Preimage sampleable function families without discrete Gaussians

Eamonn W. Postlethwaite and Filip Trenkić

King's College London

Abstract. We construct preimage sampleable function families on q-ary lattices [Gentry–Peikert–Vaikuntanathan, STOC'08] for which preimage sampling reduces to sampling uniform points of unitriangular lattices from simple polytopes: affine linear transforms on the ℓ_1 and ℓ_∞ balls, and a scaled intersection thereof. We build the necessary samplers by adapting an algorithm of [Kannan–Vempala, STOC'97] and improving its analysis. This sampling requires only uniform bits and affine linear transforms on the uniform distribution on $[0, 1]$. The collision resistance of these families relies on the short integer solutions problem [Ajtai, STOC'96] in various ℓ_p norms. Considering the Lee metric as the ℓ_1 norm in q-ary lattices, we answer an open question to construct such families for the Lee metric [Hörmann–van Woerden, CRYPTO'24]. We also answer an open question of [Plançon–Prest, PKC'21] by sampling from polytopes with inradius smaller by a factor almost square root in the lattice rank. We provide a generic framework for constructing preimage sampleable function families from polytopes with sufficient conditions for realising it. While our parameters are worse than prior discrete Gaussian based constructions, such distributions are challenging from a physical security perspective.

1 Introduction

Preimage sampleable function families were introduced in [GPV08]. If they have a collision resistance property then the authors construct strong existentially unforgeable digital signatures from them. They also instantiate such a family over an ensemble of q-ary lattices with two important properties. First, a party in possession of a trapdoor, some special basis of a lattice from the ensemble, is able to efficiently compute samples from a discrete Gaussian distribution with arbitrary centre and somewhat small width over this lattice. Second, the short integer solution problem is hard over this ensemble. These two properties are ensured via a trapdoor generation procedure that samples a uniform lattice from the ensemble along with a trapdoor basis. The trapdoor basis has geometric properties that allow the discrete Gaussian distribution to be efficiently sampled. Worst case lattice problems reduce to the short integer solution problem over this lattice ensemble, supporting its presumed hardness [MR07,Pei08,GPV08,MP13].

This paradigm for digital signatures schemes has been of great importance, for example the NIST standard FALCON [PFH+22] realises it (see [FGdG+26]

for details). One difficulty that is clarified by practice is the requirement to sample from discrete Gaussian distributions over lattices [GMRR22,GR24].

We use an algorithm inspired by Kannan and Vempala [KV97] to, assuming certain geometric conditions are met, sample uniform lattice points from a polytope, and instantiate new collision resistant preimage sampleable function families over ensembles of q-ary lattices. These families maintain the connection to worst case lattice problems but do not require the ability to sample from discrete Gaussian distributions. In particular preimage sampling requires only uniform bits, simple transforms on the uniform distribution on $[0, 1]$, and Babai's nearest plane algorithm. We instantiate our collision resistant preimage sampleable function families with three polytopes which have good properties, and provide a generic framework for other polytopes satisfying these properties.

Acknowledgements. We thank the anonymous reviewers for alerting us to prior work, and for their detailed comments. We thank Sasha Lapiha for helpful discussions.

1.1 Technical overview and contributions

Effective polytopes. A polytope in $\mathbb{R}^d$ may be represented as f linear inequalities in d variables. However, this representation is unsuitable for computation when f is large. To remedy this, we define an *effective* polytope: a polytope for which we can check membership and sample the uniform distribution, in polynomial time in d. We show a large class of polytopes are effective, in particular allowing us to work with certain polytopes where f is superpolynomial in d.

Kannan–Vempala algorithm. In [KV97] the authors give an algorithm to sample, almost uniformly, integer points in polytopes with sufficiently large *inradius* – the maximum radius of a ball contained within the polytope. In other words, for a polytope P with large inradius, one can sample close to uniform on $P \cap \mathbb{Z}^d$.

However, the algorithm only achieves constant statistical distance to uniform, and requires sampling from a modified polytope P' which is not necessarily effective. We introduce a variant of their algorithm, and subsequent analysis, that satisfies our requirements.

- We extend the algorithm from the integer lattice to all *unitriangular lattices*, a broader class of lattices, by generalising a randomised rounding procedure and without degrading the quality of the sampler.
- We change the construction of the modified polytope P', so that if P is effective then so is P', again without degrading the quality of the sampler.
- We present explicit bounds on the statistical distance to uniform achieved by the algorithm. Additionally, the algorithm may internally repeat and we bound the expected number of repetitions.
- We give concrete and asymptotic conditions on the inradius of P such that one can achieve *negligible* statistical distance within a *constant* number of expected repetitions.

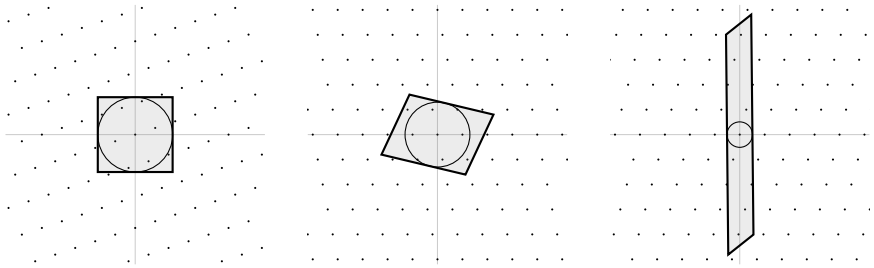

(a) A polytope P and a lattice Λ.

(b) $\widetilde{\mathbf{T}}^{-1}P$ and $\widetilde{\mathbf{T}}^{-1}\Lambda$ where $\mathbf{T}$ is a trapdoor basis of Λ.

(c) $\widetilde{\mathbf{B}}^{-1}P$ and $\widetilde{\mathbf{B}}^{-1}\Lambda$, where $\mathbf{B}$ is an arbitrary basis of Λ.

Fig. 1: Illustration of the core principle of our construction. Knowledge of a trapdoor basis $\mathbf{T}$ of Λ yields a polytope with large inradius, in which efficiently sampling a unitriangular lattice point is possible; but in general for an arbitrary basis, the resulting inradius is too small.

Our analysis is more explicit and flexible than that of [KV97] and we believe it may be of independent interest.

Trapdoor sampling for polytopes. Given the above we show that one can sample almost uniformly from the intersection of an effective polytope P and any lattice $\Lambda = \Lambda(\mathbf{B})$, or any coset thereof, given that P has sufficiently large inradius and an appropriate trapdoor basis is known. Our method is comparable to that of [PP21], but does not rely on heuristic reasoning regarding the number of expected internal repetitions, and allows sampling from smaller P (see §4.1) which was left as an open question.

In short, sampling $\mathsf{U}(P \cap \Lambda)$ is equivalent to sampling $\mathsf{U}(\widetilde{\mathbf{B}}^{-1}P \cap \widetilde{\mathbf{B}}^{-1}\Lambda)$, where $\widetilde{\mathbf{B}}$ is the Gram–Schmidt orthogonalisation of the basis $\mathbf{B}$. The transformed lattice is unitriangular, and the transformed polytope $\widetilde{\mathbf{B}}^{-1}P$ has inradius smaller than that of P by a factor at most $\|\widetilde{\mathbf{B}}\|$.

