $\lambda_1^\infty(\mathsf{D}_n^\vee) = \frac{1}{2}$, so that the usual bound[10] obtained from the shortest vector of the dual *in the ℓ_∞ norm* would give an overestimate by a factor ≈ 2.

5.2. **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 instanciating our framework of Section 3.3.2 and Section 4.2 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.

[10] From e.g. [28, Lemma 3.5], we have $\eta_\varepsilon(\Lambda) \leqslant \lambda_1^\infty(\Lambda^\vee)^{-1} \cdot \eta_\varepsilon(\mathbf{Z}^n)$ 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 [28, Lemma 3.5], *unconditionnally* on ε contrary to [30, Lemma 2.6], but involving the kissing number of the dual.

5.2.1. *Trivial instantiations: Peikert and Klein samplers.* The Gram matrix of the standard basis $(\mathbf{e}_i - \mathbf{e}_{i+1})_{1\leqslant i\leqslant n}$ of A_n is

$$G_n = \begin{pmatrix} 2 & -1 & 0 & 0 & \cdots & 0 & 0 \\ -1 & 2 & -1 & 0 & \cdots & 0 & 0 \\ \vdots & \vdots & \ddots & \ddots & \ddots & \vdots & \vdots \\ 0 & 0 & 0 & \cdots & -1 & 2 & -1 \\ 0 & 0 & 0 & \cdots & 0 & -1 & 2 \end{pmatrix} \in \mathbf{Z}^{n\times n}. \tag{6}$$

Unrolling the Cholesky algorithm on this matrix 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 to perform Gaussian sampling at standard deviation above $\sqrt{2}\eta_\varepsilon(\mathbf{Z})$. As G_n is a tridiagonal Toeplitz matrix, its eigenvalues are of the form $2+2\cos(k\pi/(n+1))$ for $1 \leqslant k \leqslant n$ [20]. Consequently, the largest singular value of this basis is $(2+2\cos(\pi/(n+1)))^{1/2} \geqslant 2\sqrt{1-\frac{\pi^2}{2n^2}}$, a worse reachable standard deviation. The other classic basis $(\mathbf{e}_1 - \mathbf{e}_i)_{2\leqslant i\leqslant n+1}$ of A_n has a largest singular value of $\sqrt{n+1}$, which has an even worse geometry.

5.2.2. *Constructing a better filtration.* To showcase possible trade-offs using Algorithm 2, 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 11)—as subroutines. This new family of algorithms allows to sample 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 be 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 reuseability.

Proposition 5.1. *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 $0 = \Lambda_0 \subset \Lambda_1 \subset \Lambda_2 \subset \cdots \subset \Lambda_q \subset \mathsf{A}_n$, where for all $1 \leqslant i \leqslant q$, Λ_i is isometric to an orthogonal direct sum of i copies of A_k.*

Sketch. 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. All details are presented in Supplementary Material G.9. ■

The remaining vectors have to be dealt with, but it turns out not impacting what follows. This allows to sample over the A_n lattice using Algorithm 2. Let us call $\mathbf{B}_n$ the basis corresponding to the filtration of Proposition 5.1. At the deepest level of recursion, we sample in the lattice Λ/Λ_q, using for example Klein sampler, or equivalently, Algorithm 2 with the filtration corresponding to the projection of the last $q+r$ columns of $\mathbf{B}_n$ orthogonally to $V_q^\perp = \mathrm{Span}(\Lambda_q)^\perp$. Then, all subsequent samplings happen in (a copy of) A_k, and for example, when $k = 8$, one calls Algorithm 11 for these last q steps. For the sake of clarity, we restrict ourselves to $k = 8$ and give an equivalent *iterative* algorithm. The result is proved in Supplementary Material G.

Algorithm 4: A_n sampler

```
Input: $\sigma \geqslant \max\left\{\sqrt{\frac{9}{8}}\eta_\varepsilon(\mathbf{Z}^8), \eta_\varepsilon(\mathsf{A}_8)\right\}$, a center $\mathbf{t} \in \mathrm{Span}_{\mathbf{R}}(\mathsf{A}_n)$, a filtration
       $(\Lambda_i)_i$ of $\mathsf{A}_n$ in the form of $\mathbf{B}_n$, as in Proposition 5.1.
Output: $\mathbf{v}$ following distribution statistically close to $\mathcal{D}_{\mathsf{A}_n,\mathbf{t},\sigma^2}$
Compute $\mathbf{c}_i = \pi_{V_q^\perp}(\mathbf{b}_{i+kq})$ for $1 \leqslant i \leqslant q+r$
$\mathbf{t}_{q+1} \leftarrow \pi_{V_q^\perp}(\mathbf{t})$
$\mathbf{x}_{q+1} \leftarrow$ Algorithm 2$(\{\mathbf{c}_1, \dots, \mathbf{c}_{q+r}\}, \sigma, \mathbf{t}_{q+1})$
$\mathbf{u} \leftarrow$ Lift$(\mathbf{x}_{q+1}, V_q)$
$\mathbf{t}' \leftarrow \mathbf{t} - \mathbf{x}_{q+1}$
Compute the orthogonal projections $\mathbf{t}'_j$ of $\mathbf{t}'$ on Span$(\mathbf{b}_{jk+1}, \dots \mathbf{b}_{jk})$ for
  $0 \leqslant j < q$
$\mathbf{x}_1, \dots, \mathbf{x}_q \leftarrow$ Algorithm 11$(\sigma, \mathbf{t}'_1), \dots,$ Algorithm 11$(\sigma, \mathbf{t}'_q)$   /* can be done
  in parallel */
return $\mathbf{x}_1 + \dots + \mathbf{x}_q + \mathbf{x}_{q+1}$
```

Theorem 5.2. *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 $\sigma \geqslant \max\{\sqrt{9/8} \cdot \eta_\varepsilon(\mathbf{Z}^8), \eta_\varepsilon(\mathsf{A}_8)\}$. Then for small enough ε, the statistical distance between $\mathcal{D}$ and $\mathcal{D}_{\mathsf{A}_n,\mathbf{t},\sigma^2}$ is at most $(q+1)\varepsilon$.*

6. Practical applications to lattice-based hash-and-sign signatures

FALCON and its variant MITAKA [15] are two efficient instantiations of the GPV framework [19]. The former has recently been selected for post-quantum standardization, and the latter enjoys better versatility as 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 ring 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 be find further use in designs (for instance in [5, 24, 27]), providing more flexible parameters, more efficient samplers, and tighter security.

Rationale of Hash-and-sign over lattices: In the GPV framework[19], signing amounts to sampling a discrete Gaussian in some lattice close to an arbitrary target (corresponding to a hash of the message). In efficient instantiations, the underlying lattice is a so-called NTRU lattice over a cyclotomic ring $\mathcal{R}_m$, and corresponds to $\mathcal{M}_{\mathrm{ntru}} = (f,g)^t\mathcal{R}_m \oplus (F,G)^t\mathcal{R}_m$ for some basis $\mathbf{b}_0 = (f,g)^t, \mathbf{b}_1 = (F,G)^t \in \mathcal{R}_m^2$ of the module. This basis is carefully selected [11, 15] to enable short Gaussians sampling in $\mathcal{M}_{\mathrm{ntru}}$, and acts as the secret key. As being an underlying Approx-CVP instance, the security against forgery is driven by the standard deviation σ_{sig} that one can achieve with a given sampler: one wants to minimize it. In this section, we focus on the *hybrid sampler* [13, 32] that MITAKA uses to sign.

