A Hybrid Gaussian Sampler for Lattices over Rings

Léo Ducas[1] and Thomas Prest[2]

Abstract. Gaussian sampling over lattices is a cornerstone of lattice-based cryptography as it allows to build numerous cryptographic primitives. There are two main algorithms performing this task. The first one is due to Klein (SODA 2000) and Gentry, Peikert and Vaikuntanathan (STOC 2008), and outputs vectors of good quality but runs rather slowly, in quadratic time. The second one is due to Peikert (CRYPTO 2010) and outputs vectors of slightly worse quality, but can be made to run in quasilinear time in the ring setting.
We present a Gaussian Sampler optimized for lattices over the ring of integer of a cyclotomic number field. At a high-level it works as Klein's sampler but uses an efficient variant of Peikert's sampler as a subroutine. The result is a new sampler that samples vectors with a quality close to Klein's sampler and achieves the same quasilinear complexity as Peikert's sampler. In practice, we get close to the best of both worlds.

Key words: Lattice-based Cryptography, Gaussian Sampling, Klein's sampler, Peikert's Sampler.

1 Introduction

Sampling lattice points is an essential primitive in lattice-based cryptography, as several constructions are based on it. It is used in hash-and-sign signatures [GPV08], standard-model signatures [ABB10,Boy10], (hierarchical) identity-based encryption schemes [GPV08,CHKP10,ABB10], attribute-based encryption [BGG$^+$14], and many other constructions. It can also be used to solve variants of the closest (or shortest) vector problem [Kle00,ADRS14] .

In a nutshell, the shorter the vectors that one can sample given a good basis, the more secure schemes we may build. This holds as long as the distribution of the sampled vectors is a discrete Gaussian (even though other simulatable distributions also yield secure schemes , [LW15]).

Currently, two main algorithms allow to do Gaussian Sampling over arbitrary lattices. The first one is a randomized version of Babai's nearest-plane algorithm [Bab86] due to Klein [Kle00] and Gentry, Peikert, Vaikuntanathan [GPV08]. Currently, it is also the one that produces the shortest vectors but it has a sequential structure and runs in time $\tilde{O}(n^2)$, even on ideal lattices.

[1] Cryptology Group, CWI, Amsterdam, The Netherlands. Supported by an NWO Free Competition Grant. Email: ducas@cwi.nl
[2] École Normale Supérieure, Paris, France and Thales Communications & Security, Gennevilliers, France. Email: prest@di.ens.fr

The second one is due to Peikert [Pei10] and can be seen as a randomized version of Babai's rounding algorithm [Bab86]. In essence, Peikert's sampler is very different from Klein's as its underlying idea essentially is that the convolution of two Gaussians is also a Gaussian. In practice, it performs very differently from Klein's sampler as it is parallelizable. If the underlying lattice is endowed with a ring structure, Peikert's sampler can take advantage of it and run in quasi-linear time $\tilde{O}(n)$, where n is the dimension of the lattice. However, its quality is worse than Klein's sampler—that is, the size of the outputted vector is significantly longer—which undermines the security of underlying cryptographic constructions.

1.1 Our Contribution

The main contribution of our paper is a new discrete Gaussian sampler over ideal lattices. Our algorithm is constructed in two steps. First, we give ring variants for Klein's and Peikert's samplers. Given a ring $\mathcal{R}$, the ring variant of Klein's sampler is a generalization from sampling in lattices in $\mathbb{Z}^m$ to lattices in $\mathcal{R}^m$, whereas the ring variant of Peikert's sampler is simply an instanciation over $\mathcal{R}$.

The ring variant of Klein's sampler may run faster over ring bases than the original algorithm. However, in order to do so it needs to invoke a fast sampler over $\mathcal{R}$ as an internal oracle. So we use the ring variant of Peikert's sampler to be this internal oracle. The resulting algorithm is a hybrid sampler: at high-level it operates as Klein's sampler, but at low level it uses Peikert's algorithm. Our hybrid sampler allows a trade-off between the slow but high-quality sampler of Klein, and the fast but lower-quality sampler of Peikert. As a by-product, we also obtain a hybrid between Babai's nearest plane and rounding algorithms.

We test the practical value of our hybrid sampler by comparing it to Klein's and Peikert's sampler. This is done by evaluating the security of a Full-Domain Hash signature scheme over NTRU lattices using either of these samplers as the core part of the signing procedure. For lattices of fixed dimension and modulus, using the hybrid sampler allows to get a signature scheme a bit less secure than using Klein's sampler (160 bits of security for our sampler versus 192 for Klein's and 120 for Peikert's), while having the same quasilinear complexity as Peikert's algorithm.

1.2 Choice and Implementation of the Ring

In essence, our techniques are independent of the choice of the number field and its ring of integers, and they are relevant as long as the chosen field admits a fast multiplication algorithm.

Our results are tested with typical rings used in lattice-based cryptography so far, that is the ring of integers of the m-th cyclotomic field for m a power of 2. Other choices are of course possible, but choosing a ring comes with the task of studying its geometry (that is essentially finding a good $\mathbb{Z}$-basis of that ring). The cyclotomic cases have been treated extensively [LPR10,DD12,LPR13].

In the light of recent cryptanalytic developments [Ber14a,CDPR15], it is worth noting that our result also applies to the so-called NTRU-prime ring $\mathcal{R} = \mathbb{Z}[x]/(x^p - x - 1)$ as proposed by Bernstein [Ber14a], yet the geometric aspect of such rings remains to be studied.

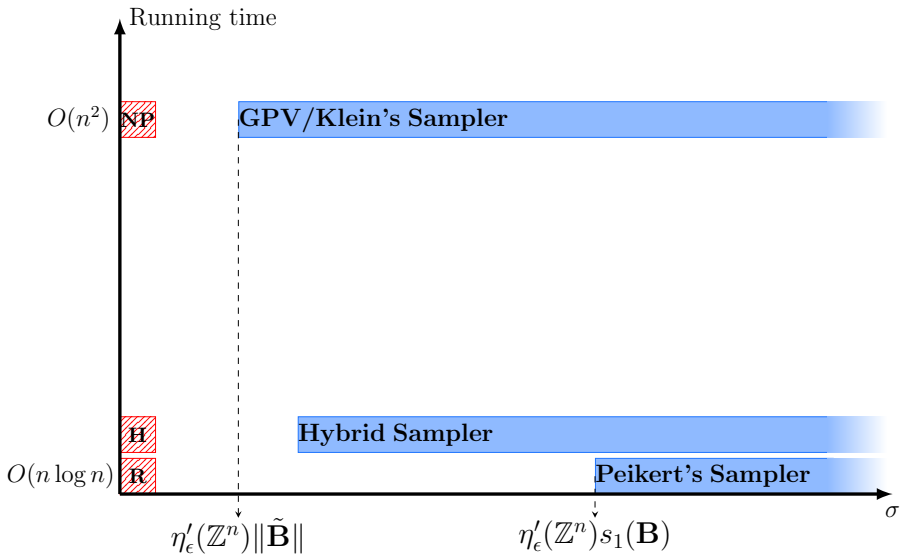

Fig. 1. Our contributions and their performances compared to existing algorithms. On the left are Babai's algorithms: nearest plane (NP), hybrid (H) and rounding (R). On the right are the Gaussian Samplers. For samplers, the lower σ can be set, the better.

1.3 Related Work

The seminal work of [GPV08] spawned a lot of papers trying to improve it. The most notable may be [Pei10], which created a completely different sampler.

An important contribution is the work of [MP12], which introduces the use of a public matrix $\mathbf{G}$ to generate a random matrix $\mathbf{A}$ along with some trapdoor information allowing to do very efficient Gaussian sampling on $\Lambda^\perp(\mathbf{A})$. However, to the best of our knowledge these techniques do not apply to specific families of lattices such as NTRU lattices. More importantly, these techniques imply a huge blowup in the parameters of the lattice: the public key is a matrix $\mathbf{A} \in \mathbb{Z}^{n \times m}$, where n is rank of the lattice and $m > n \ln q$. In comparison, NTRU lattice only have $m = 2n$.

Floating point arithmetic issues were addressed in [DN12], which speeds up both Klein's and Peikert's samplers to $\tilde{O}(n^2)$. For lattices over rings, [DN12] also

uses the ring structure of the lattice to get Peikert's offline phase to reach time and space complexity $\tilde{O}(n)$.

The work of [DLP14] uses a measure called KL Divergence as an alternative to statistical distance in order to assess the security of cryptographic constructions based on Klein's sampler. In addition, they evaluate (experimentally and heuristically) which NTRU bases are the most suitable for these constructions. Rejection sampling techniques are used in [BLP+13] in order to use Klein's sampler with an even shorter standard deviation. The geometric properties of lattices over rings are used in [LP15] to reduce the space requirement of Klein's sampler to $\tilde{O}(n)$.

1.4 Roadmap

In Section 2, we set up the definitions and notations that we will use throughout the paper. In Section 3, we define the ring variants of Klein's and Peikert's samplers. In Section 4, we introduce our hybrid sampler and prove its correctness. In section 5, we assess the practical interest of our new sampler by comparing its efficiency and time complexity to those of Klein's and Peikert's algorithms.

2 Preliminaries

2.1 Notations

Vectors will be written with non-capital bold letters, matrices and bases in capital bold letters, and scalars (which includes ring and field elements) in non-bold letters. A matrix $\mathbf{B} \in \overline{\mathbb{K}}^{m \times n}$ may be viewed as the set of its row vectors $\mathbf{B} = \{\mathbf{b}_1, ..., \mathbf{b}_n\}$. If $\mathbf{B}$ is non-singular, then its set of row vectors form a basis.

2.2 Algebraic Background

Let $\mathbb{K}$ be a number field, i.e. an algebraic extension of $\mathbb{Q}$ of finite degree, and let d be its degree over $\mathbb{Q}$. Let $\bar{\cdot}$ denote the complex conjugation over $\mathbb{K}$, it is an involution and an automorphism of $\mathbb{K}$, which collapses to the identity if and only if $\mathbb{K}$ is a real number field. We let $\mathcal{R}$ be the ring of integers of $\mathbb{K}$ (its maximal order), $\mathbb{K}^+$ be the maximal real subfield of $\mathbb{K}$ and $\mathcal{R}^+ \subset \mathcal{R}$ be the ring of integers of $\mathbb{K}^+$ (we denote d^+ the degree of $\mathbb{K}^+$). We also define the completions $\overline{\mathbb{K}} = \mathbb{R} \otimes_{\mathbb{Q}} \mathbb{K}$ and $\overline{\mathbb{K}^+} = \mathbb{R} \otimes_{\mathbb{Q}} \mathbb{K}^+$. We note that those completions are not necessarily fields.

The number field $\mathbb{K}$ comes with d complex embeddings $\mathbb{K} \mapsto \mathbb{C}$ (forming a set $\mathfrak{S}$) , indexed by $i \in \{1, \ldots, d\}$. Similarly $\mathbb{K}^+$ comes with d^+ real embeddings $\mathbb{K}^+ \mapsto \mathbb{R}$ (forming a set $\mathfrak{S}^+$). Each σ in $\mathfrak{S}$ (resp. $\mathfrak{S}^+$) can be extended to the completion $\overline{\mathbb{K}}$ (resp. $\overline{\mathbb{K}^+}$). An element $e \in \overline{\mathbb{K}}$ (or $\overline{\mathbb{K}^+}$) is invertible if and only if all its embeddings are non-zero. Otherwise, e is said to be singular.

An element e of $\overline{\mathbb{K}^+}$ is said totally positive (and we write $e > 0$) if for all the real embeddings $\sigma \in \mathfrak{S}^+$ we have $\sigma(e) > 0$. Note that if e is totally positive, then

it is invertible. If e is totally positive, it admits 2^{d^+} square roots in $\overline{\mathbb{K}^+}$, and we define its canonical square root $\sqrt{e} \in \overline{\mathbb{K}^+}$ as its unique square root that is totally positive : $\sqrt{e} > 0$. Note that this implies $\sigma(\sqrt{e}) = \sqrt{\sigma(e)}$ for all real embeddings $\sigma \in \mathfrak{S}^+$. This extends naturally to a definition of totally non-negative elements, noted $e \geqslant 0$. This also equips the field $\mathbb{K}^+$ (and its completion $\overline{\mathbb{K}^+}$) with a partial order: $e \geqslant e' \Leftrightarrow e - e' \geqslant 0$.