A trapdoor basis $\mathbf{T}$ produces $\widetilde{\mathbf{T}}^{-1}P$ with inradius not too much smaller than that of P, hence efficient sampling is possible by our modified Kannan–Vempala algorithm. In particular a trapdoor basis is one with short Gram–Schmidt orthogonalisation, and generating random lattices with such trapdoors is well studied [Ajt99,AP11,MP12]. On the other hand, an arbitrary basis $\mathbf{B}$ will in general lead to an inradius too small for efficient sampling. This core geometric principle is illustrated in Figure 1.

Generic constructions for CR-PSFs. We provide a generic construction of a collision resistant preimage sampleable function family based on an arbitrary effective polytope P. The construction assumes a list of conditions on P, most importantly that the inradius is large enough for trapdoor sampling while the associated collision resistance problem remains hard.

In the construction, we use our trapdoor sampling to sample a uniformly random solution for the following ISIS style problem

$$\text{Find } \mathbf{z} \in P \cap \mathbb{Z}^m \text{ such that } \mathbf{A}\mathbf{z} = \mathbf{u}, \text{ for } \mathbf{A} \leftarrow \mathsf{U}(\mathbb{Z}_q^{n \times m})$$

given a trapdoor basis of the corresponding q-ary lattice $\Lambda_q^{\perp}(\mathbf{A})$.

The construction has collision resistance if the following SIS style problem

$$\text{Find } \mathbf{z} \in ((P - P) \cap \mathbb{Z}^m) \setminus \{\mathbf{0}\} \text{ such that } \mathbf{A}\mathbf{z} = \mathbf{0}, \text{ for } \mathbf{A} \leftarrow \mathsf{U}(\mathbb{Z}_q^{n \times m})$$

is hard, where $P - P = \{\mathbf{x} - \mathbf{y} : \mathbf{x}, \mathbf{y} \in P\}$ is the *difference body* of P.

Specific instantiations. We instantiate the construction when P is a *hypercube*, and the hardness of collision resistance admits a reduction from SIS^{∞}. We also instantiate the construction when P is an *orthoplex*, reducing from SIS^{1}. Our constructions are simple from an implementation point of view, notably they do not require sampling discrete Gaussians on lattices.

These constructions allow uniform preimage sampling for ISIS^{∞} and ISIS^{1} respectively, which may be of interest in applications which naturally operate in the ℓ_∞ or ℓ_1 norm. Since the ℓ_1 norm on a q-ary lattice corresponds to the *Lee metric* of the underlying code, this solves an open question in [HvW24]:

> *"... generally it remains an open question how to adapt the GPV framework to the Lee metric."*

Finally, to illustrate the generality of our approach, we let P be a polytope considered in [BBRS24] which is an intersection of the hypercube and a scaled orthoplex. This polytope approximates the ℓ_2-ball, so we provide a reduction from SIS^{2} as well as SIS^{1}.

1.2 Open questions

1. *Better polytopes.* In §6 we give three instantiations for P and in §7 compare the resulting CR-PSFs, but our construction is generic; do there exist better choices for P? One approach to answering this question would be to find effective polytopes with few facets, i.e. $f \in \text{poly}(d)$, which are good approximations of the Euclidean ball. By good approximation one may imagine polytopes with close in and out radii, or more generally good solutions to the isoperimetric inequality for polytopes [Gru07, §11.2], which allow for tighter reductions via SIS^{2}.
2. *Improved trapdoor sampling.* Does there exists a better algorithm to sample unitriangular lattice points in polytopes? For the integer lattice, tightness of the inradius condition is known only up to a factor $\sqrt{\log f}$, see [KV97, §4]. Removing this factor would in our context allow sampling from smaller polytopes, especially when $f \notin \text{poly}(d)$.

1.3 Related works

The definition of CR-PSFs was first given in [GPV08]. A line of recent works in the Fiat–Shamir with aborts paradigm for lattice based digital signatures has considered optimal sets of integer points and distributions over which to sample them [DFPS22,CCD+24,BBRS24]. In particular [BBRS24] introduced the polytope $\mathcal{H}$ and several geometric facts about it that we make extensive use of in one of our instantiations of our polytope CR-PSF framework. The signature scheme FuLeeca [RBK+23] provided a form of CR-PSF using the Lee metric on codes (corresponding to the ℓ_1 norm in lattices) with a heuristic proof that preimages did not leak information about the trapdoor, which was shown to be false in [HvW24].

The CR-PSFs given in [GPV08] rely on trapdoor sampling for discrete Gaussian distributions, based on an algorithm of Klein [Kle00]. In [MP12] a new type of trapdoor is introduced which also allows for discrete Gaussian sampling.

Trapdoor sampling for distributions other than discrete Gaussians has been studied [LW15,PP21]. The first yields samplers for a wide range of distributions, but cannot sample uniformly from standard polytopes due to the structure of the trapdoors in [MP12] that it requires. The second yields sampling for 'square-monic' distributions, which include uniformly from polytopes, but relies on a heuristic; our method is provable and yields sampling from smaller polytopes. We elaborate, with detailed comparisons, in §4.1.

1.4 Organisation

In §2 we introduce preliminaries on polytopes (including a definition of effective polytopes and sampling routines), lattices, the short integer solutions problem and CR-PSFs. In §3 we describe our variant of the Kannan–Vempala algorithm with generalised randomised rounding along with our improved analysis. In §4 we describe how to use the algorithm of §3 to perform trapdoor sampling for polytopes and give minimal geometric conditions on polytopes for achieving negligible statistical distance from uniform and some expected number of internal repeats. We also compare more fully to prior art [GPV08,LW15,PP21]. In §5 we give a leftover hash lemma variant and our generic construction of CR-PSFs for polytopes. In §6 we realise our generic construction with three polytopes, the ℓ_1 and ℓ_∞ balls and $\mathcal{H}$ from [BBRS24]. Finally, in §7 we compare the three constructions to each other.

2 Preliminaries

For $n \in \mathbb{N}$ let $[n] = \{1, \ldots, n\}$. Given a Lebesgue measurable set $X \subset \mathbb{R}^d$ let $\mathrm{vol}(X)$ denote its volume. For $p \in [0,1]$ let $X \sim \mathrm{Ber}(p)$ be the Bernoulli random variable on p, where $\mathbb{P}[X = 0] = 1 - p$ and $\mathbb{P}[X = 1] = p$. For random variable X its min entropy is $H_\infty(X) = -\log_2(\max_{x \in \mathrm{supp}(X)} \mathbb{P}[X = x])$ and its expectation is $\mathbb{E}[X]$. If X has $\mathrm{supp}(X) \subseteq [0, \infty)$ and $a > 0$ then Markov's

inequality gives $\mathbb{P}[X \geq a] \leq \mathbb{E}[X]/a$. Let X, Y be discrete random variables and $Z = \text{supp}(X) \cup \text{supp}(Y)$ then the statistical distance between X and Y is