6.1. **Hybrid sampling and representation of cyclotomic numbers.** As seen in Section 3.4, this sampler leverages the filtration $\{0\} \subset \psi(\mathbf{b}_0\mathcal{R}_m) \subset \psi(\mathcal{M}_{\mathrm{ntru}})$, where ψ denotes the canonical embedding extended to vectors. The calls to Algorithm 3 consider $\mathbf{b}_0\mathcal{R}_m$ and $\mathcal{M}_{\mathrm{ntru}}/\mathbf{b}_0\mathcal{R}_m$ as linear transformations[11] of $\mathcal{R}_m$. Under this identification,

[11] In practice, this second call is encoded by the orthogonalization $\mathbf{b}_1^*$ of $\mathbf{b}_1$ *in the cyclotomic field*; such details are not our focus here, we let the interested reader refer to Supplementary Material B for a complete presentation.

Section 3.3.1 and Lemma 4.1 show that the sampler of MITAKA reaches standard deviation as

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

where $\alpha > 1$ encodes how close we are able to sample from the smooting parameter of the base ring $\mathcal{R}_m$. For Algorithm 2 to reach the stated covariance, it requires two *elliptic* samples[12] in $\mathcal{R}_m$. In [15], this is handled by Peikert's sampler in $\psi(\mathcal{R}_m)$, or equivalently, Algorithm 3 with $\mathbf{C}$ being the (canonical embedding of) the power basis. This choice is made partly because of square roots computation *in the field* and the use of a continuous perturbation: both steps can ben handled in quasi-linear time in the canonical embedding (equivalently, in Fourier domain). In particular, the covariance of the perturbation is represented by a diagonal matrix — this avoids costly Cholesky decompositions.

The next requirement of Algorithm 3 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 $\mathcal{M}_{ntru}$. Another approach is to sample directly in 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 $\mathcal{M}_{\mathrm{ntru}}$ 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. First, the codifferent ideal in a prime cyclotomic is the principal ideal $\mathcal{R}_p^\vee = \langle \frac{1-\zeta_p}{p} \rangle$. Using Algorithm 2 over the filtration induced by the so-called *decoding basis* [21] $\zeta_p^i - \zeta_p^{i+1}$ 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 can leverage this observation by relying on the next result.

Proposition 6.1 (Adapted from [34, 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 linear maps $\tau : \mathcal{R}_p \longrightarrow \mathsf{A}_{p-1}^\vee$*

[12] They are scalar in the (completion of) the ring, but not when acting as transformation of $\mathbf{R}^d$.

and $\phi : \langle 1-\zeta_p\rangle \longrightarrow \mathsf{A}_{p-1}$ such that, for all $x \in \mathcal{R}_p$ and $y \in \langle 1-\zeta_p\rangle$, we have

$$\|\psi(x)\|^2 = p\|\tau(x)\|^2 \quad \textit{and} \quad \|\psi(y)\|^2 = p\|\phi(y)\|^2.$$

The second map is not described in [34], see Supplementary Material B for details. Recall that ψ can be computed using the Vandermonde matrix V_p associated to 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 $\frac{1}{p}\|x\|^2 \leqslant \|\tau(x)\|^2 \leqslant \|x\|^2$ for all $x = \sum_i x_i\zeta_p^i \in \mathcal{R}_p$. Identical inequalities are obtained with ϕ and elements in $\langle 1-\zeta_p\rangle$. Thanks to our samplers in root lattices, we could use either of these maps to get Gaussian cyclotomic integers.

Which map to use? The question is perhaps more subtle than it seems, and may ultimately boil down to how the resulting scheme is implemented. Providing a complete implementation is of course out of the scope of our work. Instead, we focus only on the expected size of signature vectors achievable with these representations, and discuss their impact on security. Recall that the resistance against forgery is mainly driven by the ratio of σ_{sig} by the volume of the representation of $\mathcal{M}_{ntru}$ used. When looking at the coefficients, this means we want to take the minimum of

$$E^\vee = \frac{\eta_\varepsilon(\mathsf{A}_{p-1}^\vee)}{\det(\mathsf{A}_{p-1}^\vee)^{1/p-1}} \quad \text{and} \quad E = \frac{\eta_\varepsilon(\mathsf{A}_{p-1})}{\det(\mathsf{A}_{p-1})^{1/p-1}}.$$

The normalizing factor also encodes that the volume of $\phi(\mathcal{M}_{ntru})$ or $\tau(\mathcal{M}_{ntru})$ is changed. When $p=3$, we know that $\mathsf{A}_2^\vee$ and A_2 are similar lattices. Using one map or another can be expected to be equivalent, and we will see that it covers the smooth conductor $m = 2^\ell 3^k$ case. When p grows, Lemma 5.1 shows that the smoothing of $\mathsf{A}_{p-1}^\vee$ and A_{p-1} behave differently, since $\lambda_1(\mathsf{A}_{p-1})$ stays constant while $\lambda_1(\mathsf{A}_{p-1}^\vee)$ increases to 1. A bit less informally, using Lemma 5.1 again and $\det(\mathsf{A}_{p-1}) = \sqrt{p}$, we make the approximation

$$\frac{E^\vee}{E} \approx p^{\frac{1}{p-1}} \cdot \sqrt{\frac{p}{2(p-1)}} \cdot \left(\frac{\ln(p^2/\varepsilon)}{\ln(p/\varepsilon)}\right)^{1/2}.$$

The ratio of logarithms decreases to 1 when ε goes to 0, and the remaining factors are very close to $\sqrt{1/2} < 1$. This suggests that considering the map $\tau : \mathcal{R}_p \to \mathsf{A}_{p-1}^\vee$ is a better choice in the vacuum (once again, depending on the implementation solutions chose, this might not be the case on specific architectures). For completeness, we describe both possibilities.

6.2. **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 focus [13] on such conductors. From e.g. [21, 34] or Supplementary Material B, 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* [21])

$$\mathcal{R}_m \cong \mathbf{Z}^{\frac{\ell}{2}} \otimes \left(\bigoplus_{i=1}^{3^{k-1}} \mathcal{R}_3 \right) \cong \bigoplus_{i=1}^{\frac{m}{6}} \mathcal{R}_3.$$

By orthogonality, sampling in $\mathcal{R}_3$ amounts to $m/6$ independent sampling in $\mathsf{A}_2^\vee$, which are done as explained in Section 5 by projecting samples from $\mathbf{Z}^3$ (see also Algorithm 8 in Supplementary Material G). Alternatively, we have an orthogonal decomposition $\langle 1 - \zeta_m^{m/3} \rangle = \frac{m}{2} \mathcal{R}_m^\vee \cong \bigoplus_{i=1}^{m/6} \langle 1 - \zeta_3 \rangle$ (see also [21, Cor. 2.18]). We could use instead Algorithm 9 with this decomposition.