Hermitian structure of $\overline{\mathbb{K}}$. Seen as a $\mathbb{Q}$-vector space, $\mathbb{K}$ can be equipped with the sesquilinear map $\langle \cdot, \cdot \rangle : \mathbb{K} \times \mathbb{K} \to \mathbb{C}$, $(a, b) = \mathrm{Tr}(a\bar{b}) = \sum_\sigma \sigma(a)\sigma(\bar{b})$. This sesquilinear map extends to $\overline{\mathbb{K}}$. The associated norm $x \mapsto \sqrt{\langle x, x \rangle} \in \mathbb{R}$ is noted $\| \cdot \|$.

Hermitian vector space over $\mathbb{K}$. For vector spaces $H = \mathbb{K}^n$, we can also define an sesquilinear product $\langle \cdot, \cdot \rangle_{\mathbb{K}} : H \times H \to \mathbb{K}$,

$$\langle \boldsymbol{a}, \boldsymbol{b} \rangle_{\mathbb{K}} = \sum a_i \bar{b}_i.$$

One indeed verifies that $\langle \boldsymbol{a}, \boldsymbol{a} \rangle_{\mathbb{K}} \geqslant 0$ for any vector $\boldsymbol{a} \in \mathbb{K}$, and that $\langle \boldsymbol{a}, \boldsymbol{a} \rangle_{\mathbb{K}} \neq 0$ for any non-zero vector $\boldsymbol{a} \in \mathbb{K} \setminus \{\boldsymbol{0}\}$ (we carefully note that it does not imply that $\langle \boldsymbol{a}, \boldsymbol{a} \rangle_{\mathbb{K}} > 0$). The associated norm is given by $\| \cdot \|_{\mathbb{K}} : \boldsymbol{a} \mapsto \sqrt{\langle \boldsymbol{a}, \boldsymbol{a} \rangle}$. This map can be completed to $\overline{H} \times \overline{H} \to \overline{\mathbb{K}}$ where $\overline{H} = \mathbb{R} \otimes_{\mathbb{Q}} V$.

The two sesquilinear maps compose nicely: denoting $\langle \cdot, \cdot \rangle^{\oplus n}$ (resp. $\| \cdot \|^{\oplus n}$) the component-wise application of $\langle \cdot, \cdot \rangle$ (resp. $\| \cdot \|$), for all $a \in \mathbb{K}^n$ we have the following commutative diagrams

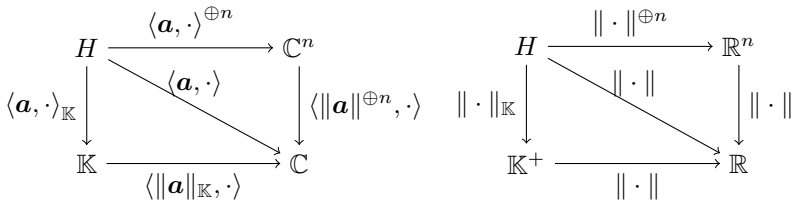

that naturally defines a sesquilinear map $H \times H \to \mathbb{C}$ together with a norm $\| \cdot \| : H \to \mathbb{R}$.

2.3 Lattice over a Ring

A lattice over the ring R is a discrete R-module of $H = \mathbb{K}^n$ equipped with the Euclidean norm described above.

2.4 Gram-Schmidt Orthogonalization over Number Fields

Equipped with this algebraic background, we may now generalize to number fields notions related to the Gram-Schmidt orthogonalization of a basis. For a matrix $\mathbf{B} \in \mathbb{K}^{n \times m}$, the conjugate transpose of $\mathbf{B}$ is, as the name suggests, the $m \times n$ matrix $\mathbf{B}^*$ whose coefficients verify $(\mathbf{B}^*)_{ji} = \overline{\mathbf{B}_{ji}}$.

Definition 1. *Let $\mathbf{B} = \{\mathbf{b}_1, ..., \mathbf{b}_k\} \in H^n$ and $X \subseteq \mathbb{K}$. We note $Span_X(\mathbf{b}_1, ..., \mathbf{b}_k)$ (or $Span_X(\mathbf{B})$) the set $\{\sum_{1 \leqslant i \leqslant k} x_i \mathbf{b}_i, x_i \in X\}$. In particular, $Span_{\mathcal{R}}(\mathbf{B})$ is an $\mathcal{R}$-module and and $Span_{\mathbb{K}}(\mathbf{B})$ is a $\mathbb{K}$-vector space. If the vectors of $\mathbf{B}$ are linearly independent as elements of a $\mathbb{K}$-vector space, we say that $\mathbf{B}$ is a ($\mathbb{K}$-)basis (of $Span_{\mathbb{K}}(\mathbf{B})$).*

Definition 2. *Let $n \leqslant m$, V be a n-dimensional subspace of H and $\mathbf{B} \in \mathbb{K}^{n \times m}$ be a basis of V. For any $\mathbf{x} \in H$, the projection of $\mathbf{x}$ over V is:*

$$\mathbf{Proj}(\mathbf{x}, Span_{\mathbb{K}}(\mathbf{B})) = \mathbf{x}\mathbf{B}^*(\mathbf{B}\mathbf{B}^*)^{-1}\mathbf{B}$$

In particular, if $\mathbf{y} \in \mathbb{K}^m$, then the projection of $\mathbf{x}$ over $Span_{\mathbb{K}}(\mathbf{y})$ is:

$$\mathbf{Proj}(\mathbf{x}, Span_{\mathbb{K}}(\mathbf{y})) = \frac{\langle \mathbf{x}, \mathbf{y} \rangle_{\mathbb{K}}}{\langle \mathbf{y}, \mathbf{y} \rangle_{\mathbb{K}}} \mathbf{y}$$

With the convention $\frac{\langle \mathbf{x}, \mathbf{0} \rangle_{\mathbb{K}}}{\langle \mathbf{0}, \mathbf{0} \rangle_{\mathbb{K}}} \mathbf{0} = \mathbf{0}$.

Noting $p = \mathbf{Proj}(\cdot, Span_{\mathbb{K}}(\mathbf{y}))$, one can check that p is the usual orthogonal projection over V: in particular $p(H) = V, \ker p = V^{\perp}, p \circ p = p$ and $p \mid_V = id \mid_V$.

We now have all the tools to easily define and analyze the (generalized) Gram-Schmidt orthogonalization (or GSO). Unlike most definitions of the GSO, this one doesn't consider vectors of $\mathbb{R}^m$, but of $\mathbb{K}^m$. We first state in Lemma 1 the equivalent properties verified by the GSO of a set, and then formally define the GSO in Definition 3.

Lemma 1. *Let $\mathbf{B} = \{\mathbf{b}_1, ..., \mathbf{b}_n\} \in H^n$ be a basis. For any $k \in [\![1, n]\!]$, we note $V_k \overset{\Delta}{=} Span_{\mathbb{K}}(\mathbf{b}_1, ..., \mathbf{b}_k)$. There is a unique basis $\tilde{\mathbf{B}} = \{\tilde{\mathbf{b}}_1, ..., \tilde{\mathbf{b}}_n\} \in H^n$ verifying any of these equivalent properties:*

1. $\forall k \in [\![1, n]\!], \tilde{\mathbf{b}}_k = \mathbf{b}_k - \mathbf{Proj}(\mathbf{b}_k, V_{k-1})$
2. $\forall k \in [\![1, n]\!], \tilde{\mathbf{b}}_k = \mathbf{b}_k - \sum_{j=1}^{k-1} \frac{\langle \mathbf{b}_k, \tilde{\mathbf{b}}_j \rangle_{\mathbb{K}}}{\langle \tilde{\mathbf{b}}_j, \tilde{\mathbf{b}}_j \rangle_{\mathbb{K}}} \tilde{\mathbf{b}}_j$
3. $\forall k \in [\![1, n]\!], \tilde{\mathbf{b}}_k \perp V_{k-1}$ *and* $(\mathbf{b}_k - \tilde{\mathbf{b}}_k) \in V_{k-1}$

Noting $\tilde{V}_k \overset{\Delta}{=} Span_{\mathbb{K}}(\tilde{\mathbf{b}}_1, ..., \tilde{\mathbf{b}}_k)$, we also have: $\forall k \in [\![1, n]\!], \tilde{V}_k = V_k$.

Proof. We first prove the equivalence of the conditions. They are equivalent at step 1. We suppose it is the case up to step $k-1$ and prove the equivalence at step k:

- $\boxed{1 \Leftrightarrow 2}$ First let us notice that $V_{k-1} = \tilde{V}_{k-1}$ (see condition 2). Observing that for any $j < k$, the $\tilde{\mathbf{b}}_j$'s are pairwise orthogonals (see condition 3), we get for any $\mathbf{v} \in H$:

$$\mathbf{Proj}(\mathbf{v}, V_{k-1}) = \mathbf{Proj}(\mathbf{v}, \tilde{V}_{k-1}) = \sum_{j=1}^{k-1} \frac{\langle \mathbf{v}, \tilde{\mathbf{b}}_j \rangle_{\mathbb{K}}}{\langle \tilde{\mathbf{b}}_j, \tilde{\mathbf{b}}_j \rangle_{\mathbb{K}}} \tilde{\mathbf{b}}_j$$

  Where the second equality comes from the first equality of Definition 2.

- $\boxed{2 \Rightarrow 3}$ $\tilde{\mathbf{b}}_k \perp \tilde{V}_{k-1}$ is a simple computation, and $(\mathbf{b}_k - \tilde{\mathbf{b}}_k) \in V_{k-1}$ is straightforward.
- $\boxed{3 \Rightarrow 1}$ We proved that $(1 \Rightarrow 3)$, so $\mathbf{b}_k = \mathbf{Proj}(\mathbf{b}_k, V_{k-1}) + (\mathbf{b}_k - \mathbf{Proj}(\mathbf{b}_k, V_{k-1}))$ yields a decomposition of $\mathbf{b}_k$ over V_{k-1} and $V_{k-1}^{\perp}$. Since V_{k-1} and $V_{k-1}^{\perp}$ are in direct sum, such a decomposition is unique.

The uniqueness of $\tilde{\mathbf{B}}$ comes from the deterministic formula in condition 2, as does the fact that $V_k = \tilde{V}_k$. □

As the matrix $\tilde{\mathbf{B}}$ in the previous lemma is uniquely defined, this allows us to generalize the notion commonly known as Gram-Schmidt orthogonalization.

Definition 3. *Let $\mathbf{B} = \{\mathbf{b}_1, ..., \mathbf{b}_n\} \in H^n$. We call Gram-Schmidt orthogonalization (or GSO) of $\mathbf{B}$ and note $\tilde{\mathbf{B}}$ the unique set $\{\tilde{\mathbf{b}}_1, ..., \tilde{\mathbf{b}}_n\} \in H^n$ verifying one of the equivalent properties of Lemma 1.*

We now define the generalized Gram-Schmidt norm of a basis. Just as the usual Gram-Schmidt norm is very useful for Klein's sampler, ours will be useful for the ring version of Klein's sampler.

Definition 4 (Generalized Gram-Schmidt norm). *Let $\mathbf{B} = \{\mathbf{b}_1, ..., \mathbf{b}_n\} \in H^n$ be a basis, and $\tilde{\mathbf{B}} = \{\tilde{\mathbf{b}}_1, ..., \tilde{\mathbf{b}}_n\}$ its Gram-Schmidt orthogonalization. We call ($\mathbb{K}$-)Gram-Schmidt norm, and note $|\tilde{\mathbf{B}}|_{\mathbb{K}}$, the smallest value in $\mathbb{R}^+$ such that*

$$\forall i \in [\![1, n]\!], |\tilde{\mathbf{B}}|_{\mathbb{K}} \geqslant \|\tilde{\mathbf{b}}_i\|_{\mathbb{K}}$$

This definition allows us to subsume and encompass notions used in distinct Gaussian samplers.

For $\overline{\mathbb{K}} = \mathbb{R}$, it matches the usual definition of the Gram-Schmidt norm as in e.g. [ABB10], [DLP14]. And for $n = m = 1$, the Gram-Schmidt norm coincides with the definition of the largest singular value $s_1(\mathbf{B})$, used in [Pei10] to quantify the standard deviation of the output of Peikert's sampler. We briefly recall that the largest singular value $s_1(\mathbf{B})$ of a real matrix $\mathbb{R}^{n \times m}$ is defined by $s_1(\mathbf{B}) = \max_{\mathbf{x} \neq \mathbf{0}} \frac{\|\mathbf{xB}\|}{\|\mathbf{x}\|}$.

2.5 Gaussians

The norm defined at the end of Subsection 2.2 allows us to define Gaussians directly over $\overline{\mathbb{K}}^n$ (as opposed to $\mathbb{R}^n$ usually).