$$\Delta(X,Y) = \frac{1}{2}\sum_{z \in Z} |\mathbb{P}[X=z] - \mathbb{P}[Y=z]|.$$

Throughout, log is the natural logarithm, and $\log_x$ for $x > 1$ is the base x logarithm. For positive integer q let $\mathbb{Z}_q = \mathbb{Z}/q\mathbb{Z}$ and $[\cdot]_q \colon \mathbb{Z} \to \mathbb{Z}_q$ be the (unique) map sending an integer x to its modulo class $[x]_q$. We call any map $\phi \colon \mathbb{Z}_q \to \mathbb{Z}$ a (modulo q) lift whenever $[\phi(x)]_q = x$ for all $x \in \mathbb{Z}_q$. Both $[\cdot]_q$ and a lift ϕ may be extended entrywise to vectors or matrices. For $p \in [1, \infty)$ we define the ℓ_p norm $\|\cdot\|_p \colon \mathbb{R}^d \to \mathbb{R}_{\geq 0}$ as $\mathbf{x} \mapsto (\sum_i |x_i|^p)^{1/p}$ and for $p = \infty$ as $\|\mathbf{x}\|_\infty = \max_i |x_i|$. For matrix $\mathbf{M} = (\mathbf{x}_i)_i$ we let $\|\mathbf{M}\| = \max_i \|\mathbf{x}_i\|$. For $p \in [1, \infty]$, $r \in \mathbb{R}$ and $\mathbf{y} \in \mathbb{R}^d$ we let $B_d^p(r; \mathbf{y}) = \{\mathbf{x} \in \mathbb{R}^d \colon \|\mathbf{x} - \mathbf{y}\|_p \leq r\}$ denote the closed ℓ_p ball of radius r around $\mathbf{y}$ in $\mathbb{R}^d$. For norms and balls, $p = 2$ if omitted, and for balls $\mathbf{y} = \mathbf{0}$ if omitted and B_d^p implies $r = 1$. Note that $p \in \{1, \infty\}$ are the only such that $B_d^p(r; \mathbf{y})$ is a polytope (see §2.1). Let $\text{GL}_d(\mathbb{R}) = \{\mathbf{S} \in \mathbb{R}^{d \times d} \colon \det(\mathbf{S}) \neq 0\}$. For $\mathbf{T} \in \mathbb{R}^{d' \times d}$ its first singular value is $s_1(\mathbf{T}) = \max\{\|\mathbf{T}\mathbf{u}\| \colon \mathbf{u} \in B_d(1)\}$.

	$p=1$	$p=2$	$p=\infty$
$\text{vol}(B_d^p(r))$	$\frac{(2r)^d}{d!}$	$\frac{\pi^{d/2}}{\Gamma\left(\frac{d}{2}+1\right)} r^d$	$(2r)^d$

Table 1: Volumes of p-balls, for some p.

2.1 Polytopes

Definition 1 (Polytope). *A (closed, convex)* polytope $P \subset \mathbb{R}^d$ *is the convex hull of finitely many points in* $\mathbb{R}^d$*. We say* P *is full dimensional if it does not lie in any proper affine subspace of* $\mathbb{R}^d$*, i.e. it has positive volume,* $\text{vol}(P) > 0$.

Lemma 1. *A closed, convex, full dimensional polytope* P *may be represented as the set of solutions to a system of* f *linear inequalities*

$$P = \{\mathbf{x} \in \mathbb{R}^d \colon \mathbf{M}^t\mathbf{x} \leq \mathbf{v}\} = \{\mathbf{x} \in \mathbb{R}^d \colon \wedge_{i=1}^f \langle \mathbf{m}_i, \mathbf{x} \rangle \leq v_i\},$$

where $\mathbf{M} \in \mathbb{R}^{d \times f}$ *has* ith *column* $\mathbf{m}_i$ *and* $\mathbf{v} \in \mathbb{R}^f$*. When* P *is full dimensional and* f *is minimal the representation is unique (up to permutation and rescaling).*

Henceforth we work with such polytopes and their minimal representations, and use the shorthand $P = (\mathbf{M}, \mathbf{v})$. We may view P as the region of $\mathbb{R}^d$ bounded by the f hyperplanes $H_i = \{\mathbf{x} \in \mathbb{R}^d \colon \langle \mathbf{m}_i, \mathbf{x} \rangle = v_i\}$. Each set $P \cap H_i$ is called a *facet* of the polytope, and so P has f facets. A polytope must have at least $f \geq d + 1$ facets (attained by the *d-simplex*).

Definition 2. *Given $\mathbf{c} \in P$ write $r(P;\mathbf{c})$ for the radius of the largest Euclidean ball centred at $\mathbf{c}$ and contained within P,*

$$r(P;\mathbf{c}) = \max\{r \geq 0 : B_d(r;\mathbf{c}) \subseteq P\}.$$

The inradius of P is the maximum over all $\mathbf{c} \in P$,

$$r(P) = \max_{\mathbf{c}\in P} r(P;\mathbf{c}).$$

We give some facts about scaling and transforming polytopes, then some examples of some common polytopes and their representations.

Lemma 2. *Let $P = (\mathbf{M}, \mathbf{v})$ be a polytope, $\alpha > 0$, $\mathbf{y} \in \mathbb{R}^d$ and $\mathbf{S} \in \mathrm{GL}_d(\mathbb{R})$, then $\alpha P = (\mathbf{M}, \alpha\mathbf{v})$, $\mathbf{y} + P = (\mathbf{M}, \mathbf{v} + \mathbf{M}^t\mathbf{y})$ and $\mathbf{S}P = (\mathbf{S}^{-t}\mathbf{M}, \mathbf{v})$.*

Lemma 3. *If $\mathbf{c} \in P = (\mathbf{M}, \mathbf{v})$ then*

$$r(P;\mathbf{c}) = \min_{i\in[f]} \frac{(\mathbf{v} - \mathbf{M}^t\mathbf{c})_i}{\|\mathbf{m}_i\|}$$

Proof. For the case $\mathbf{c} = \mathbf{0}$, recall P is the region bounded by the hyperplanes $H_i = \{\mathbf{x} \in \mathbb{R}^d \colon \langle \mathbf{m}_i, \mathbf{x}\rangle = v_i\}$, with $v_i \geq 0$ when $\mathbf{0} \in P$, and each has minimal distance $v_i/\|\mathbf{m}_i\|$ from the origin. For the general case, note $r(P;\mathbf{c}) = r(P-\mathbf{c};\mathbf{0})$ where $P - \mathbf{c} = (\mathbf{M}, \mathbf{v} - \mathbf{M}^t\mathbf{c})$ by Lemma 2. □

Lemma 4. *Let $\mathbf{c} \in P = (\mathbf{M}, \mathbf{v})$, and $\mathbf{S} \in \mathrm{GL}_d(\mathbb{R})$, then*

$$r(\mathbf{S}P;\mathbf{S}\mathbf{c}) \geq \frac{r(P;\mathbf{c})}{s_1(\mathbf{S}^{-1})} \quad \textit{and} \quad r(\mathbf{S}P) \geq \frac{r(P)}{s_1(\mathbf{S}^{-1})}.$$

Proof. For the first statement it suffices to proof the case $\mathbf{c} = \mathbf{0}$, as the general case follows by translation.[1] If $\mathbf{0} \in \partial P$ then $r(P;\mathbf{0}) = 0$ and the statement is true. Else, $B_d(\varepsilon) \subset P$ for some $\varepsilon > 0$ and all $v_i \neq 0$. By Lemmas 2 and 3,

$$r(\mathbf{S}P;\mathbf{0}) = \min_{i\in[f]} \left(\frac{v_i}{\|\mathbf{S}^{-t}\mathbf{m}_i\|}\right) = \frac{1}{\max_i \|\mathbf{S}^{-t}\mathbf{m}_i/v_i\|}.$$