6.2.1. *Efficiency and signature quality.* Being a combination of Algorithm 2 with either Algorithm 8 or Algorithm 9, but because we want to control the statistical properties of the output, we can reasonably go as low as

$$\sigma' = \eta_{6\varepsilon/m}(\mathsf{A}_2^\vee) \approx \sqrt{\frac{1}{2}} \cdot \eta_{2\varepsilon/3}(\mathbf{Z}^{\frac{m}{3}}) \quad \text{or} \quad \sigma' = \eta_{6\varepsilon/m}(\mathsf{A}_2) \approx \sqrt{\frac{3}{2}} \cdot \eta_{2\varepsilon/3}(\mathbf{Z}^{\frac{m}{3}}), \tag{7}$$

where approximations come from Section 5.1. We emphasize again here that the volume of the involved lattice is changed: after normalization, even though the values are different, their impact on the security is the same.

The running time is linear in the conductor, and the sampler requires $m/2$ samples from $\mathcal{D}_{\mathbf{Z},*,\sigma^2}$. The approach is completely parallelizable, and also memory-efficient: we only needs to store a table for integer Gaussians of a small width. Moreover, thanks to the small Gaussian parameter, the constant-time implementation is easy and efficient.

6.2.2. *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 $\sigma' \geqslant \sqrt{2} \cdot \eta_\varepsilon(\mathbf{Z}^{m/3})$. This ≈ 2 factor in our favour translates quantitatively into a NIST

[13] FALCON 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 its high technicality. Such rings are also the focus of the implementation in [22].

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

security[14] 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 the sample through table look-up. However, the size of the table for $\mathcal{D}_{\mathsf{A}_2,\sigma'^2}$ 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.

6.3. **Sampling over prime cyclotomic fields.** Proposition 6.1 gives two immediate approaches. With the map τ, we can directly project samples of $\mathbf{Z}^p$ onto samples of $\mathsf{A}_{p-1}^\vee$. For the map ϕ, the situation is less nice, but we can nevertheless use our efficient Algorithm 4.

6.3.1. *Efficiency and signature quality.* Both approaches are linear in p, with Algorithm 4 main cost coming from the sampling in A_8 (Algorithm 11). Similarly as the smooth case, we want to control the statistical properties of the output. Thus the approach using Algorithm 8 and the approach of Algorithm 4 can reach respectively

$$\sigma' = \eta_\varepsilon(\mathsf{A}_p^\vee) \approx \sqrt{\frac{1}{2}} \cdot \eta_{2\varepsilon/p}(\mathbf{Z}^{p-1}) \quad \text{or} \quad \sigma' = \eta_{\varepsilon/q}(\mathsf{A}_8) \approx \sqrt{\frac{9}{8}} \cdot \eta_{2\varepsilon/9}(\mathbf{Z}^{8q}), \tag{8}$$

where approximations come from Section 5.1 and we have $q = \lfloor p/9 \rfloor$.

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

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

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 (τ)	$\eta_\varepsilon(\mathsf{A}_{p-1}^\vee) \approx \sqrt{\frac{1}{2}} \cdot \eta_{2\varepsilon/p}(\mathbf{Z}^{p-1})$	$O(p)$
Ours (ϕ)	$\eta_{\varepsilon/q}(\mathsf{A}_8) \approx \sqrt{\frac{9}{8}} \cdot \eta_{2\varepsilon/9}(\mathbf{Z}^{8q})$	$O(p)$

samplings. They are both highly parallelizable, thanks to the filtration shown in Proposition 5.1; 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.

6.3.2. *Comparisons with other methods.* In [21, 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 **Z**-basis of the ideal. To the best of our knowledge, we give the first concrete Gaussian sampler in prime-power cyclotomic rings. From the map ϕ, we can directly identify the Gram matrix of $\mathcal{R}_p$ as a scaling by p of that of $\mathsf{A}_{p-1}^\vee$:

$$G_p^\vee = \begin{pmatrix} p-1 & -1 & \cdots & -1 & -1 \\ -1 & p-1 & \cdots & -1 & -1 \\ \vdots & \vdots & \vdots & \vdots & \vdots \\ -1 & -1 & \cdots & -1 & p-1 \end{pmatrix}.$$

The largest element in the diagonal of the Cholesky of G_p is $\sqrt{p-1}$, which drives the quality of a Klein approach (as done in [15]). An approach *à la Peikert* with Algorithm 3 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 (6) (equivalently, using the map ϕ), we have identified the meaningful quantities in Section 5.2. Comparisons with our approaches are displayed in Table 2.

6.3.3. *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 reducing both the public key and signature sizes. All in all, after compression, we gain between **50** and **200** bytes at each security level on the size using our new sampler (for fairness of the comparison, the signature size is not optimized using the recent results of [17]).

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

	Conductor	Modulus	Quality	Security	Pub. key size	Sig. size
	$m : \varphi(m)$	q	α	(C/Q/NIST level)	(bytes)	(bytes)
FALCON	1024 : 512	12289	1.17	124/112/NIST-I		666
MITAKA	1994 : 648	3889	2.13	136/123/NIST-I		827
This work				***/***/NIST-I		
FALCON	N/A	N/A	N/A	N/A	N/A	N/A
MITAKA	2304 : 768	18433	2.20	167/151/NIST-II		1080
This work	683 : 682	1367	2.125	**157/138/NIST-II**		**799**
FALCON	N/A	N/A	N/A	N/A	N/A	N/A
MITAKA	2592 : 864	10369	2.25	192/174/NIST-III		1176
This work	857 : 856	6857	2.215	**207/182/NIST-III**		**1120**
FALCON	N/A	N/A	N/A	N/A	N/A	N/A
MITAKA	2916 : 972	17497	2.30	220/199/NIST-IV		1359
This work	919 : 918	3677	2.247	**223/196/NIST-IV**		**1155**
FALCON	2048 : 1024	12289	1.17	285/258/NIST-V	1792	1280
MITAKA	2048 : 1024	12289	2.33	233/211/NIST-V	1792	1405
This work	1009 : 1008	10091	2.30	**250/219/NIST-V**		**1360**

References

[1] Divesh Aggarwal and Noah Stephens-Davidowitz. Just take the average! an embarrassingly simple 2ˆ 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] Jean-Benoît Bost. Theta invariants of euclidean lattices and infinite-dimensional hermitian vector bundles over arithmetic curves, 2015.

[5] Yilei Chen, Nicholas Genise, and Pratyay Mukherjee. Approximate trapdoors for lattices and smaller hash-and-sign signatures. In Steven D. Galbraith and Shiho Moriai, editors, *ASIACRYPT 2019, Part III*, volume 11923 of *LNCS*, pages 3–32. Springer, Heidelberg, December 2019.

[6] 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.

[7] 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.

[8] J. Conway and N. Sloane. A fast encoding method for lattice codes and quantizers. *IEEE Transactions on Information Theory*, 29(6):820–824, 1983.

[9] J. H. Conway and N. J. A. Sloane. *Sphere Packings, Lattices and Groups*. Grundlehren der Mathematischen Wissenschaften 290. Springer-Verlag, New York, 1988.

[10] Léo Ducas, Eike Kiltz, Tancrède 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.

[11] 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.

[12] 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.

[13] Léo Ducas and Thomas Prest. Fast Fourier Orthogonalization. In *ISSAC 2016*, pages 191–198, 2016.