The Gaussian function $\rho : \overline{\mathbb{K}}^n \to (0, 1]$ is defined as follows:

$$\rho(\mathbf{x}) \stackrel{\Delta}{=} \exp\left(-\|\mathbf{x}\|^2/2\right)$$

If $\mathbf{B} \in \overline{\mathbb{K}}^{n \times n}$ is a nonsingular matrix, and $\mathbf{c} \in \overline{\mathbb{K}}^n$, then we extend this definition:

$$\rho_{\mathbf{B},\mathbf{c}}(\mathbf{x}) \stackrel{\Delta}{=} \rho\left((\mathbf{x} - \mathbf{c})\mathbf{B}^{-1}\right)$$

Let $\Sigma = \mathbf{B}^*\mathbf{B}$. Then for any orthogonal matrix $\mathbf{Q}$, $\mathbf{B}^*\mathbf{B} = (\mathbf{QB})^*(\mathbf{QB}) = \Sigma$. We therefore also use the notation $\rho_{\sqrt{\Sigma},\mathbf{c}}$ for $\rho_{\mathbf{B},\mathbf{c}}$, where $\sqrt{\Sigma}$ is an arbitrary matrix verifying $\sqrt{\Sigma}^*\sqrt{\Sigma} = \Sigma$. For a finite set $S \subseteq \overline{\mathbb{K}}^n$, we note $\rho_{\mathbf{B},\mathbf{c}}(S) \stackrel{\Delta}{=} \sum_{\mathbf{x}\in S} \rho_{\mathbf{B},\mathbf{c}}(\mathbf{x})$. For a lattice Λ, normalizing $\rho_{\sqrt{\Sigma},\mathbf{c}}$ by $\rho_{\sqrt{\Sigma},\mathbf{c}}(\Lambda)$ yields the discrete Gaussian distribution $D_{\Lambda,\sqrt{\Sigma},\mathbf{c}}$ over Λ. We also note $D_{\Lambda,\sigma,\mathbf{c}} \stackrel{\Delta}{=} D_{\Lambda,\sigma\cdot I_n,\mathbf{c}}$, $D_{\Lambda,\sqrt{\Sigma}} \stackrel{\Delta}{=} D_{\Lambda,\sqrt{\Sigma},0}$ and $D_\Lambda \stackrel{\Delta}{=} D_{\Lambda,1}$. Note that these definitions also hold if $\Lambda = \overline{\mathbb{K}}$, except that the Gaussian distribution is then no longer discrete but continuous.

We also recall the definition of the smoothing parameter (and of a scaled version better suited to our purposes), as well as two lemmas that will be very useful through this paper.

Definition 5 (Smoothing parameter [MR07]). *Let Λ be any n-dimensional lattice and $\epsilon > 0$, the smoothing parameter $\eta_\epsilon(\Lambda)$ is the smallest $s > 0$ such that $\rho_{1/s\sqrt{2\pi},\mathbf{0}}(\Lambda^* \setminus 0) \leqslant \epsilon$. We also define a scaled version $\eta'_\epsilon(\Lambda) \stackrel{\Delta}{=} \frac{1}{\sqrt{2\pi}}\eta_\epsilon(\Lambda)$.*

Lemma 2 (Corollary of [MR07], Lemma 4.4). *Let Λ be any n-dimensional lattice. Then for any $\epsilon \in (0,1), \sigma \geqslant \eta'_\epsilon(\Lambda)$ and $\mathbf{c} \in \mathbb{R}^n$, we have*

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

Lemma 3 (Special case of [MR07], Lemma 4.4). *For any $\epsilon \in (0,1)$:*

$$\eta'_\epsilon(\mathbb{Z}^n) \leqslant \frac{1}{\pi}\sqrt{\frac{1}{2}\log\left(2n\left(1+\frac{1}{\epsilon}\right)\right)}$$

2.6 The Kullback-Leibler Divergence

We now present the notion of Kullback-Leibler divergence (or KL Divergence) that works as a replacement of the more familiar notion of statistical distance. For Gaussian distributions ot oftens provides tighter proofs, therefore allowing sampling with a smaller variance, as already done in several works [DLP14,PDG14].

Definition 6 (Kullback-Leibler Divergence). *Let $\mathcal{P}$ and $\mathcal{Q}$ be two distributions over a common countable set Ω, and let $S \subset \Omega$ be the strict support of $\mathcal{P}$ ($\mathcal{P}(i) > 0$ iff $i \in S$). The Kullback-Leibler divergence, noted Δ_{KL} of $\mathcal{Q}$ from $\mathcal{P}$ is defined as:*

$$\Delta_{KL}(\mathcal{P}\|\mathcal{Q}) = \sum_{i\in S} \ln\left(\frac{\mathcal{P}(i)}{\mathcal{Q}(i)}\right)\mathcal{P}(i)$$

with the convention that $\ln(x/0) = +\infty$ for any $x > 0$.

To conclude a security argument using KL Divergence one can rely on the following lemma.

Lemma 4 (Bounding Success Probability Variations [DLP14,PDG14]). *Let $\mathcal{E}^{\mathcal{P}}$ be an algorithm making at most q queries to an oracle sampling from a distribution $\mathcal{P}$ and returning a bit. Let $\epsilon \geqslant 0$, and $\mathcal{Q}$ be a distribution such that $\Delta_{KL}(\mathcal{P}\|\mathcal{Q}) \leq \epsilon$. Let x (resp. y) denote the probability that $\mathcal{E}^{\mathcal{P}}$ (resp. $\mathcal{E}^{\mathcal{Q}}$) outputs 1. Then, $|x-y| \leq \sqrt{q\epsilon/2}$.*

Finally, to bound the KL Divergence, one can apply the following Lemma.

Lemma 5 (KL Divergence for bounded relative error [DLP14,PDG14]). *Let $\mathcal{P}$ and $\mathcal{Q}$ be two distributions of same countable support. Assume that for any $i \in S$, there exists some $\delta(i) \in (0, 1/4)$ such that we have the relative error bound $|\mathcal{P}(i) - \mathcal{Q}(i)| \leq \delta(i)\mathcal{P}(i)$. Then*

$$\Delta_{KL}(\mathcal{P}\|\mathcal{Q}) \leq 2\sum_{i\in S}\delta(i)^2\mathcal{P}(i).$$

3 Ring Variants of Klein's and Peikert's Samplers

In this section, we present ring variants of Klein's and Peikert's samplers.

3.1 A Ring Variant of Klein's Sampler

Here we present a ring generalization of Klein's algorithm, where $\mathbb{Z}$ (resp. $\mathbb{R}$) is replaced by $\mathcal{R}$ (resp. $\overline{\mathbb{K}}$). To get an intuition of why this could be faster than Klein's sampler, see that each output is of the form $\mathbf{v} = \sum_i z_i\mathbf{b}_i \in \mathcal{R}^m$, where the $z_i \in \mathcal{R}$. Then this algorithm samples an entire z_i at each step, whereas the original algorithm from Klein can only sample one coordinate of one z_i at each step.

Algorithm 1 `Ring_Klein`$(\mathcal{R}, \mathbf{B}, \tilde{\mathbf{B}}, \sigma, \mathbf{c})$

Require: Basis $\mathbf{B} = \{\mathbf{b}_1, ..., \mathbf{b}_n\} \in \mathcal{R}^{n\times m}$, its GSO $\tilde{\mathbf{B}} = \{\tilde{\mathbf{b}}_1, ..., \tilde{\mathbf{b}}_n\} \in \overline{\mathbb{K}}^{n\times m}$, $\sigma \in \overline{\mathbb{K}}^+$, target $\mathbf{c} \in \overline{\mathbb{K}}^m$

Ensure: $\mathbf{v}$ sampled in a distribution close to $\mathcal{D}_{Span_{\mathcal{R}}(\mathbf{B}),\sigma,\mathbf{c}}$

$\mathbf{c}_n \leftarrow \mathbf{c}$ $\quad \in \overline{\mathbb{K}}^m$
$\mathbf{v}_n \leftarrow \mathbf{0}$ $\quad \in \mathcal{R}^m$
for $i \leftarrow n, ..., 1$ **do**
  $d_i \leftarrow \langle \mathbf{c}_i, \tilde{\mathbf{b}}_i\rangle_{\mathbb{K}}/\|\tilde{\mathbf{b}}_i\|_{\mathbb{K}}^2$ $\quad \in \overline{\mathbb{K}}$
  $\Sigma_i \leftarrow \sigma^2/\|\tilde{\mathbf{b}}_i\|_{\mathbb{K}}^2$ $\quad \in \overline{\mathbb{K}}$
  $z_i \leftarrow$ `SampleR`$(\mathcal{R}, \Sigma_i, d_i)$ $\quad \in \mathcal{R}$
  $\mathbf{c}_{i-1} \leftarrow \mathbf{c}_i - z_i\mathbf{b}_i$ $\quad \in \overline{\mathbb{K}}^m$
  $\mathbf{v}_{i-1} \leftarrow \mathbf{v}_i + z_i\mathbf{b}_i$ $\quad \in \mathcal{R}^m$
end for
return $\mathbf{v}_0$

In Algorithm 1, `SampleR` is assumed to be a perfect discrete Gaussian sampler over $\mathcal{R}$: given a ring $\mathcal{R}$, a covariance $\Sigma \in \overline{\mathbb{K}^+}$ and a center $d \in \overline{\mathbb{K}}$, we assume

$\texttt{SampleR}(\mathcal{R}, \Sigma, d) = D_{\mathcal{R},\sqrt{\Sigma},d}$. This mirrors the sampler in [GPV08], which uses a discrete Gaussian sampler over $\mathbb{Z}$ as an oracle. Here $\mathbb{Z}$ is replaced by $\mathcal{R}$, which of course raises practicality issues, but these questions will be addressed in Section 4.

The rest of this subsection is devoted to analyzing the correctness of Algorithm 1. Since the lemmas and their proofs are mostly identical to similar counterparts in [GPV08,DLP14], readers interested only in the practical applications may wish to acknowledge Theorem 1 and then skip to the next section.

Lemma 6. *For any input* $(\mathcal{R}, \mathbf{B}, \tilde{\mathbf{B}}, \sigma, \mathbf{c})$ *and output* $\mathbf{v} = \sum_i \hat{z}_i \mathbf{b}_i \in Span_{\mathcal{R}}(\mathbf{B})$ *of* `Ring_Klein`*,*

$$\mathbf{v} - \mathbf{c} = \sum_i (\hat{z}_i - d_i)\tilde{\mathbf{b}}_i$$

where the values $c_i, \hat{z}_i$ *are as in* $\texttt{Ring\_Klein}(\mathcal{R}, \mathbf{B}, \tilde{\mathbf{B}}, \sigma, \mathbf{c}) \to \mathbf{v}$*.*

Proof. The proof is identical to the proof of Lemma 4.4 in [GPV08]. The only difference is that $\mathbb{Z}$ (resp. $\mathbb{R}$) is replaced with $\mathcal{R}$ (resp. $\overline{\mathbb{K}}$), and therefore $\Lambda(\mathbf{B})$ (resp. $Span_{\mathbb{R}}(\mathbf{b}_1, ..., \mathbf{b}_k)$) has to be replaced with $Span_{\mathcal{R}}(\mathbf{B})$ (resp. $Span_{\overline{\mathbb{K}}}(\mathbf{b}_1, ..., \mathbf{b}_k)$). □

Lemma 7. *For any input* $(\mathcal{R}, \mathbf{B}, \tilde{\mathbf{B}}, \sigma, \mathbf{c})$ *and output* $\mathbf{v} = \sum_i \hat{z}_i \mathbf{b}_i \in Span_{\mathcal{R}}(\mathbf{B})$ *of* `Ring_Klein`*, the probability that* $\mathbf{v}$ *is output is exactly*

$$\rho_{\sigma,\mathbf{c}}(\mathbf{v}) \cdot \prod_{1 \leqslant i \leqslant n} \frac{1}{\rho_{\sqrt{\Sigma_i},d_i}(\mathcal{R})}$$

Proof. For each i, the probability that $z_i = \hat{z}_i$ (conditioned on $z_j = \hat{z}_j$ for all $j > i$) is exactly $D_{\mathcal{R},\sqrt{\Sigma_i},d_i}(\hat{z}_i)$. Therefore the probability that $\mathbf{v}$ is output is

$$\prod_{1 \leqslant i \leqslant n} D_{\mathcal{R},\sqrt{\Sigma_i},d_i}(\hat{z}_i) = \frac{\prod_{1 \leqslant i \leqslant n} \rho_{\sqrt{\Sigma_i},d_i}(\hat{z}_i)}{\prod_{1 \leqslant i \leqslant n} \rho_{\sqrt{\Sigma_i},d_i}(\mathcal{R})}$$