We upper bound the latter denominator via

$$\max_i \|\mathbf{S}^{-t}\mathbf{m}_i/v_i\| \leq \max_i \frac{\|\mathbf{S}^{-t}\mathbf{m}_i\|}{\|\mathbf{m}_i\|} \cdot \max_i \frac{\|\mathbf{m}_i\|}{v_i} \leq s_1(\mathbf{S}^{-t})/r(P;\mathbf{0})$$

and use $s_1(\mathbf{S}^{-t}) = s_1(\mathbf{S}^{-1})$. For the second statement, take the maximum of both sides over $\mathbf{c} \in P$ and note $\mathbf{S}$ is a bijection $P \to \mathbf{S}P$. □

Definition 3 (Common polytopes). *The d-orthoplex[2] is B_d^1, the d-hypercube is B_d^∞ and a parallelotope[3] is any set of the form $\mathbf{S} \cdot B_d^\infty$ with $\mathbf{S} \in \mathrm{GL}_d(\mathbb{R})$. In particular*

[1] $r(\mathbf{S}P;\mathbf{S}\mathbf{c}) = r(\mathbf{S}(P-\mathbf{c});\mathbf{0}) \geq r(P-\mathbf{c};\mathbf{0})/s_1(\mathbf{S}^{-1}) = r(P;\mathbf{c})/s_1(\mathbf{S}^{-1})$.

[2] Commonly called the *cross-polytope*.

[3] Sometimes 'parallelepiped', though this can mean $d = 3$ specifically.

1. $P = B_d^1 = (\mathbf{M}, \mathbf{v})$ *with* $f = 2^d$ *facets defined by* $\{\mathbf{m}_i\}_i = \{-1, 1\}^d$ *and each* $v_i = 1$,
2. $P = B_d^\infty = (\mathbf{M}, \mathbf{v})$ *with* $f = 2d$ *facets defined by* $\{\mathbf{m}_i\}_i = \{\pm \mathbf{e}_i\}_i$ *(the positive and negative canonical basis vectors) and each* $v_i = 1$,
3. $P = \mathbf{S} \cdot B_d^\infty = (\mathbf{M}, \mathbf{v})$ *with* $f = 2d$ *facets defined by* $\{\mathbf{m}_i\}_i = \{\pm \mathbf{S}^{-t} \mathbf{e}_i\}_i$ *and each* $v_i = 1$ *by Lemma 2.*

We require the ability to efficiently sample $\mathsf{U}(P)$ for such common polytopes. We assume we can sample $U = \mathsf{U}([0, 1])$ to a finite but arbitrary precision. To sample B_d^∞ it is enough to return $2 \cdot (U - 1/2)^d$. Sampling B_d^1 requires somewhat more care. We note [BGMN05, Thm. 1] gives a uniform sampler for B_d^p for $p \geq 1$ however we give a method that is potentially simpler from an implementation perspective that combines the following definition and two facts.

Definition 4 (d-simplex). *Let $d \in \mathbb{N}$ then the d-simplex is*

$$S_d = \left\{ \mathbf{x} \in \mathbb{R}^{d+1} : \ \wedge_i x_i \geq 0 \wedge \sum_i x_i = 1 \right\} \subseteq [0, 1]^{d+1}.$$

Lemma 5 ([BGMN05, Cor. 4] with $p = 1$). *Let $(x_1, \ldots, x_{d+1})$ be sampled according to $\mathsf{U}(\partial B_{d+1}^1)$ then $\mathbf{x} = (x_1, \ldots, x_d) \sim \mathsf{U}(B_d^1)$.*

Lemma 6 ([Dev86, Thm. 2.1]). *Let $U_0 = 0$ and $U_{d+1} = 1$. For $i \in [d]$ let $U_i \sim U$ and permute the indices such that $U_i \geq U_{i-1}$. For $i \in [d+1]$ let $s_i = U_i - U_{i-1}$. Then $\mathbf{s} = (s_1, \ldots, s_{d+1})$ is distributed according to $\mathsf{U}(S_d)$.*

The boundary ∂B_{d+1}^1 is formed of 2^{d+1} facets that are each a d-simplex reflected into one of the 2^{d+1} orthants. In particular, to sample ∂B_{d+1}^1 first sample $\mathsf{U}(S_d)$ per Lemma 6 and let $d+1$ fair coins via $\lfloor U \rceil^{d+1}$ determine the orthant. Dropping the final entry gives a sample from $\mathsf{U}(B_d^1)$ by Lemma 5.

2.2 Effective polytopes

Polytopes of interest may have $f \notin \text{poly}(d)$ so the $P = (\mathbf{M}, \mathbf{v})$ representation is not efficient to read, nevermind compute on. For example B_d^1 has $f = 2^d$ per Definition 3. As a computational model for polytopes we instead assume a limited number of efficient operations, and realise these assumptions for particular P, which we call effective. We may still prove geometric statements about P and related polytopes using the $P = (\mathbf{M}, \mathbf{v})$ representation.

Definition 5 (Effective P). *We call polytope $P \subset \mathbb{R}^d$ effective if the following operations can be realised in polynomial time in d.*

1. *Determining whether $\mathbf{x} \in P$, for any $\mathbf{x} \in \mathbb{R}^d$.*
2. *Sampling a uniform point $\mathbf{x} \leftarrow \mathsf{U}(P)$.*

In practice, we may realise the uniform sampler up to some prescribed statistical distance. We show that effectiveness is preserved under dilation, translation, and the action of $\text{GL}_d(\mathbb{R})$, and hence a large class of polytopes are effective.

Lemma 7. *Let $\alpha > 0$, $\mathbf{y} \in \mathbb{R}^d$ and $\mathbf{S} \in \mathrm{GL}_d(\mathbb{R})$. If P is an effective polytope then so are αP, $P + \mathbf{y}$, and $\mathbf{S}P$.*

Proof. Note $\mathbf{y} \in \alpha P \iff \alpha^{-1}\mathbf{y} \in P$ which is efficient to check by hypothesis. Analogous statements hold for $P + \mathbf{y}$ and $\mathbf{S}P$.