[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] Thomas Espitau, Mehdi Tibouchi, Alexandre Wallet, and Yang Yu. Shorter hash-and-sign lattice-based signatures. In *Crypto 2022*, 2022.

[18] Craig Gentry. Fully homomorphic encryption using ideal lattices. In Michael Mitzenmacher, editor, *41st ACM STOC*, pages 169–178. ACM Press, May / June 2009.

[19] 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.

[20] M.J.C. Gover. *The eigenproblem of a tridiagonal 2-Toeplitz matrix*. Linear algebra and its applications. 1994.

[21] 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.

[22] 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.

[23] Jacques Martinet. *Perfection and Eutaxy*, pages 67–108. 2003.

[24] 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.

[25] 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.

[26] Daniele Micciancio and Oded Regev. Worst-case to average-case reductions based on Gaussian measures. *SIAM Journal on Computing*, 37(1):267–302, 2007.

[27] Daniele Micciancio and Jessica Sorrell. Simpler statistically sender private oblivious transfer from ideals of cyclotomic integers. In Shiho Moriai and Huaxiong Wang, editors, *ASIACRYPT 2020, Part II*, volume 12492 of *LNCS*, pages 381–407. Springer, Heidelberg, December 2020.

[28] Chris Peikert. Limits on the hardness of lattice problems in l_p norms, 2008.

[29] 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.

[30] 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.

[31] 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.

[32] Thomas Prest. *Gaussian Sampling in Lattice-Based Cryptography*. PhD thesis, École Normale Supérieure, Paris, France, 2015.

[33] 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.

[34] Wessel PJ van Woerden. *The closest vector problem in cyclotomic lattices*. PhD thesis, Leiden University, 2016.

Supplementary Material

Appendix A. Additional content on lattices

A.1. **On effective lifting.** Below is the pseudo-code for Nearest Plane based lifting.

Algorithm 5: Lift (by Babai's nearest plane)

```
Input: A lattice basis $\mathbf{B} = (\mathbf{b}_1, \ldots, \mathbf{b}_n)$ of $\Lambda'$ in $\Lambda$, a vector $\mathbf{t} \in \Lambda$.
Result: A vector $\mathsf{s} \in \Lambda$ in the coset $\mathbf{t} + \Lambda'$.
Compute the Gram-Schmidt orthogonalization $(\mathbf{b}_1^*, \ldots, \mathbf{b}_k^*)$ of $\mathbf{B}$
$\mathbf{s} \leftarrow \mathbf{t}$
for $i = k$ downto $1$ do $\mathbf{s} \leftarrow \mathbf{s} - \left\lfloor \frac{\langle \mathbf{s}, \mathbf{b}_i^* \rangle}{\|\mathbf{b}_i^*\|^2} \right\rceil \mathbf{b}_i$
return $\mathbf{s}$
```

A.2. **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 $\mathbf{Z}$.

Appendix B. Additional background on cyclotomic fields

B.1. **Cyclotomic fields.** 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 $\mathrm{Tr}(f) = \sum_{i\in\mathbf{Z}_m^*} \psi_i(f)$. It gives an inner product as $\mathrm{Tr}(fg^*) = \langle\psi(f), \psi(g)\rangle$ and an euclidean norm $\|\psi(f)\|^2 = \mathrm{Tr}(ff^*)$. 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 p-th primitive roots $\zeta_{p,1}, \dots, \zeta_{p,p-1}$ of 1 is

$$V_p = \begin{bmatrix} 1 & \zeta_{p,1} & \zeta_{p,1}^2 & \dots & \zeta_{p,1}^{p-2} \\ 1 & \zeta_{p,2} & \zeta_{p,2}^2 & \dots & \zeta_{p,2}^{p-2} \\ \vdots & & \dots & & \\ 1 & \zeta_{p,p-1} & \zeta_{p,p-1}^2 & \dots & \zeta_{p,p-1}^{p-2} \end{bmatrix}.$$

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

B.2. **Identifications between cyclotomic ideals and root lattices.** Most of the next facts are well-known, and can be found e.g. in [14, 21, 34]. 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.

We note $(\mathbf{e}_i)$ the canonical basis of $\mathbf{R}^p$. Let $\mathbf{b}_i = \frac{p-1}{p}\mathbf{e}_i - \frac{1}{p}\sum_{1 \leqslant i \neq j \leqslant p} \mathbf{e}_j \in \mathbf{Z}^p$. It is known that $\mathsf{A}_{p-1}^{\vee}$ has basis given by $(\mathbf{b}_i)_{1 \leqslant i \leqslant p-1}$. The key fact for what follows is the isometry between $\mathcal{R}_p$ and $\mathsf{A}_{p-1}^{\vee}$:

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

The normalization for the isometry is 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 $\tau^{\vee} : \mathcal{R}_p^{\vee} \to \mathsf{A}_{p-1}$ sending the dual basis of $1, \zeta_p, \ldots, \zeta_p^{p-1}$ for the trace product, also called *decoding basis* by [21], to the basis $(\mathbf{e}_1 - \mathbf{e}_i)_{2 \leqslant i \leqslant p}$ of 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 $(1-\zeta_p), \ldots, \zeta_p^{p-2} - \zeta_p^{p-1}$. Letting $u_i = \zeta_p^i - \zeta_p^{i+1}$, we have:

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

Up to a factor p, this is the Gram matrix G_{p-1} (6) associated to the standard basis of A_{p-1}. Hence, if we define

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

it identifies both lattices and we have $\|\psi(x)\|^2 = p\|\phi(x)\|^2$.

Appendix C. Proofs of Section 2

C.1. **Cosets mass.** 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 $\mathbf{x} \mapsto \mathbf{x}^t \Sigma^{-1} \mathbf{x}$. Then we have $\rho_\Sigma(\mathbf{x} + \Lambda) \leqslant \rho_\Sigma(P(\mathbf{x})) \cdot \rho_\Sigma(\Lambda)$. If moreover Λ is primitive in Λ', we have $\rho_\Sigma(\Lambda') \leqslant \rho_\Sigma(\Lambda)\rho_\Sigma(P(\Lambda'))$. The equality case occurs when $\Lambda' = \Lambda \perp P(\Lambda')$.