In the expression above, the numerator is

$$\prod_{1 \leqslant i \leqslant n} \rho_{\sqrt{\Sigma_i},d_i}(\hat{z}_i) = \prod_{1 \leqslant i \leqslant n} \rho_\sigma\left((\hat{z}_i - d_i)\|\tilde{\mathbf{b}}_i\|_{\mathbb{K}}\right) = \rho_\sigma\left(\Sigma_i(\hat{z}_i - d_i)\tilde{\mathbf{b}}_i\right) = \rho_{\sigma,\mathbf{c}}(\mathbf{v})$$

The first equality comes from the fact that $\Sigma_i = \sigma^2/\|\tilde{\mathbf{b}}_i\|^2_{\mathbb{K}}$, the second one from the pairwise orthogonality of the $\tilde{\mathbf{b}}_i$'s, and the last one from Lemma 6. □

Theorem 1. *Let* $\mathbf{B} = \{\mathbf{b}_1, ..., \mathbf{b}_n\} \in \mathcal{R}^{n \times m}$ *be a* $\mathcal{R}$*-basis,* $\tilde{\mathbf{B}} = \{\tilde{\mathbf{b}}_1, ..., \tilde{\mathbf{b}}_n\} \in \overline{\mathbb{K}}^{n \times m}$ *its GSO,* $\mathbf{c} \in \mathcal{R}^m$*. Let* $\epsilon \in (0, \frac{1}{2n})$ *and* $\sigma \in \overline{\mathbb{K}^+}$ *such that* $\sigma \geqslant \eta'_\epsilon(\mathcal{R}) \cdot \|\tilde{\mathbf{B}}\|$*. The statistical distance (resp. KL Divergence) between* $D_{Span_{\mathcal{R}}(\mathbf{B}),\sigma,\mathbf{c}}$ *and the output distribution of* $\texttt{Ring\_Klein}(\mathcal{R}, \mathbf{B}, \tilde{\mathbf{B}}, \sigma, \mathbf{c})$ *is upper bounded by* $2n\epsilon$ *(resp.* $2\left(1 - \left(\frac{1+\epsilon}{1-\epsilon}\right)^n\right)^2 \approx 8n^2\epsilon^2$*).*

Proof. The proof is almost to identical to the proof of Theorem 2 in [DLP14].

Let $\mathcal{P} = D_{Span_{\mathcal{R}}(\mathbf{B}),\sigma,\mathbf{c}}$, and $\mathcal{Q}$ be the output distribution of $\texttt{Ring\_Klein}(\mathcal{R}, \mathbf{B}, \tilde{\mathbf{B}}, \sigma, \mathbf{c})$. By assumption, $\sqrt{\Sigma_i} \geqslant \eta'_\epsilon(\mathcal{R})$, so from Lemma 2 we can infer that $\rho_{\sqrt{\Sigma_i},d_i}(\mathcal{R}) \in \left[\frac{1-\epsilon}{1+\epsilon}, 1\right] \cdot \rho_{\sqrt{\Sigma_i}}(\mathcal{R})$. Applying Lemma 7 then gives us a relative error between $\mathcal{P}$ and $\mathcal{Q}$:

$$\forall \mathbf{v} \in Span_{\mathcal{R}}(\mathbf{B}), |\mathcal{P}(\mathbf{v}) - \mathcal{Q}(\mathbf{v})| \leqslant \left(1 - \left(\frac{1+\epsilon}{1-\epsilon}\right)^n\right) \mathcal{P}(\mathbf{v})$$

At this point we get the statistical distance (resp. KL Divergence) using a straightforward computation (resp. 5). □

3.2 A Ring Variant of Peikert's Sampler

In this subsection, we present and analyze a ring variant of Algorithm 1 from [Pei10]. Although we do not introduce new techniques to allow such a transformation (which was started in [Pei10] and completed in [DN12]), we provide a complete description of a ring variant and give an analysis of its divergence from a perfect sampler.

For this step, we require a $\mathbb{Z}$-basis of the ring $\mathcal{R}$, that we denote by $(e_1, \ldots, e_d)$. For example, if $\mathcal{R} = \mathbb{Z}[x]/(f(x))$ one may take a power-basis $1, x, \ldots, x^{\deg(f)-1}$, but other choices may lead to better results (see [?]). One note that for the variance $\Sigma_e = \sum e_i \bar{e}_i \in \mathbb{K}^+$, it is essentially trivial to sample a distribution close to $D_{\mathcal{R},\eta\cdot\sqrt{\Sigma_e},c_i}$ (where $\eta = \eta_\epsilon(\mathbb{Z}) \in \mathbb{R}$), by computing $\sum_{1\leqslant i\leqslant d} a_i e_i$ where each a_i is drawn from $D_{\mathbb{Z},\eta,c_i}$ and $c = \sum_{1\leqslant i\leqslant d} c_i e_i$.

Algorithm 2 $\texttt{Ring\_Peikert}(\mathcal{R}, \Sigma, c)$

Require: A variance $\Sigma \in \overline{\mathbb{K}^+}$, a target $c \in \overline{\mathbb{K}}$, a precomputed value $b \in \overline{\mathbb{K}}$ such that $\Sigma_e(b\bar{b} + \eta^2) = \Sigma$

Ensure: z sampled according to $D_{\mathcal{R},\sigma,c}$

1: $p \leftarrow b \cdot D_{\overline{\mathbb{K}},\sqrt{\Sigma_e}}$ $\quad \in \overline{\mathbb{K}}$
2: $z \leftarrow D_{\mathcal{R},\eta\sqrt{\Sigma_e},c+p}$ $\quad \in \mathcal{R}$
3: Return z

Just like before, the rest of this subsection is dedicated to proving the correctness of Algorithm 2. Once again, readers interested in practical applications may wish to acknowledge Theorem 3 and skip to Section 4.

We first recall a theorem (in a simplified version) from [Pei10] which will help us analyze Algorithm 2. Our version of this theorem also adds a bound on the KL Divergence between the output of Algorithm 2 and the desired distribution.

Theorem 2 (Theorem 3.1 of [Pei10], continuous case, simplified). *Let $\mathcal{R} = \mathbb{Z}^m$, $\Sigma_1, \Sigma_2 > \mathbf{0}$ be positive definite matrices, with $\Sigma = \Sigma_1 + \Sigma_2 > \mathbf{0}$. Let*

Λ_1 be a lattice such that $\sqrt{\Sigma_1} \geqslant \eta'_\epsilon(\Lambda_1)$ for some positive $\epsilon \leqslant 1/2$, and let $\mathbf{c}$ be arbitrary. Consider the following probabilistic experiment:

$$\textit{Choose } \mathbf{x}_2 \sim D_{\sqrt{\Sigma_2}}, \textit{ then choose } \mathbf{x}_1 \sim D_{\Lambda_1+\mathbf{c}_1,\sqrt{\Sigma_1},\mathbf{x}_2}.$$

The statistical distance (resp. KL Divergence) between the distribution of $\mathbf{x}_1$ and $D_{\Lambda_1+\mathbf{c}_1,\sqrt{\Sigma}}$ is upper bounded by $\frac{\epsilon}{1-\epsilon} \approx \epsilon$ (resp. $2\left(\frac{2\epsilon}{1-\epsilon}\right)^2 \approx 8\epsilon^2$).

Proof. We only improve a specific point of the original theorem: namely, adding a bound on the KL Divergence and very slightly enhancing the bound on the statistical distance. Therefore, we focus on the part we improve and invite the reader interested in more details to read the complete proof in [Pei10].

Let $\bar{\mathbf{x}}_1 \in \Lambda_1 + \mathbf{c}_1$. From the proof of theorem 3.1 of [Pei10], we have $Pr[\bar{\mathbf{x}}_1 = \mathbf{x}_1] \propto \rho_{\sqrt{\Sigma}}(\bar{\mathbf{x}}_1) \cdot \left[1, \frac{1+\epsilon}{1-\epsilon}\right]$ (and not $\left[\frac{1-\epsilon}{1+\epsilon}, \frac{1+\epsilon}{1-\epsilon}\right]$ like in the original proof, since in the continuous case we need only in invoke Lemma 2.4 of [Pei10] once). It follows that $Pr[\bar{\mathbf{x}}_1 = \mathbf{x}_1] \in \left[\frac{1-\epsilon}{1+\epsilon}, \frac{1+\epsilon}{1-\epsilon}\right] \cdot D_{\Lambda_1+\mathbf{c}_1,\sqrt{\Sigma}}(\bar{\mathbf{x}}_1)$. A straightforward computation (resp. Lemma 5) yields the bound on the statistical distance (resp. the KL Divergence). □

We now use Theorem 2 to prove that Algorithm 2 is indeed correct by being a ring instanciation of Algorithm 1 from [Pei10].

Theorem 3. *Let $\epsilon \in (0, \frac{1}{2})$, $b \in \overline{\mathbb{K}}$ such that $\Sigma_e(b\bar{b} + \eta^2) = \Sigma$ and $\eta \geqslant \eta'_\epsilon(\mathcal{R})$. The statistical distance (resp. KL Divergence) between the output of* `Ring_Peikert`$(\mathcal{R}, \Sigma, \mathbf{c})$ *and $D_{\mathcal{R},\sqrt{\Sigma},\mathbf{c}}$ is upper bounded by $\approx 2\epsilon$ (resp. $\approx 8\epsilon^2$).*

Proof. Our goal is to prove that we fulfill the conditions necessary to apply theorem 2.

Let us take $\Lambda_1 := \mathcal{R}$, $\Sigma_1 := \eta^2\Sigma_e$, $\Sigma_2 := b\bar{b}\Sigma_e$, $\mathbf{c}_1 := c$, $\mathbf{x}_2 := -p$ and $\mathbf{x}_1 := z - c$. One can check that all the conditions are verified, in particular $\mathbf{x}_2 \sim D_{\sqrt{\Sigma_2}}$ and $\mathbf{x}_1 \sim D_{\Lambda_1+\mathbf{c}_1,\sqrt{\Sigma_1},\mathbf{x}_2}$, so this allows us to apply theorem 2 and assert that the distribution of $\mathbf{x}_1$ is close to $D_{\mathcal{R}+c,\sqrt{\Sigma}}$ as specified by the theorem: the bijective transformation $z := c + \mathbf{x}_1$ then allows us to conclude. □

4 Hybrid Algorithms for Sampling and Reduction

4.1 A Hybrid Sampler

In this section, we show that the two algorithms presented in Section 3 can be efficiently combined: more precisely, using `Ring_Peikert` as a subroutine of `Ring_Klein` will give us a new discrete Gaussian sampler.

Definition 7. *Using the notations from Section 3, we note* `Hybrid_Sampler`$(\mathcal{R}, \mathbf{B}, \tilde{\mathbf{B}}, \sigma, \mathbf{c})$ *the algorithm* `Ring_Klein`$(\mathcal{R}, \mathbf{B}, \tilde{\mathbf{B}}, \sigma, \mathbf{c})$ *where instead of calling a perfect sampler* `SampleR`$(\mathcal{R}, \Sigma_i, d_i)$ *in step 6, we call* `Ring_Peikert`$(\mathcal{R}, \Sigma_i, d_i)$.

We now show that for carefully chosen parameters, the output distribution of `Hybrid_Sampler` is close (in the sense of either the statistical distance or the KL Divergence) to a perfect discrete Gaussian. Given the nature of this sampler (it combines `Ring_Klein` and `Ring_Peikert`), we have to take into account the divergence of both `Ring_Klein` and `Ring_Peikert` from a "perfect behavior".

It is tempting to first quantify the divergence between `Hybrid_Sampler` and `Ring_Klein`, and then use the triangle inequality along with Theorem 1 to get the divergence between `Hybrid_Sampler` and a perfect discrete Gaussian. This approach works in the case of statistical distance but is useless for KL Divergence since the latter does not verify the triangle inequality. Instead, we directly compute the statistical distance (resp. KL Divergence) between both distributions.

Theorem 4. *Let $\mathbf{B} = \{\mathbf{b}_1, ..., \mathbf{b}_n\} \in \mathcal{R}^{m\times n}$ be a $\mathcal{R}$-basis, $\tilde{\mathbf{B}} = \{\tilde{\mathbf{b}}_1, ..., \tilde{\mathbf{b}}_n\} \in \overline{\mathbb{K}}^{m\times n}$ its GSO and $\mathbf{c} \in \mathcal{R}^m$. Let $\epsilon \in (0, \frac{1}{6n})$ and $\sigma \in \mathbb{K}^+$ such that $\sigma \geqslant \eta'_\epsilon(\mathcal{R}) \cdot \|\tilde{\mathbf{B}}\|$.*

The statistical distance (resp. KL Divergence) between $D_{Span_\mathcal{R}(\mathbf{B}),\sigma,\mathbf{c}}$ and the output distribution of $\texttt{Hybrid\_Sampler}(\mathcal{R}, \mathbf{B}, \tilde{\mathbf{B}}, \sigma, \mathbf{c})$ is upper bounded by $\frac{1}{2}((\frac{1+\epsilon}{1-\epsilon})^{3n} - 1) \approx 3n\epsilon$ (resp. $2((\frac{1+\epsilon}{1-\epsilon})^{3n} - 1)^2 \approx 72n^2\epsilon^2$).