For Lebesgue measureable $X \subseteq P$ we have $\mathrm{vol}(\alpha X) = \alpha^d X$, therefore $\mathrm{vol}(\alpha X)/\mathrm{vol}(\alpha P) = \mathrm{vol}(X)/\mathrm{vol}(P)$. Thus $\mathsf{U}(\alpha P) = \alpha \cdot \mathsf{U}(P)$. Sampling from the latter is efficient by hypothesis. Translations preserve the Lebesgue measure and thus uniform distributions, hence the respective statement for $P+\mathbf{y}$. Finally, write $\mathbf{S} = \gamma\mathbf{T}$ for $\gamma > 0$ and $\det(\mathbf{T}) \in \{-1, 1\}$. The action of $\mathbf{T}$ preserves the Lebesgue measure and γ is a dilation, hence the respective statement for $\mathbf{S}P$. □

2.3 Martingales

Definition 6 (Martingale). *A (discrete-time)* martingale *is a sequence of random variables $Z_0, Z_1, Z_2, \ldots$ such that for all $t > 0$,*

$$\mathbb{E}[|Z_t|] < \infty$$
$$\mathbb{E}[Z_{t+1} \mid Z_1, \ldots, Z_t] = Z_t.$$

Lemma 8 (Azuma's inequality [Azu67]). *Let $(Z_i)_{i\geq 0}$ be a martingale with its differences $|Z_t - Z_{t-1}| \leq c_t$ bounded. For every $\varepsilon > 0$ and positive integer T*

$$\mathbb{P}\left[Z_T - Z_0 > \varepsilon\right] \leq \exp\left(\frac{-\varepsilon^2}{2\sum_{t=1}^{T} c_t^2}\right).$$

2.4 Lattices

A lattice Λ is a discrete subgroup of $(\mathbb{R}^d, +)$. That is $(\Lambda, +) \leq (\mathbb{R}^d, +)$ and there exists $\ell > 0$ such that $B_d(\ell) \cap \Lambda = \{\mathbf{0}\}$. We consider only full rank lattices, therefore there exists some linearly independent $\mathbf{B} = (\mathbf{b}_i)_{i=1}^d$, called a basis, such that $\Lambda = \mathbf{B}\mathbb{Z}^d$, and we write $\Lambda = \Lambda(\mathbf{B})$. We define $\mathrm{vol}(\Lambda) = \det(\mathbf{B})$ for any basis $\mathbf{B}$ of Λ. A q-ary lattice Λ is such that $q\mathbb{Z}^d \subseteq \Lambda \subseteq \mathbb{Z}^d$.

Definition 7 (Gram–Schmidt Orthogonalisation, GSO). *Given matrix $\mathbf{B} = (\mathbf{b}_1, \ldots, \mathbf{b}_d)$, its GSO is the orthogonal vectors $\widetilde{\mathbf{B}} = (\widetilde{\mathbf{b}}_1, \ldots, \widetilde{\mathbf{b}}_d)$ with*

$$\widetilde{\mathbf{b}}_i = \mathbf{b}_i - \sum_{j<i} \mu_{ij}\widetilde{\mathbf{b}}_j \text{ where } \mu_{ij} = \frac{\langle \mathbf{b}_i, \widetilde{\mathbf{b}}_j \rangle}{\langle \widetilde{\mathbf{b}}_j, \widetilde{\mathbf{b}}_j \rangle}.$$

We define $\mathcal{P}_{1/2}(\widetilde{\mathbf{B}}) = \widetilde{\mathbf{B}} \cdot [-1/2, 1/2)^d$. If $\mathbf{B}$ is a basis of Λ then $\mathcal{P}_{1/2}(\widetilde{\mathbf{B}})$ is a Λ-tiling of $\mathbb{R}^d$. That is, $\{\mathbf{x} + \mathcal{P}_{1/2}(\widetilde{\mathbf{B}})\}_{\mathbf{x}\in\Lambda} = \mathbb{R}^d$ (it covers) and for $\mathbf{x}, \mathbf{x}' \in \Lambda$, $(\mathbf{x} + \mathcal{P}_{1/2}(\widetilde{\mathbf{B}})) \cap (\mathbf{x}' + \mathcal{P}_{1/2}(\widetilde{\mathbf{B}})) \neq \emptyset \iff \mathbf{x} = \mathbf{x}'$ (it packs).

Definition 8 (Unitriangular). *A basis $\mathbf{B} \in \mathbb{R}^{d\times d}$ is unitriangular if it is upper triangular with leading diagonal $(1, \ldots, 1)$, or equivalently $\widetilde{\mathbf{B}} = \mathbf{I}_d$. A lattice $\Lambda \subset \mathbb{R}^d$ is unitriangular if it admits a unitriangular basis.*

Unitriangular lattices are precisely those that induce a tiling of space by unit volume hypercubes. That is, $\{\mathbf{x} + [1/2, 1/2)^d\}_{\mathbf{x}\in\Lambda}$ tiles $\mathbb{R}^d$ if and only if Λ is unitriangular. This result is due to Hajós [Haj42], see [SC89, Thm. 30]. (One direction is clear since $\mathcal{P}_{1/2}(\widetilde{\mathbf{B}}) = [-1/2, 1/2)^d$.) Note $\mathbb{Z}^d$ is unitriangular, and for any basis $\mathbf{B}$ we have $\mathbf{B} = \widetilde{\mathbf{B}}\mathbf{M}$ where $\mathbf{M} = \widetilde{\mathbf{B}}^{-1}\mathbf{B}$ is unitriangular.

Definition 9 (Babai's nearest plane [Bab86]). *On input a basis $\mathbf{B} \in \mathbb{R}^{d\times d}$ of lattice Λ and a vector $\mathbf{x} \in \mathbb{R}^d$, $\mathbf{z} = \mathsf{NearestPlane}(\mathbf{B}, \mathbf{x})$ is a polynomial time algorithm that returns the unique lattice point $\mathbf{z} \in \Lambda \cap (\mathbf{x} + \mathcal{P}_{1/2}(\tilde{\mathbf{B}}))$.*

The nearest plane algorithm admits a simple recursive description in the case of unitriangular bases, given as $\mathsf{NearestPlane}_U$ in Figure 2. Therein $\mathbf{x}_{[d-1]}$ denotes the first $(d-1)$ coordinates of $\mathbf{x} \in \mathbb{R}^d$ and $\mathbf{B}_{[d-1]}$ denotes the upper left $(d-1)\times(d-1)$ submatrix of $\mathbf{B} \in \mathbb{R}^{d\times d}$.

$\mathsf{NearestPlane}_U(\mathbf{B} \in \mathbb{R}^{d\times d}, \mathbf{x} \in \mathbb{R}^d)$

$a = \lfloor x_d \rceil$

$\mathbf{x}' = \mathbf{x} - a \cdot \mathbf{b}_d$

$\mathbf{z}' = \mathsf{NearestPlane}_U(\mathbf{B}_{[d-1]}, \mathbf{x}'_{[d-1]})$

$\mathbf{z} = (\mathbf{z}', 0) + a \cdot \mathbf{b}_d$

return z

Fig. 2: The nearest plane algorithm specialised to unitriangular bases $\mathbf{B} \in \mathbb{R}^{d\times d}$. In the base case $d = 1$ the algorithm simply computes and returns a. Note $(\mathbf{z}', 0)$ denotes the column vector $\mathbf{z}'$ with a zero appended below.

Definition 10. *For $\mathbf{A} \in \mathbb{Z}_q^{n\times m}$ define*

- *the code $\mathcal{C}^\perp_{\mathbf{A}} = \{\mathbf{x} \in \mathbb{Z}_q^m : \mathbf{A}\mathbf{x} = \mathbf{0}\} \subseteq \mathbb{Z}_q^m$, and*
- *the q-ary lattice $\Lambda_q^\perp(\mathbf{A}) = \phi(\mathcal{C}^\perp_{\mathbf{A}}) + q\mathbb{Z}^m \subseteq \mathbb{Z}^m$, where $\phi\colon \mathbb{Z}_q^m \to \mathbb{Z}^m$ is any modulo q lift.*

When q is prime let $k = \operatorname{rank}_{\mathbb{Z}_q}(\mathcal{C}^\perp_{\mathbf{A}}) \in \{0\} \cup [n]$ be the rank of $\mathcal{C}^\perp_{\mathbf{A}}$ and call it an (m, k, q)-code. It is well known that the ensemble of random q-ary lattices attained by lifting uniformly random (m, n, q)-codes admits an averaging property [Loe97] which can be seen as a discrete analogue of the celebrated Minkowski–Hlawka theorem. The following variant considers uniformly sampled $\mathbf{A}$, which may give an (m, k, q)-code with $k < n$.