Proof (Adapted from [4, Lemma 2.8.2]*).* We recall the notation: let $\Lambda_{\mathbf{R}}$ be the real space spanned by Λ, and P be the orthogonal projection onto $\Lambda_{\mathbf{R}}^{\perp}$, where the orthogonality is for the quadratic form $\mathbf{x} \mapsto \mathbf{x}^t \Sigma^{-1} \mathbf{x}$. For all $\mathbf{x} \in \mathbf{R}^m$, we can write elements of $\mathbf{x} + \Lambda$ as $\mathbf{x} + \mathbf{u} = (\mathbf{x} - P(\mathbf{x}) + \mathbf{u}) + P(\mathbf{x})$, where $\mathbf{x} - P(\mathbf{x}) + \mathbf{u} \in \Lambda_{\mathbf{R}}$. By orthogonality, we write $\rho_\Sigma(\mathbf{x}+\Lambda) = \rho_\Sigma(P(\mathbf{x})) \cdot \sum_{\mathbf{u} \in \Lambda} \rho_\Sigma(\mathbf{x} - P(\mathbf{x}) + \mathbf{u})$. The right-hand sum is $\rho_\Sigma(\mathbf{x} - P(\mathbf{x}) + \Lambda)$ so we get the first claim with Lemma 2.1. When Λ is primitive in Λ', then $P(\Lambda')$ identifies to the lattice Λ'/Λ, and the sum over all $\mathbf{x} \in P(\Lambda')$ gives the second claim. ■

Appendix D. Proofs of Section 3

D.1. **Short exact sequences and sampling.** For completeness, we recall the *Poisson Summation Formula* for Gaussian functions. For any rank n lattice Λ, we have

$$\rho_s(\Lambda) = \frac{s^n}{\det \Lambda} \cdot \rho_{1/s}(\Lambda^\vee). \tag{9}$$

[Modularity of smoothing parameter] Let Λ be a lattice and $0 < \varepsilon < \sqrt{17} - 4$, then

$$\eta_{3\varepsilon}(\Lambda) \leqslant \min_{\Lambda' \subset \Lambda} \max\Big(\eta_\varepsilon(\Lambda'), \eta_\varepsilon\Big({}^{\Lambda}\!\big/_{\overline{\Lambda'}}\Big)\Big).$$

where $\eta_\varepsilon(\Lambda/\overline{\Lambda'})$ is set by convention to 0 if the quotient is of torsion, and where the minimum ranges over all possible sublattices of Λ.

Proof. Let Λ' be any sublattice of Λ and $s \geqslant \max(\eta_\varepsilon(\Lambda'), \eta_\varepsilon(\Lambda/\overline{\Lambda'}))$. Now, consider the orthogonal projection π of Λ onto the orthogonal space to $\Lambda'_{\mathbf{R}}$, for the standard inner product. Lemma 2.2.1 gives that $\rho_s(\Lambda) \leqslant \rho_s(\overline{\Lambda'}) \cdot \rho_s(\pi(\Lambda))$. Using Identity (9) and the fact that $\det(\Lambda) = \det(\pi(\Lambda)) \det(\overline{\Lambda'})$, this is equivalent to $\rho_{1/s}(\Lambda^\vee) \leqslant \rho_{1/s}(\overline{\Lambda'}^\vee)\rho_{1/s}(\pi(\Lambda)^\vee)$. Because they have the same rank, $\Lambda' \subset \overline{\Lambda'}$ is equivalent to $\overline{\Lambda'}^\vee \subset \Lambda'^\vee$, so we have $\rho_{1/s}(\overline{\Lambda'}^\vee) \leqslant \rho_{1/s}(\Lambda'^\vee)$. By assumption on s, this implies $\rho_{1/s}(\Lambda^\vee) \leqslant (1+\varepsilon)^2$, and we conclude by noting that our choice of sublattice was arbitrary. ■

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

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

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 T the torsion subgroup of Λ/Λ'. Independence of the sampling oracles of lines 3 and 4 and law of total probability yields:

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

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

$$\Pr(\mathbf{q} = \pi(\mathbf{v})) \in \left[\frac{1-\delta}{1+\delta}\frac{1-\varepsilon}{1+\varepsilon}, \frac{1+\delta}{1-\delta}\frac{1+\varepsilon}{1-\varepsilon}\right] \cdot \frac{1}{|\mathrm{T}|} \cdot \mathsf{D}_{\pi(\Lambda),\pi(\mathbf{t}),\Sigma}(\pi(\mathbf{v})). \tag{10}$$

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

$$\begin{aligned}\Pr(\mathbf{u}' = \mathbf{v} - \mathbf{u}_q | \mathbf{q} = \pi(\mathbf{v})) &= \frac{\rho_\Sigma(\mathbf{v} - \mathbf{u}_q - (\mathrm{Id} - \pi)(\mathbf{t} - \mathbf{u}_q))}{\rho_\Sigma(\Lambda' - (\mathrm{Id} - \pi)(\mathbf{t} - \mathbf{u}_q))} \\ &= \frac{\rho_\Sigma((\mathrm{Id} - \pi)(\mathbf{v} - \mathbf{t}))}{\rho_\Sigma(\Lambda' - (\mathrm{Id} - \pi)(\mathbf{t} - \mathbf{u}_q))}.\end{aligned} \tag{11}$$

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

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

We then find that:

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

Routine calculations conclude the proof. ■

D.2. **On filtration bounds and the filtered sampler.** Let $k \geqslant 1$ be an integer, Λ a lattice and $\varepsilon \leqslant e^{-k/2}/\sqrt{2}$. We have

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

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

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 $\rho_\Sigma(\Lambda) \leqslant \rho_\Sigma(P_{k-1}(\Lambda)) \cdot \rho_\Sigma(\Lambda_{k-1})$. Letting P_0 be the identity, we obtain by induction $\rho_\Sigma(\Lambda) \leqslant \prod_{i=1}^k \rho_\Sigma(P_{i-1}(\Lambda_i)) = \prod_{i=1}^k \rho_\Sigma(\Lambda_i/\Lambda_{i-1})$. Using the Poisson Summation Formula (9) and taking $\Sigma \succ \max_i \eta_{\varepsilon/(k+1)}(\Lambda_i/\Lambda_{i-1})$ gives

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

Calculations concludes the proof.

■

[Correctness of the filtered sampler] Algorithm 2 is correct. Moreover, let $\mathcal{D}$ be the distribution of its output. For small enough ε, we have

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

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

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 $\{0\} \subset \Lambda_2/\Lambda_1 \subset \cdots \subset \Lambda/\Lambda_1$. 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 $\{0\} \subset \Lambda_{i+1}/\Lambda_i \subset \cdots \subset \Lambda/\Lambda_i$, with a corresponding flag of subspaces $(V_{i+1} \cap V_i^{\perp}) \subset \ldots \subset (V \cap V_i^{\perp})$. 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, and Algorithm 2 outputs an element in the lattice thanks to the properties of Algorithm 5.

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 2 outputs some vector $\mathbf{v}$ belongs to $\left[\delta^{k-1}, \delta^{1-k}\right] \cdot \mathcal{D}_{\Lambda',\mathbf{t}',\Sigma}(\mathbf{v})$, 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 $k \geqslant 1$. By construction, we can write $\mathbf{u} = \mathbf{z} + \mathbf{v}$ for some $\mathbf{v} \in V_1$. In particular, we have $(\mathrm{Id} - \pi)(\mathbf{t} - \mathbf{u}) = (\mathrm{Id} - \pi)(\mathbf{t}) - \mathbf{v}$. Next, let $P(\mathbf{z})$ resp. $P(\mathbf{u}')$ the probability to obtain $\mathbf{z}$ resp. $\mathbf{u}'$. Orthogonality gives us $\rho_\Sigma(\mathbf{u} + \mathbf{u}' - \mathbf{t}) = \rho_\Sigma(\mathbf{z} - \pi(\mathbf{t}))\rho_\Sigma(\mathbf{u}' - (\mathrm{Id} - \pi)(\mathbf{t}) + \mathbf{v})$. The induction hypothesis then yields