Proof. This proof reprises elements from the proofs of Theorems 1 and 2.

Let $\mathcal{Q}$ be the output distribution of $\texttt{Hybrid\_Sampler}(\mathcal{R}, \mathbf{B}, \tilde{\mathbf{B}}, \sigma, \mathbf{c})$ and $\mathcal{P} = D_{Span_\mathcal{R}(\mathbf{B}),\sigma,\mathbf{c}}$. Divergence between $\mathcal{P}$ and $\mathcal{Q}$ come from both the use of `Ring_Klein` and `Ring_Peikert`. Theorem 1 (resp. 3) quantifies the difference between the output of `Ring_Klein` (resp. `Ring_Peikert`) and a perfect Gaussian, upon the condition $\sigma \geqslant \eta'_\epsilon(\mathcal{R}) \cdot \|\tilde{\mathbf{B}}\|$ (resp. $\sqrt{\Sigma_1} \geqslant \eta'_\epsilon(\Lambda_1)$). Coincidentally, in this case, the two conditions end up being exactly the same: the ϵ mentioned in Theorems 1 and 3 are actually the same here.

Let $\mathbf{v} = \sum_i \hat{z}_i \mathbf{b}_i \in Span_\mathcal{R}(\mathbf{B})$. For any i, let $\mathcal{Q}_i$ be the output distribution of $\texttt{Ring\_Peikert}(\mathcal{R}, \Sigma_i, d_i)$, where the Σ_i, d_i are as in $\mathbf{v} \leftarrow \mathcal{Q}$. For each i, $\mathcal{Q}_i(\hat{z}_i) \in \left[\frac{1-\epsilon}{1+\epsilon}, \frac{1+\epsilon}{1-\epsilon}\right] D_{\mathcal{R},\sqrt{\Sigma_i},d_i}$, as detailed in the proof of Theorem 2. Therefore

$$\begin{aligned}
\mathcal{Q}(\mathbf{v}) &= \textstyle\prod_{1\leqslant i\leqslant n} \mathcal{Q}_i(\hat{z}_i) \\
&\in \left[(\tfrac{1-\epsilon}{1+\epsilon})^n, (\tfrac{1+\epsilon}{1-\epsilon})^n\right] \textstyle\prod_i D_{\mathcal{R},\sqrt{\Sigma_i},d_i}(\hat{z}_i) \\
&\in \left[(\tfrac{1-\epsilon}{1+\epsilon})^n, (\tfrac{1+\epsilon}{1-\epsilon})^n\right] \frac{\rho_{\sigma,\mathbf{c}}(\mathbf{v})}{\prod_i \rho_{\sqrt{\Sigma_i},d_i}(\mathcal{R})} \\
&\in \left[(\tfrac{1-\epsilon}{1+\epsilon})^n, (\tfrac{1+\epsilon}{1-\epsilon})^n\right] \left[1, (\tfrac{1+\epsilon}{1-\epsilon})^n\right] \frac{\rho_{\sigma,\mathbf{c}}(\mathbf{v})}{\prod_i \rho_{\sqrt{\Sigma_i}}(\mathcal{R})}
\end{aligned}$$

Where the second equality comes from the fact that for each i, $\mathcal{Q}_i(\hat{z}_i) \in \left[\frac{1-\epsilon}{1+\epsilon}, \frac{1+\epsilon}{1-\epsilon}\right] D_{\mathcal{R},\sqrt{\Sigma_i},d_i}$, the third one from Lemma 7 and the fourth from Lemma 2.

Let $\alpha = \frac{\rho_{\sigma,\mathbf{c}}(Span_\mathcal{R}(\mathbf{B}))}{\prod_i \rho_{\sqrt{\Sigma_i}}(\mathcal{R})}$. We can then write $\mathcal{Q}(\mathbf{v}) \in \alpha \left[(\frac{1-\epsilon}{1+\epsilon})^n, (\frac{1+\epsilon}{1-\epsilon})^{2n}\right] \mathcal{P}(\mathbf{v})$. Summing $\mathcal{Q}(\mathbf{v})$ over $Span_\mathcal{R}(\mathbf{B})$ yields

$$1 \in \alpha \left[\left(\frac{1-\epsilon}{1+\epsilon}\right)^n, \left(\frac{1+\epsilon}{1-\epsilon}\right)^{2n}\right] \overbrace{\sum_{\mathbf{v}\in Span_\mathcal{R}(\mathbf{B})} \mathcal{P}(\mathbf{v})}^{=1}$$

This implies that $\alpha \in \left[(\frac{1-\epsilon}{1+\epsilon})^{2n}, (\frac{1+\epsilon}{1-\epsilon})^{n}\right]$, so we get a relative error bound between $\mathcal{P}$ and $\mathcal{Q}$:

$$|\mathcal{Q}(\mathbf{v}) - \mathcal{P}(\mathbf{v})| \leqslant \left(\left(\frac{1+\epsilon}{1-\epsilon}\right)^{3n} - 1\right)\mathcal{P}(\mathbf{v})$$

We can then conclude using a straightforward computation (resp. Lemma 5). □

The bound on the statistical distance (resp. KL Divergence) can then be used to assert the security of the scheme following a standard argument (resp. Lemma 4).

4.2 A Hybrid Babai's Algorithm

One could see Babai's rounding (resp. nearest plane) algorithm as a specific instanciation of Klein's (resp. Peikert's) sampler for a standard deviation $\sigma = 0$. While Babai's algorithms [Bab86] were invented long before the aforementioned samplers and do not serve the same purpose, it is nevertheless correct to view them like this, under the convention that a gaussian with standard deviation 0 behaves like an exact rounding. Therefore, the method used to create a hybrid sampler from Klein's and Peikert's sampler can be used to create a hybrid approximation algorithm for the closest vector problem from Babai's approximation algorithms.

Definition 8. *Let $\mathbf{B} \in \mathcal{R}^{m\times n}$ be a $\mathcal{R}$-basis, $\tilde{\mathbf{B}}$ = its GSO and $\mathbf{c} \in \mathcal{R}^m$. We define*

$$\texttt{Hybrid\_Babai}(\mathcal{R}, \mathbf{B}, \tilde{\mathbf{B}}, \cdot) \triangleq \texttt{Hybrid\_Sampler}(\mathcal{R}, \mathbf{B}, \tilde{\mathbf{B}}, 0, \cdot)$$

We now give a bound on the output of the algorithm. The proof can be done either by generalizing the proof of Babai's nearest plane algorithm (see e.g. [Gal12, chapter 18]) or by reprising the proof of Theorem 4.

Lemma 8. *For any $\mathbf{c} \in \overline{\mathbb{K}}^m$, $\texttt{Hybrid\_Babai}(\mathcal{R}, \mathbf{B}, \tilde{\mathbf{B}}, \mathbf{c})$ outputs a point $\mathbf{v}$ verifying:*

$$\mathbf{v} - \mathbf{c} = \sum_{1\leqslant i\leqslant n} \varepsilon_i \tilde{\mathbf{b}}_i$$

where $\varepsilon_i \in \overline{\mathbb{K}}$ are such that, for all $i \leq n$, $-\frac{1}{2} \leqslant \varepsilon_i \leqslant \frac{1}{2}$ (as inequalities of $\overline{\mathbb{K}}$).

5 A Practical Viewpoint: Gaussian Sampling over NTRU-like Lattices

While proper use of lattice trapdoors [GPV08] is quite recent, the design of compact lattices with quite a good trapdoor comes from the NTRU-cryptosystems

[HPS98,HHGP+03]. Unlike more recent constructions [MP12], those lattices are not necessarily uniformly random (unless one makes impractical choices of parameters [SS11]), in which case they may not benefit from theoretical worst-case hardness [Ajt96,MR07]. Nevertheless, when it comes to practical instantiation and concrete security, this type of construction remains to date the most efficient one.

In this section, we assess the practical interest of our hybrid sampler by comparing it with Klein's and Peikert's samplers. To do this, we instantiate the Full-Domain Hash signature scheme proposed in [GPV08] over NTRU lattices. For Klein's sampler, an analysis of this scheme has already been done in [DLP14].

This last section is organized as follows: first we recall what NTRU lattices are, then we explain how we will compare our Hybrid sampler with Klein's and Peikert's and then we will summarize our results in Table 1. Subsections 5.1, 5.2 and 5.3 detail briefly how the results for each sampler were obtained, and Subsection 5.4 heuristically explains the quality gap between the Hybrid and Peikert's samplers.

Definition 9. *Let Φ_m be the m-th cyclotomic polynomial, let ω_m be an arbitrary root of Φ_m, and $\Omega_m = \{\omega_m^k, k \in \mathbb{Z}_m^\times\}$ be the set of complex roots of Φ_m. We also note $\varphi(m)$ Euler's totient function evaluated on m, and $\mathbf{\Phi}_m$ the diagonal $\varphi(m) \times \varphi(m)$ matrix whose diagonal coefficients are the elements of Ω_m.*

Note that we sligthly deviate from the originally chosen ring of NTRU cryptosystems, and prefer the now standard choice of cyclotomic rings.

In the rest of the section, $m \in \mathbb{N}^*$ will be a power of two, $q \in \mathbb{N}^*$, $N = \varphi(m)$, $\mathcal{Z} = \mathbb{Z}[x]/(\phi_m(x)) = \mathbb{Z}[x]/(x^N + 1)$ and $\mathcal{K} = \mathbb{Q}[x]/(\phi_m(x))$.

Definition 10 (NTRU lattice). *Let $f, g, F, G \in \mathbb{Z}_N[x]$ such that*

$$fG - gF = q \mod (x^N + 1) \tag{1}$$

The NTRU lattice generated by f, g, F, G is the lattice generated by the rows of the block matrix

$$\mathbf{B}_{f,g,F,G} = \left[\begin{array}{c|c} \mathcal{A}(f) & \mathcal{A}(g) \\ \hline \mathcal{A}(F) & \mathcal{A}(G) \end{array}\right]$$

Where $\mathcal{A}(p)$ is the $N \times N$ matrix which i-th row is the coefficients of $x^{i-1} \cdot p(x) \mod (x^N + 1)$.

For any fixed (f, g) and distinct (F_1, G_1), (F_2, G_2) verifying equation 1, $\Lambda(\mathbf{B}_{f,g,F_1,G_1}) = \Lambda(\mathbf{B}_{f,g,F_2,G_2})$ so in the rest of the paper we simply assume that (F, G) is reduced with respect to $Span_{\mathcal{R}}((f, g))$ using Babai's rounding algorithm and note $\mathbf{B}_{f,g} \stackrel{\Delta}{=} \mathbf{B}_{f,g,F,G}$.

Depending if we take $\mathbb{Z}$, $\mathcal{Z}$ or $\mathcal{Z}^{2\times 2}$ as our base ring $\mathcal{R}$, we can do Gaussian sampling in three different ways:

1. Taking $\mathcal{R} = \mathbb{Z}$, our hybrid algorithm is simply the original Klein's sampler. This approach is developed in Subsection 5.1.

2. Taking $\mathcal{R} = \mathcal{Z}$, we get a sampler that is strictly different from Klein's and Peikert's sampler. Indeed, $\Lambda(\mathbf{B}_{f,g,F,G})$ can then be seen as a $\mathcal{R}$-module of rank 2. This approach is developed in Subsection 5.2.
3. Taking $\mathcal{R} = \mathcal{Z}^{2\times 2}$, one uses Peikert's original algorithm[1]. This approach is developed in Subsection 5.3.

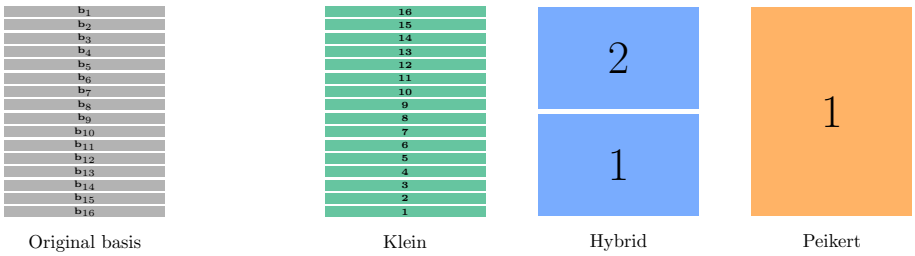

Fig. 2. Each sampler outputs a vector $v = \sum_i z_i \mathbf{b}_i$. The z_i's are computed one by one for Klein's sampler (the figure gives the order in which they are computed), all at the same time for Peikert's, and N at a time for our instantiation of the Hybrid sampler over NTRU lattices.