Lemma 9 (Lem. 11 of full version [PS24]). *Let q be prime and $m \geq n$ be integers. Let $f\colon \mathbb{Z}_q^m \to \mathbb{R}$ and $\mathbf{A} \sim \mathsf{U}(\mathbb{Z}_q^{n\times m})$. Then*

$$\mathbb{E}_{\mathbf{A}}\left[\sum_{\mathbf{v}\in\mathcal{C}^\perp_{\mathbf{A}}\setminus\{\mathbf{0}\}} f(\mathbf{v})\right] = q^{-n} \sum_{\mathbf{v}\in\mathbb{Z}_q^m\setminus\{\mathbf{0}\}} f(\mathbf{v})$$

Corollary 1. *For* $\mathbf{x} \in \mathbb{Z}^m$ *and* $\mathbf{A} \sim \mathsf{U}(\mathbb{Z}_q^{n \times m})$,

$$\mathbb{P}\left[\mathbf{x} \in \Lambda_q^\perp(\mathbf{A})\right] = \begin{cases} 1 & \mathbf{x} \in q\mathbb{Z}^m, \\ q^{-n} & \mathbf{x} \notin q\mathbb{Z}^m. \end{cases}$$

Proof. First, $q\mathbb{Z}^m \subseteq \Lambda_q^\perp(\mathbf{A})$ for all $\mathbf{A}$ by definition. Then, for $\mathbf{x} \notin q\mathbb{Z}^m$ set $f(\mathbf{v}) = \mathbb{1}(\mathbf{v} = [\mathbf{x}]_q)$ in Lemma 9. □

We note that $\mathrm{vol}(\Lambda_q^\perp(\mathbf{A})) \in \{1, q, \ldots, q^n\}$ for $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$.

Producing trapdoor bases for q-ary lattices, or more specifically generating near uniform $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$ alongside a basis $\mathbf{T}$ of $\Lambda_q^\perp(\mathbf{A})$ with minimal $\|\widetilde{\mathbf{T}}\|$, has seen much interest [Ajt99,AP11,MP12]. We use the following implicit result for simplicity, improvements are available e.g. [MP12].

Theorem 1 ([AP11, Thm. 3.2]). *Let* $n, m, q \in \mathbb{N}^+$ *such that* $6n \log_2 q \leq m \in O(n \log q)$ *and* $2 \leq q \leq 2^{\mathrm{poly}(n)}$. *There is an efficient randomised algorithm that outputs* $(\mathbf{A}, \mathbf{T}) \in \mathbb{Z}_q^{n \times m} \times \mathbb{Z}^{m \times m}$ *such that* $\Delta(\mathbf{A}, \mathsf{U}(\mathbb{Z}_q^{n \times m}))$ *is negligible,* $\Lambda(\mathbf{T}) = \Lambda_q^\perp(\mathbf{A})$, *and* $\|\widetilde{\mathbf{T}}\| \in O(\sqrt{m})$ *with overwhelming probability.*

2.5 Short integer solutions

The short integer solutions problem was introduced in [Ajt96]. The ℓ_p norm variants for $p \in [1, \infty]$ are defined below.

Definition 11 (Short integer solutions (SIS^p)). *Let* $n \in \mathbb{N}$ *and* m, q, β *be parametrised by it, let* $p \in [1, \infty]$. *A* $\mathrm{SIS}^p_{n,m,q,\beta}$ *instance is* $\mathbf{A} \sim \mathsf{U}(\mathbb{Z}_q^{n \times m})$. *A solution is* $\mathbf{x} \in (\Lambda_q^\perp(\mathbf{A}) \cup B_m^p(\beta)) \setminus \{\mathbf{0}\}$. *If* p *is omitted then* $p = 2$.

Worst case lattice problems (under ℓ_p for $p \in [1, \infty]$, see [Pei08, Thm. 6.4]) reduce to well parametrised SIS. In particular, we have the following.

Theorem 2 (Implicit in [MP13, Thm. 4]). *Let* $\delta > 0$. *Let* $n \in \mathbb{N}$ *and* m, q, β *be parametrised by it with* $m \in \mathrm{poly}(n)$ *and* $q \geq \beta \cdot n^\delta$ *prime. Let* $\gamma \in \max\{1, \beta^2/q\} \cdot O(\beta\sqrt{n}) \cdot \omega(\sqrt{\log n})$. *Then there is an efficient reduction from* SIVP_γ *to* $\mathrm{SIS}_{n,m,q,\beta}$.

Here SIVP_γ is a worst case problem over lattices that is conjectured to be hard for any $\gamma \in \mathrm{poly}(n)$. As such, provided $\beta \in \mathrm{poly}(n)$, as it will be in our constructions with small bounded degree, and the hypothesis of Theorem 2 is satisfied, we assume the respective SIS instances are hard. In particular, all else being fixed, it is a matter of choosing a sufficiently large q. The following immediate lemma allows the use of such worst case to average case reductions for SIS^p with $p \neq 2$.

Lemma 10. *Let* $n \in \mathbb{N}$ *and* m, q, β *be parametrised by it. Let* $p_1, p_2 \in [1, \infty]$ *with* $p_1 \leq p_2$. *Then*

1. $\mathrm{SIS}^{p_1}_{n,m,q,\beta}$ *is not easier than* $\mathrm{SIS}^{p_2}_{n,m,q,\beta}$,
2. $\mathrm{SIS}^{\infty}_{n,m,q,\beta}$ *is not easier than* $\mathrm{SIS}_{n,m,q,\beta\sqrt{m}}$,
3. $\mathrm{SIS}_{n,m,q,\beta}$ *is not easier than* $\mathrm{SIS}^1_{n,m,q,\beta\sqrt{m}}$, *and*
4. $\mathrm{SIS}^{\infty}_{n,m,q,\beta}$ *is not easier than* $\mathrm{SIS}^1_{n,m,q,\beta m}$.