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

Since $\Sigma \geqslant \eta_\varepsilon(\Lambda_1)$, we have $\rho_\Sigma(\Lambda_1 - (1 - \pi)(\mathbf{t}) + \mathbf{v}) \in [\delta, \delta^{-1}] \cdot \rho_\Sigma(\Lambda_1 - (\mathrm{Id} - \pi)(\mathbf{t}))$. With a last orthogonality argument, we obtain our claim on the distribution. ■

Appendix E. Complements and proof of Section 4

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

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

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 $\mathbf{p} \in \mathbf{C_R} := \mathrm{span}_{\mathbf{R}}(\mathbf{C}), \mathbf{x} \in \mathcal{L}(\mathbf{C})$ and outputting $\mathbf{y}$ is the marginalized distribution

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

Gaussian functions have also good multiplicative properties (see for instance [29, Fact 2.10]): there exists[15] a positive definite matrix Σ' and a vector $\mathbf{t}'$, both over $\mathbf{C_R}$, such that:

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

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

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

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

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

Observe that if $m = n$, we have $\mathbf{T}^\star = \mathbf{T}^{-1}$ 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

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

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

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

Appendix F. Application: sampling in tensor lattices

As a direct application of Section 4, 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.

F.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. 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) $\eta_\varepsilon(\mathcal{L}(\mathbf{I}_k\otimes\mathbf{A})) \geqslant \eta_{\varepsilon/k}(\mathcal{L}(\mathbf{A})) \geqslant \eta_{2\varepsilon}(\mathcal{L}(\mathbf{I}_k\otimes\mathbf{A}))$.
(3) $\eta_\varepsilon(\mathcal{L}(\mathbf{I}_k\otimes\mathbf{A})) = \eta_\varepsilon(\mathcal{L}(\mathbf{A}\otimes\mathbf{I}_k))$;

Proof. Since $\mathbf{I}_k\otimes\mathbf{A} = \mathrm{diag}(\mathbf{A},\ldots,\mathbf{A})$, 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

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

Taking $s = \eta_\varepsilon(\mathcal{L}(\mathbf{I}_k \otimes \mathbf{A}))$ gives $\rho_{1/s}(\mathcal{L}(\mathbf{A})) \leqslant (1+\varepsilon)^{1/k} \leqslant 1 + \varepsilon/k$, so $s \geqslant \eta_{\varepsilon/k}(\mathcal{L}(\mathbf{A}))$. Similarly, taking $s = \eta_{\varepsilon/k}(\mathcal{L}(\mathbf{A}))$ gives $\rho_{1/s}(\mathcal{L}(\mathbf{I}_k \otimes \mathbf{A})) \leqslant (1+\varepsilon/k)^k \leqslant 1 + 2\varepsilon$, 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 $\rho_{1/s}(\mathcal{L}(\mathbf{I}_k \otimes \mathbf{A})) = \rho_{1/s}(\mathcal{L}(\mathbf{A} \otimes \mathbf{I}_k))$. ■

F.1.1. *Smoothing bound for tensors.* Combining these properties with Lemma 4.1 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 [25, Corollary 2.7] is associated to the maximal length $\|\mathbf{B}\|_{GS}$ of the Gram-Schmidt vectors of that factor. It is known that $\|\mathbf{B}\|_{GS} \leqslant s_1(\mathbf{B})$, 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 [29] and Klein's randomized nearest plane (e.g., [19]).

Lemma F.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* $m_i \geqslant n_i$*. We have*

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

F.2. **A sampling algorithm for tensor lattices.** Using Algorithm 3, this bound directly translates into Algorithm 6, where it is assumed that oracles for discrete Gaussian sampling in $\mathcal{L}(\mathbf{A})$ and $\mathcal{L}(\mathbf{B})$ are given.

Algorithm 6: 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 $m_i \geqslant n_i$, giving a basis
       of $\Lambda = \mathcal{L}(\mathbf{A} \otimes \mathbf{B})$; a center $\mathbf{t} \in \Lambda \otimes \mathbf{R}$; $r \geqslant 0$; 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
        $\mathcal{D}_{\Lambda, \mathbf{t}, \Delta}$.
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
    $s_1(\mathbf{T}) \cdot \eta_\varepsilon(\mathcal{L}(\mathbf{C}))$
$\Sigma \leftarrow (\mathbf{T}^t \Delta^{-1} \mathbf{T})^{-1}$
return Algorithm 3$(\mathbf{T}, \mathbf{C}, \mathbf{t}, \Delta, \Sigma, r)$
```

Corollary F.1 (of Theorem 4)**.** *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 $m_i \geqslant n_i$. 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 $r \geqslant \eta_\varepsilon(\mathbf{C})$. Algorithm 6 is correct when $s_n(\Delta) \geqslant r^2 \cdot s_1(\mathbf{T})^2$. Moreover, let $\mathcal{D}$ be the distribution of its output. For $\varepsilon < 1/2$, we have*

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

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 Supplementary Material F.1.

Appendix G. On the root lattice samplers

In this section, we details for the ad-hoc samplers listed in Theorem 5.1.

G.1. 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 $(\mathbf{e}_i)_{i\leqslant n}$ 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 3.1, from which Proposition G.1 is obtained.

Algorithm 7: Face centered cubic sampler

```
Input: A parameter $\sigma \geqslant \eta_\varepsilon(\mathsf{D}_n)$ and a center $\mathbf{t} = \sum t_i\mathbf{e}_i \in \mathbf{R}^n$.
Output: $\mathbf{v} \in \mathcal{D}_n$ following distribution statistically close to $\mathcal{D}_{\mathsf{D}_n,\sigma,\mathbf{t}}$
repeat
    $z_1 \leftarrow \mathcal{D}_{\mathbf{Z},\sigma,t_1}, \dots, z_n \leftarrow \mathcal{D}_{\mathbf{Z},\sigma,t_n}$
until $\sum z_i \in 2\mathbf{Z}$
return $\mathbf{z} = \sum z_i\mathbf{e}_i$
```

Proposition G.1. *Let $\mathcal{D}$ the distribution of outputs of Algorithm 7. Then $\mathcal{D}$ is at statistical distance at most ε of $\mathcal{D}_{\mathsf{D}_n,\sigma,\mathbf{t}}$. Moreover, we have*

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

and it requires two tries in average.

G.2. Sampling in $\mathsf{A}_n^\vee$.

The algorithm is straightforward from Lemma 3.2 and the identification $\mathsf{A}_n^\vee = \pi_{\mathbf{1}^\perp}(\mathbf{Z}^{n+1})$, where $\mathbf{1} = (1,\dots,1)$. It is displayed below for referencing purposes. No smoothing condition is needed for correctness. Of course the statistical properties of the output could be quite far from being Gaussian-like if one would sample with very small standard deviation.

Algorithm 8: $\mathsf{A}_n^\vee$ sampler