To compare the three approaches, we fix the parameters $N = 512, q = 2^{20}$ and try to find the NTRU bases yielding the most secure Full-Domain Hash signature scheme for each sampler. We already know that Klein's will sample with the smallest standard deviation and Peikert's with the largest, due to equation 2.

$$\sqrt{q} = \det(\mathbf{B}_{f,g})^{\frac{1}{2N}} \leqslant |\tilde{\mathbf{B}}_{f,g}|_{\mathbb{R}} \leqslant |\tilde{\mathbf{B}}_{f,g}|_{\mathcal{K}} \leqslant s_1(\mathbf{B}_{f,g}) \tag{2}$$

What we want is to quantify these standard variations, which is equivalent to computing $|\tilde{\mathbf{B}}_{f,g}|_{\mathbb{R}}$ for Klein's sampler, $|\tilde{\mathbf{B}}_{f,g}|_{\mathcal{K}}$ for the Hybrid and $s_1(\mathbf{B}_{f,g})$ for Peikert's.

Once it is done, we evaluate the security of the signatures schemes by determining the hardness of the underlying lattice problem. This hardness can be quantified using the root Hermite factor introduced in [GN08]. Let $\mathbf{v} \in \Lambda$ be a vector that one is looking for in a n-dimensional lattice. If $\|\mathbf{v}\| > det(\Lambda)^{1/n}$, then the associated root Hermite factor γ verifies:

$$\gamma^n \leq \frac{\|\mathbf{v}\|}{det(\Lambda)^{1/n}} \tag{3}$$

If $\mathbf{v}$ is an unusually short vector planted in an NTRU lattice, then according to experiments in [DDLL13], γ verifies:

$$\frac{\sqrt{n/(2\pi e)} \cdot det(\Lambda)^{1/n}}{\|\mathbf{v}\|} = .4\gamma^n. \tag{4}$$

[1] To formally subsume that case with our algorithm, we would need to generalize our description to non-commutative rings.

Using the works of [GN08,CN11], one can get a very rough estimate of the hardness of the lattice problem based on the value of γ (unfortunately, lattice cryptanalysis literature on this subject is sparse so we cannot get much more than just a rough estimate).

The results of our experiments are summarized in the Table 1.

Table 1. Comparison of the three approaches for NTRU lattices of fixed dimension $2N = 1024$ and modulus $q = 2^{20}$

Sampler	**Klein**	**Hybrid**	**Peikert**
Security level	**192**	**160**	**120**
Root Hermite factor γ	1.0042	1.0049	1.0060
Ring $\mathcal{R}$	$\mathbb{Z}$	$\mathcal{Z}$	$\mathcal{Z}^{2\times 2}$
$\text{rank}_{\mathcal{R}}(\Lambda)$	$2N$	2	1
Form of $\mathbf{B}$	$\mathcal{R}^{2N\times 2N}$	$\mathcal{R}^{2\times 2}$	$\mathcal{R}^{1\times 1}$
Running time	$\tilde{O}(N^2)$	$\tilde{O}(N)$	$\tilde{O}(N)$
Smallest $\lvert\tilde{\mathbf{B}}\rvert_{\mathbb{K}}$ attained	$1.17\sqrt{q}$	$2.5\sqrt{q}$	$8\sqrt{q}$

5.1 Klein's Sampler ($\mathcal{R} = \mathbb{Z}, \mathbb{K} = \mathbb{Q}$)

Klein's Sampler has been extensively studied in [GPV08,DN12,DLP14], We recall its behavior in the case of NTRU lattices as stated in [DLP14].

Lemma 9 ([DLP14], Lemmas 2 and 3). *Let $\mathbf{B}_{f,g}$ be as defined in Definition 10. The classical Gram-Schmidt norm of $\mathbf{B}_{f,g}$ is:*

$$|\tilde{\mathbf{B}}_{f,g}|_{\mathbb{R}} = \max\left(\|(f,g)\|, \left\|\left(\frac{q\bar{g}}{f\bar{f}+g\bar{g}}, \frac{-q\bar{f}}{f\bar{f}+g\bar{g}}\right)\right\|\right)$$

In addition, [DLP14, full version] provides experiments and a heuristic which both confirm that one can sample NTRU bases with a Gram-Schmidt norm as small as $1.17\sqrt{q}$.

5.2 Hybrid Sampler ($\mathcal{R} = \mathcal{Z}, \mathbb{K} = \mathcal{K}$)

In the case where $\mathcal{R} = \mathcal{Z}$, the vectors $\mathbf{f} = (f, g), \mathbf{F} = (F, G) \in \mathcal{R}^2$ then generate $\Lambda(\mathbf{B}_{f,g})$ as a $\mathcal{R}$-module of rank 2: $\Lambda(\mathbf{B}_{f,g}) = Span_{\mathcal{R}}(\mathbf{f}, \mathbf{F})$.[2] We then get

Lemma 10. *Let $\mathbf{B}_{f,g}$ be as defined in Definition 10. The $\mathcal{K}$-Gram-Schmidt norm of $\mathbf{B}_{f,g}$ is:*

$$|\tilde{\mathbf{B}}_{f,g}|_{\mathcal{K}} = \max_{\omega\in\Omega_m} \max\left(D(\omega), \frac{q}{D(\omega)}\right)$$

[2] Note that in this setting, the equation 1 becomes $\det[\mathbf{f}, \mathbf{F}] = q \in \mathcal{R}$.

Where $D(x) \stackrel{\Delta}{=} \|(f,g)\|_{\mathcal{K}}(x) (= \sqrt{|f(x)|^2 + |g(x)|^2}$ *when* $x \in \Omega_m$*).*

Proof. $\mathbf{B} = \{\mathbf{f} = (f,g), \mathbf{F} = (F,G)\}$ is the $\mathcal{R}$-basis used for this sampler. $|\tilde{\mathbf{B}}_{f,g}|_{\mathcal{K}}$ is by definition the smallest value in $\mathbb{R}^+$ verifying $|\tilde{\mathbf{B}}_{f,g}|_{\mathcal{K}} \geqslant \mathbf{f}$ and $|\tilde{\mathbf{B}}_{f,g}|_{\mathcal{K}} \geqslant \tilde{\mathbf{F}}$. A straightforward computation using Definition 2 gives us

$$\tilde{\mathbf{F}} = \left(\frac{-q\bar{g}}{f\bar{f} + g\bar{g}}, \frac{q\bar{f}}{f\bar{f} + g\bar{g}} \right)$$

Therefore

$$\langle \tilde{\mathbf{F}}, \tilde{\mathbf{F}} \rangle_{\mathbb{K}} = \frac{q^2}{f\bar{f} + g\bar{g}} = \frac{q^2}{\langle \mathbf{f}, \mathbf{f} \rangle_{\mathbb{K}}}$$

Evaluating $\langle \mathbf{f}, \mathbf{f} \rangle_{\mathbb{K}}$ and $\langle \tilde{\mathbf{F}}, \tilde{\mathbf{F}} \rangle_{\mathbb{K}}$ on all the $\omega \in \Omega_m$ then yields the result. □

We then generate random NTRU bases to give an estimate on the maximum Gram-Schmidt norm given by Lemma 10. When $\|(f,g)\|$ increases, so do the values $(D(\omega))_{\omega \in \Omega_m}$, but the values $(\frac{q}{D(\omega)})_{\omega \in \Omega_m}$ decrease. The goal is to sample (f,g) with a standard deviation chosen so that $\max_{\omega \in \Omega_m}(D(\omega)) \approx \max_{\omega \in \Omega_m}(\frac{q}{D(\omega)})$, which should yield a reasonable small value for $|\tilde{\mathbf{B}}_{f,g}|_{\mathcal{K}}$. Figure 3 summarize these experiments for NTRU bases of size $2N = 1024$.

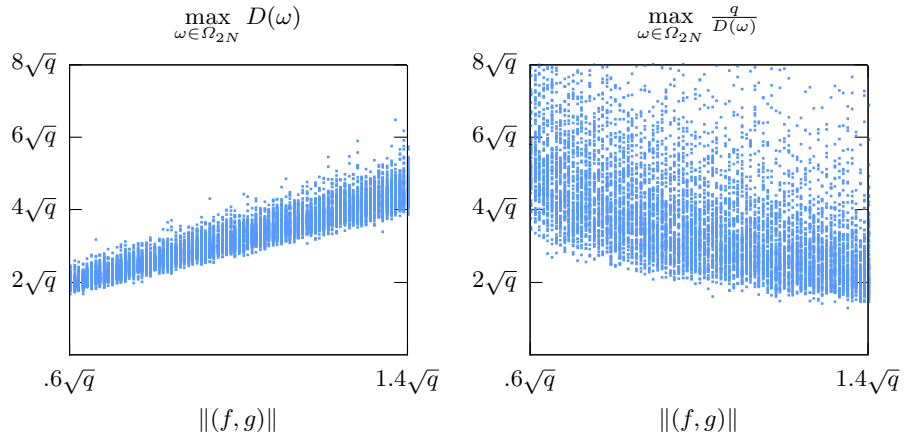

Fig. 3. The potential values of $|\tilde{\mathbf{B}}_{f,g}|_{\mathcal{K}}$, where $\mathbf{B}_{f,g}$ is a NTRU basis of dimension $2N = 1024$, for a modulus $q = 2^{20}$. On the left, $\max_{\omega \in \Omega_{2N}} D(\omega)$ and on the right, $\max_{\omega \in \Omega_{2N}} \frac{q}{D(\omega)}$, where $D(\omega)$ is defined as in Lemma 10.

In our experiments, we managed to get $|\tilde{\mathbf{B}}_{f,g}|_{\mathcal{K}} \approx 2.5\sqrt{q}$. This is just 2.5 times the theoretical smallest value that $|\tilde{\mathbf{B}}_{f,g}|_{\mathcal{K}}$ can take.

5.3 Peikert's Sampler (informally $\mathcal{R} = \mathcal{Z}^{2\times 2}, \mathbb{K} = \mathcal{K}^{2\times 2}$)

The standard deviation of the discrete Gaussian output by Peikert's sampler is $\eta'_\epsilon(\mathbb{Z}^{2N}) \cdot s_1(\mathbf{B}_{f,g})$. The following lemma gives the value of $s_1(\mathbf{B}_{f,g})$ for a NTRU basis.

Lemma 11. *Let $\mathbf{B}_{f,g}$ be as in Definition 10. The maximal singular value of $\mathbf{B}_{f,g}$ (which is also its Gram-Schmidt norm with respect to $\mathcal{Z}^{2\times 2}$) is*

$$s_1(\mathbf{B}_{f,g}) = \sqrt{\max_{\omega \in \Omega_m}(\lambda_\omega)}$$

where $C(x) \triangleq (f\bar{f}+g\bar{g}+F\bar{F}+G\bar{G})(x)$ and $\lambda_\omega \triangleq \frac{1}{2}\left(C(\omega) + (-1)^i\sqrt{C^2(\omega) - 4q^2}\right)$.

Proof. First, notice that the matrices $\mathcal{A}(f), \mathcal{A}(g), \mathcal{A}(F), \mathcal{A}(G)$ are co-diagonalizable in an orthonormal eigenbasis: there exists a matrix $\mathbf{P} \in \mathbb{C}^{N\times N}$ such that $\mathcal{A}(f) = \mathbf{P} \times f(\mathbf{\Phi}_m) \times \mathbf{P}^*$. We can therefore write

$$\mathbf{B}_{f,g} = \begin{bmatrix} \mathbf{P} & 0 \\ 0 & \mathbf{P} \end{bmatrix} \times \begin{bmatrix} f(\mathbf{\Phi}_m) & g(\mathbf{\Phi}_m) \\ F(\mathbf{\Phi}_m) & G(\mathbf{\Phi}_m) \end{bmatrix} \times \begin{bmatrix} \mathbf{P}^* & 0 \\ 0 & \mathbf{P}^* \end{bmatrix}$$

Then, see that the singular values of $\mathbf{B}_{f,g}$ are the positive square roots of the eigenvalues of $\mathbf{B}_{f,g}\mathbf{B}^t_{f,g}$. Since similar matrices share the same eigenvalues, we are therefore looking for the square roots of the eigenvalues of

$$\mathbf{M} = \begin{bmatrix} (f\bar{f} + g\bar{g})(\mathbf{\Phi}_m) & (f\bar{F} + g\bar{G})(\mathbf{\Phi}_m) \\ (\bar{f}F + \bar{g}G)(\mathbf{\Phi}_m) & (F\bar{F} + G\bar{G})(\mathbf{\Phi}_m) \end{bmatrix}$$

We compute the characteristic polynomial of $\mathbf{M}$:

$$\begin{aligned} \chi_{\mathbf{M}}(\lambda) &= \det[\mathbf{M} - \lambda I_{2N}] \\ &= \det\left[q^2 I_N - \lambda(f\bar{f} + g\bar{g} + F\bar{F} + G\bar{G})(\mathbf{\Phi}_m) + \lambda^2 I_N\right] \\ &= \prod_{\omega \in \Omega_m} \left(q^2 - \lambda(f\bar{f} + g\bar{g} + F\bar{F} + G\bar{G})(\omega) + \lambda^2\right) \end{aligned}$$

The second equality first uses the fact that $\det\begin{bmatrix} \mathbf{A} & \mathbf{B} \\ \mathbf{C} & \mathbf{D} \end{bmatrix} = \det[\mathbf{AD} - \mathbf{BC}]$ when $\mathbf{B}$ and $\mathbf{C}$ commute, then the equation 1, at at last the fact that for any $\omega \in \Omega_m$, $(f\bar{f})(\omega) = |f(\omega)|^2$.