2.6 Collision resistant preimage sampleable function families

Strongly existentially unforgeable under chosen message attack (sEUF-CMA) secure signature schemes may be built from from collision resistant preimage sampleable functions (CR-PSFs) in a stateful manner [GPV08, Prop. 6.1] or probabilistically [GPV07, Prop. 6.2]. In Definition 12 we use a random variable defined as

$$X_n(a, y) = \begin{cases} \perp & \text{if } f_a^{-1}(y) = \emptyset, \\ x \leftarrow \mathsf{SampDom}(1^n) \text{ given } f_a(x) = y & \text{otherwise.} \end{cases}$$

Definition 12 (CR-PSFs). *A* $\mathsf{CR\text{-}PSF} = (\mathsf{TrapGen}, \mathsf{SampDom}, \mathsf{SampPre})$ *is a triple of polynomial time algorithms such that*

1. *(Trapdoor generation.)* $\mathsf{TrapGen}(1^n)$ *is a randomised algorithm that outputs the pair* (a, t)*. Here* a *is a description of an efficient* $f_a \colon D_n \to R_n$ *where* D_n *and* R_n *are efficiently recognisable domains and codomains depending on* n*, and* t *is a trapdoor for* f_a*.*
2. *(Domain sampling with uniform output.)* $\mathsf{SampDom}(1^n)$ *is a randomised algorithm that outputs an element* $x \in D_n$ *with* $\Delta(f_a(x), \mathsf{U}(R_n)) \in \mathrm{negl}(n)$ *except for a negligible fraction of* $(a, \cdot) \leftarrow \mathsf{TrapGen}(1^n)$*.*
3. *(Preimage sampling with trapdoor.)* $\mathsf{SampPre}(1^n, a, t, y)$ *is a randomised algorithm such that for all but a negligible fraction of* $(a, t) \leftarrow \mathsf{TrapGen}(1^n)$*,* $\mathsf{SampPre}(1^n, a, t, y)$ *outputs an element* $x \in D_n \cup \{\perp\}$ *with* $\Delta(X_n(a, y), x) \in \mathrm{negl}(n)$ *for all* $y \in R_n$*.*
4. *(Preimage min-entropy.) For all* $y \in R_n$ *we have* $H_\infty(X_n(a, y)) \in \omega(\log n)$ *except for a negligible fraction of* $(a, \cdot) \leftarrow \mathsf{TrapGen}(1^n)$*.*
5. *(Collision resistance without trapdoor.) For any probabilistic polynomial time* $\mathcal{A}$ *the probability* $(x, x') \leftarrow \mathcal{A}(1^n, a)$ *with* $x \neq x'$ *and* $f_a(x) = f_a(x')$ *is negligible over* $(a, \cdot) \leftarrow \mathsf{TrapGen}(1^n)$ *and the coins of* $\mathcal{A}$*.*

3 Sampling unitriangular lattice points in polytopes

Kannan and Vempala [KV97] showed that, provided a polytope $P \subset \mathbb{R}^d$ has inradius $\Omega(d\sqrt{\log f})$ where f is the number of facets, one may efficiently sample a distribution with small constant statistical distance from the uniform distribution on $P \cap \mathbb{Z}^d$. The algorithm begins by forming an enlarged polytope $P' \supset P$ and assumes the ability to sample points from $\mathsf{U}(P')$. One then samples $\mathbf{x} \leftarrow \mathsf{U}(P')$ and performs a randomised rounding procedure[4] $R(\mathbf{x})$ that yields an integer point. If $R(\mathbf{x})$ does not lie in P, one samples a new $\mathbf{x}$ and repeats, and it is shown that the expected number of repetitions is constant. We build upon their results in several ways.

[4] First proposed by Raghavan and Thompson [RT87], who demonstrated its utility in algorithmic combinatorics via the probabilistic method.

Firstly, rather than only considering the integer lattice $\mathbb{Z}^d$, we extend the algorithm to allow sampling from $P \cap \Lambda$ where Λ is any unitriangular lattice. To achieve this, we generalise the randomised rounding procedure of [KV97] to obtain a randomised variant of Babai's nearest plane algorithm.

Secondly, we alter the enlarged polytope P'. Rather than constructing P' by shifting each facet of P a constant distance outwards from the origin – which does not necessarily provide an effective P' – we instead construct P' as a dilation of P. We therefore do not need to assume we can sample from $\mathsf{U}(P')$ so long as we can sample from $\mathsf{U}(P)$, by Lemma 7.

Finally, we generalise their asymptotic regime and elucidate full details of their proof techniques, presenting explicit bounds on the statistical distance and expected number of repetitions, see §3.2. This more flexible analysis of the algorithm aids its use in our cryptographic setting.

3.1 Sampling algorithm

Definition 13. *Given a basis* $\mathbf{B}$ *for a lattice* $\Lambda \subset \mathbb{R}^d$ *and a target* $\mathbf{x} \in \mathbb{R}^d$*, we define the randomised nearest plane algorithm* $\mathsf{RNP}(\mathbf{B}, \mathbf{x})$ *as*

1. *sample* $\mathbf{e} \leftarrow \mathsf{U}(\mathcal{P}_{1/2}(\widetilde{\mathbf{B}}))$,
2. *return* $\mathbf{z} = \mathsf{NearestPlane}(\mathbf{B}, \mathbf{x} + \mathbf{e})$.

In the case of the integer lattice with $\mathbf{B} = \mathbf{I}_d$, our algorithm RNP coincides exactly with the randomised rounding procedure from $\mathbb{R}^d$ to $\mathbb{Z}^d$ used in [KV97].

Lemma 11. *Let* $\mathbf{B} \in \mathbb{R}^{d \times d}$ *be a basis for* $\Lambda \subset \mathbb{R}^d$*. For any* $\mathbf{z} \in \Lambda$ *and* $\mathbf{x} \in \mathbb{R}^d$

$$\mathbb{P}\left[\mathsf{RNP}(\mathbf{B}, \mathbf{x}) = \mathbf{z}\right] = \begin{cases} \prod_{i=1}^{d} \left(1 - \left|(\widetilde{\mathbf{B}}^{-1}(\mathbf{x} - \mathbf{z}))_i\right|\right) & \text{if } \mathbf{x} \in \mathbf{z} + 2 \cdot \mathcal{P}_{1/2}(\widetilde{\mathbf{B}}), \\ 0 & \text{otherwise.} \end{cases}$$

Proof. First, $\mathsf{RNP}(\mathbf{B}, \mathbf{x}) = \mathbf{z}$ if and only if $\mathbf{x} + \mathbf{e} \in \mathbf{z} + \mathcal{P}_{1/2}(\widetilde{\mathbf{B}})$. Equivalently, $-\widetilde{\mathbf{B}}^{-1}\mathbf{e} \in \widetilde{\mathbf{B}}^{-1}(\mathbf{x} - \mathbf{z}) + (-1/2, 1/2]^d$. Given $\widetilde{\mathbf{B}}^{-1}\mathbf{e} \sim \mathsf{U}([-1/2, 1/2))^d$, the result follows from computing the per coordinate probability

$$\mathbb{P}\left[\mathsf{U}((-1/2, 1/2]) \in (\widetilde{\mathbf{B}}^{-1}(\mathbf{x} - \mathbf{z}))_i + (-1/2, 1/2]\right]. \qquad \square$$

We use RNP as a subroutine to define an algorithm that samples from the intersection of a lattice Λ and effective polytope P, see pseudocode in Figure 3.

Definition 14. *Given a basis* $\mathbf{B}$ *for a lattice* Λ*, an effective polytope* P*, a real parameter* $\delta \in (0, 1)$ *and a centre of dilation* $\mathbf{c} \in P$*, define* $\mathsf{SampPoly}(\mathbf{B}, P, \mathbf{c}, \delta)$ *as the following procedure which returns a point in* $P \cap \Lambda$*.*

1. *Form the enlarged polytope* $P' = (1 + \delta) \cdot (P - \mathbf{c}) + \mathbf{c}$.