```
Input: A parameter $\sigma > 0$ and a center $\mathbf{t} \in \mathbf{1}^\perp$.
Output: $\mathbf{v} \in \mathsf{A}_n^\vee$ following distribution statistically close to $\mathcal{D}_{\mathsf{A}_n^\vee,\mathbf{t},\sigma^2}$
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 $\mathbf{y}$
```

G.3. **Sampling in A_2.** It is known that a rank 2 lattice is always similar to its dual; the similarity in the case of A_2 is the composition of the rotation of $\mathbf{1}^\perp$ by $-\pi/2$ and fixing $\mathbf{1}$ with a scaling by $\sqrt{3}$. Algorithm 9 below follows this idea: project a sample $\mathbf{z} \leftarrow \mathcal{D}_{\mathbf{Z}^3,\mathbf{t}}$ orthogonally onto the hyperplane $\mathbf{1}^\perp$ then apply the similarity.

Algorithm 9: Hexagonal sampler

```
Input: A parameter $\sigma > 0$ and a center $\mathbf{t} \in \mathbf{1}^\perp$.
Output: $\mathbf{v} \in \mathsf{A}_2$ following distribution statistically close to $\mathcal{D}_{\mathsf{A}_2,\sigma,\mathbf{t}}$
1 $\mathbf{y} \leftarrow$ $\mathsf{A}_n^\vee$ sampler$(\sigma/\sqrt{3}, \mathbf{t})$ // Algorithm 8 with $n = 2$.
2 return $\mathbf{x} \leftarrow \sqrt{3}R_{-\pi/2}(\mathbf{y})$
```

Again, no smoothness condition is needed, thanks to the orthogonal decomposition.

Theorem G.1. *Algorithm 9 is correct. The distribution of its output is as close to $\mathcal{D}_{\mathsf{A}_2,\mathbf{t},\sigma}$ as the Gaussian integer sampler is close to $\mathcal{D}_{\mathbf{Z}^3,\mathbf{t}',\sigma/\sqrt{3}}$, for any lift $\mathbf{t}'$ of $\mathbf{t}/\sqrt{3}$ in $\mathbf{R}^3$.*

Proof. At Step 3, $\mathbf{y}$ is computed by orthogonal projection onto $\mathbf{1}^\perp$ and thus belongs to $\mathsf{A}_2^\vee$. The probability to obtain $\mathbf{y}$ is therefore

$$P(\mathbf{y}) = P(\mathbf{z} \in \mathbf{y} + (\mathbf{Z}^3 \cap \mathbf{1})) = \frac{\rho_{\sigma/\sqrt{3}}((\mathbf{Z}^3 \cap \mathbf{1}) + \mathbf{y} - \mathbf{t}')}{\rho_{\sigma/\sqrt{3}}(\mathbf{Z}^3 - \mathbf{t}')}.$$

By orthogonality, we can write $\rho_{\sigma/\sqrt{3}}(\mathbf{y} - \mathbf{t}' + (\mathbf{Z}^3 \cap \mathbf{1})) = \rho_{\sigma/\sqrt{3}}(\mathbf{y} - \mathbf{t}/\sqrt{3})\rho_{\sigma/\sqrt{3}}((\mathbf{Z}^3 \cap \mathbf{1}) - \mathbf{t}' + \mathbf{t}/\sqrt{3})$, and decompose similarly the coset at the denominator. This gives us

$$P(\mathbf{y}) = \frac{\rho_\sigma(\sqrt{3}\mathbf{y} - \mathbf{t})}{\rho_\sigma(\sqrt{3}\mathsf{A}_2^\vee - \mathbf{t})},$$

and the result follows from the radiality of the Gaussian function. ■

G.4. **Sampling in the A_3 lattice.** There exists an isometry between D_3 and A_3 (coming from the exceptional Lie isomorphisms in small dimnensions), 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 $\sigma \geqslant \eta_\varepsilon(\mathsf{A}_3)$ using (an expected number of) three samples in $\mathbf{Z}^4$

Remark. *A similar algorithm would enable to sample in the dual lattice* $\mathsf{D}_3^\vee = \mathsf{D}_3 \cup ((1/2,1/2,1/2)+\mathsf{D}_3)$ *(disjoint union).*

G.5. **Sampling in the E_8 lattice.** The E_8 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, however, isometric to $\mathbf{Z}^8$. As seen e.g. in [23, 7], it can be written as a disjoint union of cosets as $\mathsf{E}_8 = \mathsf{D}_8 \cup (1/2,\dots,1/2)+\mathsf{D}_8$, 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 3.1, from which Proposition G.2 is obtained.

Algorithm 10: E8 sampler

```
Input: A parameter $\sigma \geqslant \eta_\varepsilon(\mathsf{D}_8)$ and a center $\mathbf{t} \in \mathbf{R}^8$.
Output: $\mathbf{v}$ following distribution statistically close to $\mathcal{D}_{\mathsf{D}_8,\sigma,\mathbf{t}}$
$b \leftarrow$ Bernoulli
$\mathbf{z} \leftarrow$ Algorithm 7$(\sigma, \mathbf{t} - b\cdot\mathbf{h})$
return $b\cdot\mathbf{h} + \mathbf{z}$
```

Proposition G.2. *let $\mathcal{D}$ the distribution of outputs of Algorithm 10. Then $\mathcal{D}$ is at statistical distance at most 2ε of $\mathcal{D}_{\mathsf{E}_8,\sigma,\mathbf{t}}$. Moreover, we have*

$$\sup_{\mathbf{z}\in\mathsf{E}_8}\left|\frac{\mathcal{D}(\mathbf{z})}{\mathcal{D}_{\mathsf{E}_8,\sigma,\mathbf{t}}(\mathbf{z})} - 1\right| \leqslant 4\varepsilon$$

G.6. **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 $1 \leqslant i \leqslant 7$ all belong to E_8, and their Gram matrix is identical to the Gram matrix of the basis $(\mathbf{f}_1 - \mathbf{f}_i)_{1 \leqslant i \leqslant 8}$ of A_8, where $(\mathbf{f}_i)_i$ is the canonical basis of $\mathbf{R}^9$. As $\det \mathsf{A}_8 = 3$ 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 3.2.

```
Algorithm 11: A8 sampler

Input: A parameter $\sigma \geqslant \max(\eta_\varepsilon(\mathsf{D}_8), \eta_\varepsilon(\mathsf{A}_8))$ and a center $\mathbf{t} \in \mathbf{R}^3$.
Output: $\mathbf{v}$ following distribution statistically close to $\mathsf{D}_{\mathsf{A}_8,\sigma,\mathbf{t}}$