Noting $C(\omega) \triangleq (f\bar{f} + g\bar{g} + F\bar{F} + G\bar{G})(\omega)$, the eigenvalues are, for $(\omega, i) \in (\Omega_m, \{0,1\})$, the values

$$\lambda_{\omega,i} = \frac{1}{2}\left(C(\omega) + (-1)^i\sqrt{C^2(\omega) - 4q^2}\right)$$

Noticing that the maximal $\lambda_{\omega,i}$ must have $i = 0$, this concludes the proof. □

Just like in Subsection 5.2, we ran experiments on random NTRU bases to evaluate how small we could get the singular value $s_1(\mathbf{B}_{f,g})$ to be. The results are summarized in Figure 4.

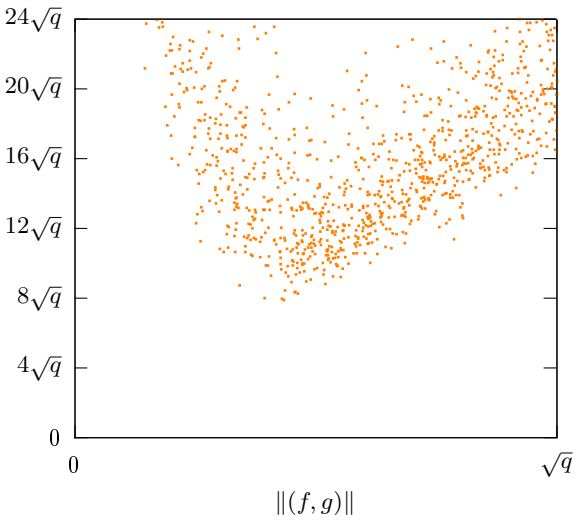

Fig. 4. Values of $s_1(\mathbf{B}_{f,g})$, where $\mathbf{B}_{f,g}$ is a NTRU basis of dimension $2N = 1024$, for a modulus $q = 2^{20}$.

The best values we managed to get in our experiments were $s_1(\mathbf{B}_{f,g}) \approx 8\sqrt{q}$. Although this could seem to be "not too big" compared to $|\tilde{\mathbf{B}}_{f,g}|_{\mathbb{R}}$ and $|\tilde{\mathbf{B}}_{f,g}|_{\mathcal{K}}$ (which condition the quality of the vectors output by Klein's and our Hybrid sampler), in practice it undermines the security parameter of the Full-Domain Hash signature scheme using Peikert's sampler.

5.4 Explaining the Gap between the Hybrid and Peikert's Samplers

The goal of this section is to explain why our Hybrid sampler performs better than Peikert's sampler over NTRU lattices.

We first recall properties of the singular norm. For two matrices $\boldsymbol{X}$ and $\boldsymbol{Y}$, and their vertical concatenation $\{\boldsymbol{X}, \boldsymbol{Y}\}$, we have:

$$\max(s_1(\boldsymbol{X})^2, s_1(\boldsymbol{Y})^2) \leqslant s_1(\{\boldsymbol{X}, \boldsymbol{Y}\})^2 \leqslant s_1(\boldsymbol{X})^2 + s_1(\boldsymbol{Y})^2. \tag{5}$$

The first inequality is an equality if and only if $\boldsymbol{X}$ and $\boldsymbol{Y}$ span orthogonal spaces. We also have euclidean sub-additivity: $s_1(\boldsymbol{X} + \boldsymbol{Y})^2 \leq s_1(\boldsymbol{X})^2 + s_1(\boldsymbol{Y})^2$.

We also recall that we write the NTRU basis as $\mathbf{B} = \{\mathbf{f}, \mathbf{F}\}$, where

$$\mathbf{f} = (f, g) \text{ and } \mathbf{F} = (F, G) = r\mathbf{f} + \tilde{\mathbf{F}}$$

for some $r \in \overline{\mathbb{K}}$ and $\tilde{\mathbf{F}}$ orthogonal to $\mathbf{f}$. The key generation process of NTRU guarantees that r has coefficients in $[-1/2, 1/2]$, and we modelize them as uniformly random in this range. Using our Hybrid sampler gives a quality of $s_1(\{\mathbf{f}, \tilde{\mathbf{F}}\}) = \max(s_1(\mathbf{f}), s_1(\tilde{\mathbf{F}}))$, against a quality of $s_1(\{\mathbf{f}, \mathbf{F}\})$ for Peikert's.

We now give three heuristics (one new and two already tested) that we will use to estimate the qualities of the samplers:

1. For $\mathbf{x} \in \{\mathbf{f}, \mathbf{F}, \tilde{\mathbf{F}}\}$, $s_1(\mathbf{x}) \approx \alpha\sqrt{\ln N}\|\mathbf{f}\|$ for some constant α. To justify this, we observe that if $y \in \overline{\mathbb{K}}$'s coefficients are gaussian (as in our experiments), then its embeddings $\sigma_i(y)$ are gaussians as well. The singular $s_1(y)$ is the maximum absolute value of these embeddings $s_1(y) = \max_\sigma |\sigma_i(y)|$, which is expected, according to order statistics to be $\approx \alpha\sqrt{\ln N}\|y\|$ for some constant α. Assuming $\mathbf{F}$, $\mathbf{f}$ and $\tilde{\mathbf{F}}$ behave like y, this gives the heuristic.
2. $\|\tilde{\mathbf{F}}\| \approx \frac{qe}{2\|\mathbf{f}\|}$. This is implicit in the heuristic of [DLP14, full version, Section 3.3], which also provides experimental confirmation.
3. $\|r\mathbf{f}\| \approx \sqrt{\frac{N}{12}}\|\mathbf{f}\|$. The heuristic can be found in [HHGP+03].

The quality of the Hybrid sampler is now discussed in points 1 and 2, and the one of Peikert's sampler in points 3 and 4.

1. Using heuristic 2, $\max(\|\mathbf{f}\|, \|\tilde{\mathbf{F}}\|)$ is minimized when $\|\mathbf{f}\| \approx \|\tilde{\mathbf{F}}\| \approx \sqrt{\frac{qe}{2}}$.
2. For the Hybrid sampler, the quality verifies:

$$s_1(\{\mathbf{f}, \tilde{\mathbf{F}}\}) = \max(s_1(\mathbf{f}), s_1(\tilde{\mathbf{F}})) \approx \alpha\sqrt{\ln N}\max(\|\mathbf{f}\|, \|\tilde{\mathbf{F}}\|)$$

   where the equality is coming from the orthogonality of $\mathbf{f}$ and $\mathbf{F}$. How small we can expect $s_1(\{\mathbf{f}, \tilde{\mathbf{F}}\})$ to be then comes from point 1.
3. We can assume that $\max(\|\mathbf{f}\|^2, \|\mathbf{F}\|^2) = \|\mathbf{F}\|^2 \approx \|\mathbf{f}\|^2 + \|\mathbf{F}\|^2$ since

$$\|\mathbf{F}\|^2 = \|\tilde{\mathbf{F}}\|^2 + \|r\mathbf{f}\|^2 \approx \frac{q^2e^2}{4\|\mathbf{f}\|^2} + \frac{N}{12}\|\mathbf{f}\|^2$$

   where the approximation comes from heuristics 2 and 3. $\|\mathbf{F}\|$ is expected to be minimal for $\|\mathbf{f}\|^2 \approx qe\sqrt{\frac{3}{N}}$, which yields $\|\mathbf{F}\| \approx \sqrt{\frac{qe\sqrt{N}}{2\sqrt{3}}}$.
4. Using equation 5 in conjunction with heuristic 1 and point 3, we approximate the quality of Peikert's sampler:

$$s_1(\{\mathbf{f}, \mathbf{F}\}) \approx \alpha\sqrt{\ln N}\|\mathbf{F}\|$$

We can now estimate the ratio between the best qualities obtained for a given dimension using the Hybrid and Peikert's samplers:

$$\frac{\sigma_{\texttt{Peikert}}}{\sigma_{\texttt{Hybrid}}} = \frac{s_1(\{\mathbf{f}, \mathbf{F}\})}{s_1(\{\mathbf{f}, \tilde{\mathbf{F}}\})} \approx \frac{\min_{\mathbf{f}} \|\mathbf{F}\|}{\min_{\mathbf{f}} \max(\|\mathbf{f}\|, \|\tilde{\mathbf{F}}\|)} \approx \left(\frac{N}{3}\right)^{1/4}$$

For $N = 512$, this ratio is about 3.6, which is close to the ratio of 3.2 that we observed in our experiments when optimizing both samplers.

However, if the samplers are ran on the same lattice, the difference can be much higher; in particular, if $\|\mathbf{f}\| \gg \|\tilde{\mathbf{F}}\|$ then from the heuristic we can expect Peikert's sampler to have a quality worse than the Hybrid by a factor $\sqrt{N/12}$, but in practice this factor is even bigger, around $\sqrt{N/4}$ (see Figure 5).

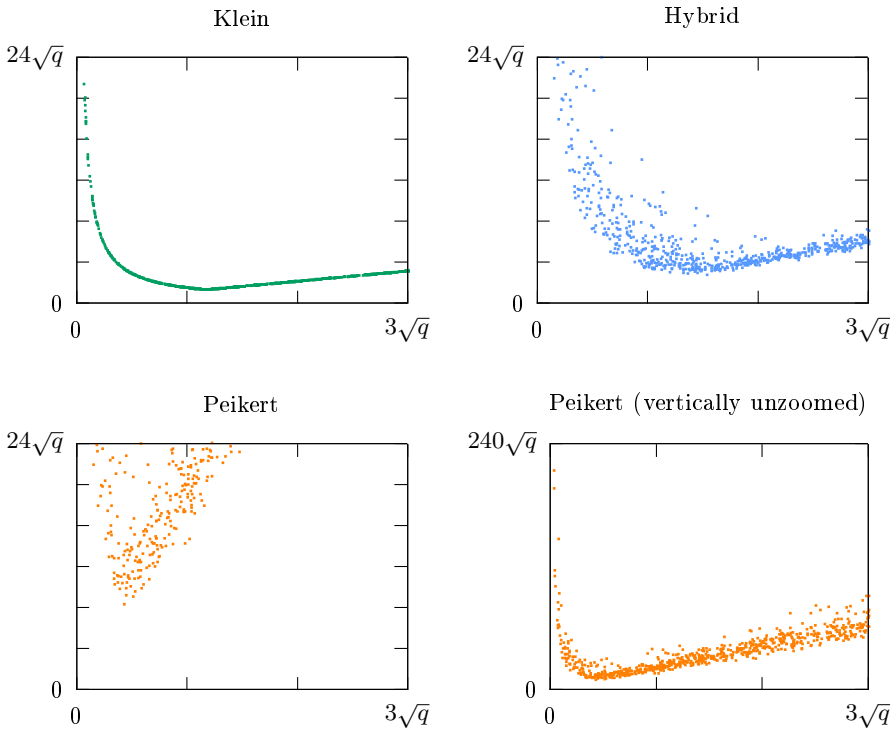

Fig. 5. Comparing the qualities of Klein's/Hybrid/Peikert's samplers for a NTRU basis of dimension $2N = 1024$ and modulus $q = 2^{20}$, with $\|(f, g)\|$ (in abscissa) ranging in $(0, 3\sqrt{q})$.

6 Trade-offs Using Subfields

Let us assume that we are working in a number field $\mathbb{K}$ of degree n that admits a subfield $\mathbb{J}$ of degree $b|n$, and let $d = n/b$ be the relative degree of $\mathbb{K}$ over $\mathbb{J}$. The field $\mathbb{K}$ may then be seen as a sub-algebra of the algebra of matrices $\mathbb{J}^{d\times d}$. That way, a basis $\mathbf{B} \in \mathbb{K}$ may be seen as a matrix of $\mathbb{J}^{dk\times dk}$, and one may run our algorithm over $\mathbb{J}$ rather than $\mathbb{K}$ to improve the quality of our sampler, at the price of slowing it down. At the extreme case, we may choose $\mathbb{J} = \mathbb{Q}$ in which case we are simply back to running the original algorithm of Klein.

Such a trade-off exists for all cyclotomic number fields of non-prime conductor. Indeed, if $\mathbb{K}_n$ denotes the n-th cyclotomic number field, then $\mathbb{K}_d$ is a subfield of $\mathbb{K}_n$ if and only if d divide n.

Because such an algebraic description is not exactly straightforward to translate into an implementation, in the next section we show explicitly how to realize this trade-off when n is a power of 2 and $d = n/2$.

6.1 Explicit Trade-off for power-of-two Cyclotomic Fields

In this subsection, m will be a power of two, $N = \varphi(m) = m/2$, $\mathbb{K}_m = \mathbb{Q}[x]/(\phi_m) = \mathbb{Q}[x]/(x^N + 1)$ and $\mathcal{Z}_m = \mathbb{Z}[x]/(\phi_m)$. The reader will notice, that, in this case, the decomposition as a matrix over a subfield is extremely similar to the steps of a Fast Fourier Transform.

Let $f \in \mathbb{K}_m$ and suppose we want to sample a spherical Gaussian over the lattice $Span_{\mathcal{Z}_m}(f)$. This lattice is either generated by the N-dimensional $\mathbb{Z}$-basis $\mathbf{B} \triangleq \mathcal{A}_N(f)$ or by the 1-dimensional $\mathcal{Z}_m$-basis f.

Using the Hybrid sampler here seems pointless at first sight, since we end up with Peikert's sampler if we use $\mathcal{Z}_m$ as base ring (i.e. $\mathbb{K}_m$ as base field).

Fortunately $\mathbb{K}_m$ is a field extension of $\mathbb{K}_{m/2}$: $\mathbb{K}_m \cong (\mathbb{K}_{m/2})^2$. More practically, if one notes $f(x) = f_1(x^2) + xf_2(x^2)$, a simple permutation of rows and columns transforms $\mathbf{B}$ into another matrix $\mathbf{B}'$ with a structure over $\mathbb{K}_{m/2}$:

$$P_\pi^t \times \mathbf{B} \times P_\pi = \left[\begin{array}{c|c} \mathcal{A}_{N/2}(f_1) & \mathcal{A}_{N/2}(f_2) \\ \hline \mathcal{A}_{N/2}(xf_2) & \mathcal{A}_{N/2}(f_1) \end{array}\right] \triangleq \mathbf{B}'$$

Where P_π is the permutation matrix associated to π, defined as:

$$\begin{cases} \pi(2k+1) = k+1 \\ \pi(2k) = \frac{N}{2} + k + 1 \end{cases}$$

Now we can use the Hybrid Sampler with the basis $\mathbf{B}'$ using $\mathcal{Z}_{m/2}$ as a base ring instead of $\mathcal{Z}_m$. This process can of course be recursively iterated: one then obtains a basis with a ring structure over $\mathcal{Z}_{m/4}$, then $\mathcal{Z}_{m/8}$, and so on, down to $\mathbb{Z}$, which corresponds to Klein's sampler.

This is of course a trade-off: while breaking down $\mathcal{Z}_m$ in smaller rings allows to sample with a smaller standard deviation, the running time of the new sampler can be up to twice longer, as $\mathcal{Z}_{m/2}$ is twice as small as $\mathcal{Z}_{m/2}$ but the basis now contains four times more ring elements. In practice, one would want to use this technique when Peikert's sampler doesn't give enough security, Klein's is too slow, and there are heavy constraints over the dimension of the lattice used (which per example might be fixed).

Figure 6 illustrates this trade-off: f is split between its even and odd coefficients, and commuting the basis elements separates them.

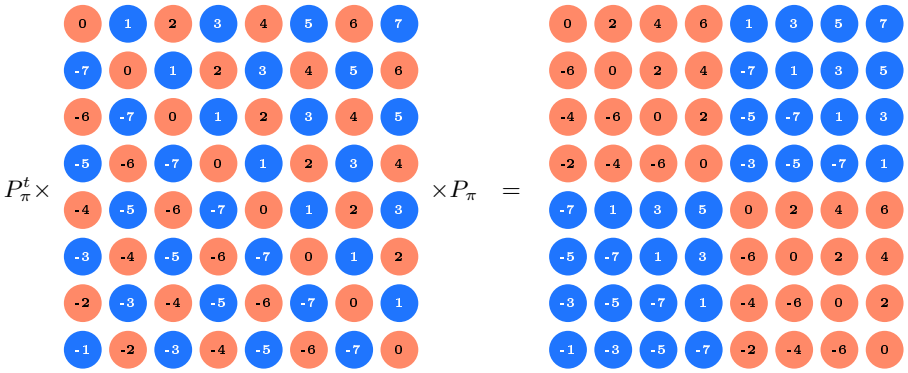

Fig. 6. An example of trade-off: using P_π as a change-of-basis matrix, with $\pi = \left(\begin{smallmatrix} 1 & 2 & 3 & 4 & 5 & 6 & 7 & 8 \\ 1 & 5 & 2 & 6 & 3 & 7 & 4 & 8 \end{smallmatrix}\right)$, allows us to turn a $\mathbb{Z}[x]/(x^8+1)$-basis into a $\mathbb{Z}[x]/(x^4+1)$-basis with twice as many elements.

References

ABB10. Shweta Agrawal, Dan Boneh, and Xavier Boyen. Lattice basis delegation in fixed dimension and shorter-ciphertext hierarchical IBE. In *Advances in Cryptology - CRYPTO 2010, 30th Annual Cryptology Conference, Santa Barbara, CA, USA, August 15-19, 2010. Proceedings*, pages 98–115, 2010.

ADRS14. Divesh Aggarwal, Daniel Dadush, Oded Regev, and Noah Stephens-Davidowitz. Solving the shortest vector problem in $2^n time via discrete gaussian sampling. *CoRR*, abs/1412.7994, 2014.

Ajt96. Miklós Ajtai. Generating hard instances of lattice problems (extended abstract). In *STOC*, pages 99–108, 1996.

Bab86. László Babai. On lovász' lattice reduction and the nearest lattice point problem. *Combinatorica*, 6(1):1–13, 1986.

Ber14a. Dan Bernstein. A subfield-logarithm attack against ideal lattices. `http://blog.cr.yp.to/20140213-ideal.html`, Febuary 2014.

Ber14b. Dan Bernstein. A subfield-logarithm attack against ideal lattices, 2014.

BGG+14. Dan Boneh, Craig Gentry, Sergey Gorbunov, Shai Halevi, Valeria Nikolaenko, Gil Segev, Vinod Vaikuntanathan, and Dhinakaran Vinayagamurthy. Fully key-homomorphic encryption, arithmetic circuit ABE and compact garbled circuits. In *Advances in Cryptology - EUROCRYPT 2014 - 33rd Annual International Conference on the Theory and Applications of Cryptographic Techniques, Copenhagen, Denmark, May 11-15, 2014. Proceedings*, pages 533–556, 2014.

BLP+13. Zvika Brakerski, Adeline Langlois, Chris Peikert, Oded Regev, and Damien Stehlé. Classical hardness of learning with errors. In *Symposium on Theory of Computing Conference, STOC'13, Palo Alto, CA, USA, June 1-4, 2013*, pages 575–584, 2013.

Boy10. Xavier Boyen. Lattice mixing and vanishing trapdoors: A framework for fully secure short signatures and more. In *Public Key Cryptography*, pages 499–517, 2010.

CDPR15. Ronald Cramer, Léo Ducas, Chris Peikert, and Oded Regev. Recovering short generators of principal ideals in cyclotomic rings. Cryptology ePrint Archive, Report 2015/313, 2015. `http://eprint.iacr.org/`.

CHKP10. David Cash, Dennis Hofheinz, Eike Kiltz, and Chris Peikert. Bonsai trees, or how to delegate a lattice basis. In *EUROCRYPT*, pages 523–552, 2010.

CN11. Yuanmi Chen and Phong Q. Nguyen. Bkz 2.0: Better lattice security estimates. In *ASIACRYPT*, pages 1–20, 2011.

DD12. Léo Ducas and Alain Durmus. Ring-lwe in polynomial rings. In *Public Key Cryptography–PKC 2012*, pages 34–51. Springer, 2012.

DDLL13. Léo Ducas, Alain Durmus, Tancrède Lepoint, and Vadim Lyubashevsky. Lattice signatures and bimodal gaussians. In *CRYPTO (1)*, pages 40–56, 2013.

DLP14. Léo Ducas, Vadim Lyubashevsky, and Thomas Prest. Efficient identity-based encryption over NTRU lattices. In *Advances in Cryptology - ASIACRYPT 2014 - 20th International Conference on the Theory and Application of Cryptology and Information Security, Kaoshiung, Taiwan, R.O.C., December 7-11, 2014, Proceedings, Part II*, pages 22–41, 2014.

DN12. Léo Ducas and Phong Q. Nguyen. Faster gaussian lattice sampling using lazy floating-point arithmetic. In *Advances in Cryptology - ASIACRYPT 2012 - 18th International Conference on the Theory and Application of Cryptology and Information Security, Beijing, China, December 2-6, 2012. Proceedings*, pages 415–432, 2012.

Gal12. Steven D. Galbraith. *Mathematics of Public Key Cryptography*. Cambridge University Press, New York, NY, USA, 1st edition, 2012.

GN08. Nicolas Gama and Phong Q. Nguyen. Predicting lattice reduction. In *EUROCRYPT*, pages 31–51, 2008.

GPV08. Craig Gentry, Chris Peikert, and Vinod Vaikuntanathan. Trapdoors for hard lattices and new cryptographic constructions. In *STOC*, pages 197–206, 2008.

HHGP+03. Jeffrey Hoffstein, Nick Howgrave-Graham, Jill Pipher, Joseph H. Silverman, and William Whyte. Ntrusign: Digital signatures using the ntru lattice. In *CT-RSA*, pages 122–140, 2003.

HPS98. Jeffrey Hoffstein, Jill Pipher, and Joseph H. Silverman. NTRU: A ring-based public key cryptosystem. In *ANTS*, pages 267–288, 1998.

Kle00. Philip N. Klein. Finding the closest lattice vector when it's unusually close. In *SODA*, pages 937–941, 2000.

LP15. Vadim Lyubashevsky and Thomas Prest. Quadratic time, linear space algorithms for gram-schmidt orthogonalization and gaussian sampling in structured lattices. In *Advances in Cryptology - EUROCRYPT 2015 - 34th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Sofia, Bulgaria, April 26-30, 2015, Proceedings, Part I*, pages 789–815, 2015.

LPR10. Vadim Lyubashevsky, Chris Peikert, and Oded Regev. On ideal lattices and learning with errors over rings. In *EUROCRYPT*, pages 1–23, 2010.

LPR13. Vadim Lyubashevsky, Chris Peikert, and Oded Regev. A toolkit for ring-lwe cryptography. In *EUROCRYPT*, pages 35–54, 2013.

LW15. Vadim Lyubashevsky and Daniel Wichs. Simple lattice trapdoor sampling from a broad class of distributions. In *Public-Key Cryptography - PKC 2015 - 18th IACR International Conference on Practice and Theory in Public-Key Cryptography, Gaithersburg, MD, USA, March 30 - April 1, 2015, Proceedings*, pages 716–730, 2015.

MP12. Daniele Micciancio and Chris Peikert. Trapdoors for lattices: Simpler, tighter, faster, smaller. In *EUROCRYPT*, pages 700–718, 2012.

MR07. Daniele Micciancio and Oded Regev. Worst-case to average-case reductions based on gaussian measures. *SIAM J. Comput.*, 37(1):267–302, 2007.

PDG14. Thomas Pöppelmann, Léo Ducas, and Tim Güneysu. Enhanced lattice-based signatures on reconfigurable hardware. In *Cryptographic Hardware and Embedded Systems–CHES 2014*, pages 353–370. Springer, 2014.

Pei10. Chris Peikert. An efficient and parallel gaussian sampler for lattices. In *CRYPTO*, pages 80–97, 2010.

SS11. Damien Stehlé and Ron Steinfeld. Making ntru as secure as worst-case problems over ideal lattices. In *Advances in Cryptology–EUROCRYPT 2011*, pages 27–47. Springer, 2011.