repeat
    $\mathbf{z} \leftarrow$ E8sampler$(\sigma, \mathbf{t})$
until $\mathbf{z}_1 - \mathbf{z}_2 - \cdots - \mathbf{z}_8 \in 3\mathbf{Z}$
return $\mathbf{z}$
```

Proposition G.3. *let $\mathcal{D}(\mathbf{z})$ the distribution of outputs of Algorithm 11. Then $\mathcal{D}(\mathbf{z})$ is at statistical distance at most 10ε of $\mathsf{D}_{\mathsf{A}_8,\sigma,\mathbf{t}}$. Moreover, we have*

$$\sup_{\mathbf{z} \in \mathsf{A}_8} \left| \frac{\mathcal{D}(\mathbf{z})}{\mathsf{D}_{\mathsf{A}_8,\sigma,\mathbf{t}}(\mathbf{z})} - 1 \right| \leqslant 20\varepsilon$$

Proof. Let Λ be the sublattice of E_8 generated by $(\mathbf{h}, \mathbf{b}_1, \dots, \mathbf{b}_7)$ and $\mathbf{u} = \mathbf{e}_1 - \mathbf{e}_2 - \cdots \mathbf{e}_8$. We show that $\Lambda = \{\mathbf{x} \in \mathsf{E}_8 \ : \ \langle \mathbf{x}, \mathbf{u} \rangle \in 3\mathbf{Z}\}$, which means that Algorithm 11 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 $\mathbf{x} \in \mathsf{E}_8$ in particular, we also know that $\mathbf{x}$ is either in D_8 either in $\mathbf{h} + \mathsf{D}_8$. 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 $2 \leqslant i \leqslant 8, a_i = x_i - a_1/2$ 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 $\mathbf{x} = a_1 \mathbf{h} + a_2 \mathbf{b}_1 + \cdots a_8 \mathbf{b}_7$, or equivalently, $\mathbf{x} \in \Lambda$. The rest of the proof comes from from Proposition 3.1. ■

G.7. **Sampling in the E_7 lattice.** We make use of the (dual covering) $L = (\mathsf{E}_7^\vee \perp \mathbf{Z}) = \mathsf{E}_8 \cup \left(\frac{1}{2}f + \mathsf{E}_8\right)$ where $f = (0.\dots,0,1,1)^T \in \mathbf{Z}^8$ (see [23, Chap.4, p.118]). Hence using Algorithm 10 we can use our E_8 sampler to sample into L and project orthogonaly onto $\mathsf{E}_7^\vee$. We now use the fact that since E_7 is integral it is contained in its dual and use the Proposition 3.1. As $|\mathsf{E}_7| = 2$, the quotient $\mathsf{E}_7^\vee/\mathsf{E}_7$ is of cardinality 4, meaning that we expect 4 repetitions of this process on average.

G.8. **Smoothing parameter estimates for some exceptional latttices.** We have the following estimates for $\varepsilon > 0$:

- $\eta_\varepsilon(\mathbf{Z}^n) = \sqrt{\frac{1}{\pi}(\ln(\frac{2n}{\varepsilon}) + o(1/\varepsilon))}$;
- for $n \geqslant 5$, $\eta_\varepsilon(\mathbf{Z}^n) \leqslant \eta_\varepsilon(\mathsf{D}_n) = \sqrt{\frac{1}{\pi}(\ln(\frac{2n}{\varepsilon}) + o(1/\varepsilon))} \approx \eta_\varepsilon(\mathbf{Z}^n)$;
- $\eta_\varepsilon(\mathsf{A}_n) = \sqrt{\frac{n+1}{n}} \cdot \sqrt{\frac{1}{\pi}(\ln(\frac{2(n+1)}{\varepsilon}) + o(1/\varepsilon))} \approx \lambda_1(\mathsf{A}_n^\vee)^{-1} \cdot \eta_\varepsilon(\mathbf{Z}^n)$;
- $\eta_\varepsilon(\mathsf{A}_n^\vee) = \frac{1}{\sqrt{2}} \cdot \sqrt{\frac{1}{\pi}(\ln(\frac{n(n+1)}{\varepsilon}) + o(1/\varepsilon))} \approx \lambda_1(\mathsf{A}_n)^{-1} \cdot \eta_{2\varepsilon/(n+1)}(\mathbf{Z}^n)$

The estimate of $\eta_\varepsilon(\mathbf{Z}^n)$ is not new, but the proof strategy we give is.

Sketch. The three estimates follow 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 $\theta(z) - 1 = \varepsilon$. A standard fact is that if $h = h_1 z + h_2 z^2 + \cdots$ is a power series, then there always exists another power series $g = g_1 z + g_2 z^2 + \cdots$ 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:

$$\begin{aligned} g_1 &= h_1^{-1} & g_3 &= -h_1^{-5}(2h_2^2 - h_1 h_3) \\ g_2 &= -h_1^{-3} h_2 & g_4 &= -h_1^{-7}(5h_1h_2h_3 - h_1^2 h_4 - 5h_2^3) \end{aligned} \tag{12}$$

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 Lemma G.1 below).

In our context where $z = \exp(-\pi s^2)$, we only really care about the first few terms in the expansion. As explained in Section 5, if S is the reversion of $\rho_{1/s}(\Lambda^\vee) - 1 = \kappa z^{\lambda_1^2} + \cdots$, 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

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

Checking the well-known theta series of these exceptional lattices in [9] is enough to obtain our estimates. ■

Lemma G.1. *Let $\Lambda \subset \mathbf{Q}^m$ be any lattice, and let $T_\Lambda(q) = \theta_\Lambda(q) - 1 = \kappa_1 q^{\lambda_1(\Lambda^\vee)^2} + \ldots$. There exists a unique series $S(q) = (q/\kappa_1)^{1/\lambda_1(\Lambda^\vee)^2} + \ldots$ 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, all exponents are positive integers. We let $n = d^2\lambda_1(\Lambda^\vee)^2$, and we consider now the formal series $T(X) = \kappa X^n + \ldots = \kappa X^n(1 + \tilde{T}(X)) \in \mathbf{R}[[X]]$ 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 $S(T(X)) = X$. ■

G.9. **Proofs of Section 5.**

Proof of Proposition 5.1. Recall that for any $n > 1$, the A_n lattice admits the basis

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

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

$$\mathbf{B}_n = \left(\begin{array}{c|c|c|c|c|c} \mathsf{A}_k & & & & \mathbf{C}_1 & \\ \hline & \mathsf{A}_k & & & \mathbf{C}_2 & \\ \hline & & \ddots & & \vdots & \\ \hline & & & \mathsf{A}_k & \mathbf{C}_q & \\ \hline & & \cdots & & \mathbf{C}_{q+1} & \mathsf{A}_r \end{array}\right) \in \mathbb{Z}^{(n+1)\times n},$$

where

$$\mathbf{C}_1 = (\varepsilon_{i,j}^{(1)}) \in \mathbb{Z}^{(k+1)\times q} : \varepsilon_{i,j}^{(1)} = \begin{cases} 1 & (i,j) = (k+1,1) \\ 0 & \text{otherwise,} \end{cases}$$

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

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

Immediately, the result follows. ■

Proof of Theorem 5.2. Let $0 = \Lambda_0 \subset \Lambda_1 \subset \Lambda_2 \subset \cdots \subset \Lambda_q \subset \mathsf{A}_n$ be the filtration given by Proposition 5.1. Since we use the Klein sampler as a base sampler in the quotient Λ_q, the result holds using Theorem 6 when

$$\sigma \geqslant \max\left\{\sigma', \eta_\varepsilon(\mathsf{A}_k), \max_{1\leqslant i\leqslant q+r} \|\mathbf{b}_{kq+i}^*\| \eta_\varepsilon(\mathbf{Z})\right\}$$

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 11 that just needs $\sigma' \geqslant \eta_\varepsilon(\mathsf{A}_8)$. Therefore it suffices to estimate $\|\mathbf{b}_i^*\|$. A routine computation verifies that the Gram-Schmidt orthogonalization of $\mathbf{A}_n$ is

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

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

[⁎] NTT Corporation, Tokyō, Japan, [⋆] Rennes University, Rennes, France, [†] Tsinghua University, Beijing, China

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