$\mathsf{SampPoly}(\mathbf{B}, P, \mathbf{c}, \delta)$	$\mathsf{RNP}(\mathbf{B}, \mathbf{x})$
$P' = (1+\delta) \cdot (P - \mathbf{c}) + \mathbf{c}$	$\mathbf{e} \leftarrow \mathsf{U}(\mathcal{P}_{1/2}(\widetilde{\mathbf{B}}))$
repeat	$\mathbf{z} = \mathsf{NearestPlane}(\mathbf{B}, \mathbf{x} + \mathbf{e})$
$\quad \mathbf{x} \leftarrow \mathsf{U}(P')$	**return z**
$\quad \mathbf{z} \leftarrow \mathsf{RNP}(\mathbf{B}, \mathbf{x})$	
until $\mathbf{z} \in P$	
return z	

Fig. 3: Pseudocode for SampPoly and the subroutine RNP.

2. *Sample* $\mathbf{x} \leftarrow \mathsf{U}(P')$ *and let* $\mathbf{z} = \mathsf{RNP}(\mathbf{B}, \mathbf{x})$.
3. *If* $\mathbf{z} \in P$ *then return it. Else repeat the previous step.*

P' is constructed by dilation centred at $\mathbf{c}$ by a factor $(1+\delta)$, so P' is effective by Lemma 7. Hence, each step of SampPoly is efficient, although *a priori* the repeat context may never terminate. It remains to analyse the expected number of iterations, and the output distribution achieved.

3.2 Analysis for unitriangular lattices

We show the techniques of [KV97] can be used to analyse SampPoly applied to unitriangular lattices. Lemma 12 bounds the probability assigned to each lattice point in P. Lemma 13 bounds the probability that an iteration succeeds. Combining these, Theorem 3 bounds the statistical distance between the output and the uniform distribution on $P \cap \Lambda$ and gives the expected number of iterations.

Notation. Throughout this subsection $\Lambda \subset \mathbb{R}^d$ is a unitriangular lattice with unitriangular basis $\mathbf{B}$, and $P = (\mathbf{M}, \mathbf{v})$ is a polytope in $\mathbb{R}^d$ with f facets. We let $\delta \in (0, 1)$, $\mathbf{c} \in P$, and let P' be the enlarged polytope $P' = (1+\delta)(P - \mathbf{c}) + \mathbf{c}$.

Lemma 12 (Adapted from [KV97, Thm. 2]). *For any* $\mathbf{z} \in P \cap \Lambda$,

$$\frac{1-\varepsilon}{\mathrm{vol}(P')} \leq \mathop{\mathbb{P}}_{\mathbf{x} \leftarrow \mathsf{U}(P')} [\mathsf{RNP}(\mathbf{B}, \mathbf{x}) = \mathbf{z}] \leq \frac{1}{\mathrm{vol}(P')}$$

where $\varepsilon = f \cdot \exp(-\delta^2\, r(P; \mathbf{c})^2/2)$.

Proof. Write $P = (\mathbf{M}, \mathbf{v})$, then $P' = (\mathbf{M}, \mathbf{v}')$ with $\mathbf{v}' = \mathbf{v} + \delta(\mathbf{v} - \mathbf{M}^t\mathbf{c})$ by Lemma 2. Let Y be a random variable on $\mathbb{R}^d$ with components $(Y_i)_{i=1}^d$ independently identically distributed on $[-1, 1]$ according to the probability density function $y \mapsto 1 - |y|$. Let $g_{\mathbf{z}}$ be the probability density function of $\mathbf{z} + Y$.

By Lemma 11, and the fact $\widetilde{\mathbf{B}} = \mathbf{I}_d$,

$$\begin{aligned}\mathop{\mathbb{P}}_{\mathbf{x}\leftarrow\mathsf{U}(P')}[\mathsf{RNP}(\mathbf{B},\mathbf{x}) = \mathbf{z}] &= \frac{1}{\mathrm{vol}(P')}\int_{P'\cap(\mathbf{z}+B_d^\infty)}\prod_{i=1}^n (1-|x_i-z_i|)\,\mathrm{d}\mathbf{x}\\ &= \frac{1}{\mathrm{vol}(P')}\int_{P'\cap(\mathbf{z}+B_d^\infty)} g_{\mathbf{z}}(\mathbf{x})\,\mathrm{d}\mathbf{x}\\ &= \frac{1}{\mathrm{vol}(P')}\mathbb{P}\left[\mathbf{z}+Y\in P'\right].\end{aligned}$$

The upper bound in the statement follows. We now prove the lower bound. By definition, $\mathbf{z}+Y\in P'$ if and only if $\mathbf{M}^tY \le \mathbf{v}'-\mathbf{M}^t\mathbf{z}$. Since $\mathbf{z}\in P$ we obtain the sufficient condition $\mathbf{M}^tY\le\mathbf{v}'-\mathbf{v}$, thus

$$\begin{aligned}\mathbb{P}\left[\mathbf{z}+Y\in P'\right] &\ge \mathbb{P}\left[\mathbf{M}^tY\le\mathbf{v}'-\mathbf{v}\right]\\ &= 1-\mathbb{P}\left[\langle\mathbf{m}_1,Y\rangle > v_1'-v_1\vee\cdots\vee\langle\mathbf{m}_f,Y\rangle > v_f'-v_f\right]\\ &\ge 1-\sum_{i=1}^f \mathbb{P}\left[\langle\mathbf{m}_i,Y\rangle\ge v_i'-v_i\right].\end{aligned}$$

It suffices to show each summand is at most $e^{-\delta^2 r(P;\mathbf{c})^2/2}$. Let $j\in\{0\}\cup[d]$ and for each i consider the random variables $Z_j=\sum_{k=1}^j m_{ik}Y_k$ where $\mathbf{m}_i = (m_{i1}\cdots m_{id})^t$. These form a martingale with $|Z_j-Z_{j-1}|\le m_{ij}$. Using Lemma 8 to bound $Z_d=\langle\mathbf{m}_i,Y\rangle$ we obtain

$$\mathbb{P}\left[\langle\mathbf{m}_i,Y\rangle\ge v_i'-v_i\right]\le\exp\left(\frac{-(v_i'-v_i)^2}{2\|\mathbf{m}_i\|^2}\right)\le e^{-\delta^2 r(P;\mathbf{c})^2/2}$$

where we used that

$$\frac{v_i'-v_i}{\|\mathbf{m}_i\|}=\frac{\delta\cdot(\mathbf{v}+\mathbf{M}^t\mathbf{c})_i}{\|\mathbf{m}_i\|}\le\delta\cdot r(P;\mathbf{c})$$

by Lemma 3. □

Lemma 13 (Adapted from [KV97, Lem. 1]). *The probability any iteration within* $\mathsf{SampPoly}(\mathbf{B},P,\mathbf{c},\delta)$ *succeeds, called the acceptance rate, is*

$$\mathop{\mathbb{P}}_{\mathbf{x}\leftarrow\mathsf{U}(P')}[\mathsf{RNP}(\mathbf{B},\mathbf{x})\in P]\ge(1-\varepsilon)^2\left(\frac{1-\delta}{1+\delta}\right)^d$$

where $\varepsilon = f\cdot\exp(-\delta^2\,r(P;\mathbf{c})^2/2)$.

Proof. Note $\delta\in(0,1)$ and consider the smaller polytope $P''=(1-\delta)(P-\mathbf{c})+\mathbf{c}$, that is, a dilation centred at $\mathbf{c}$ by a factor $1-\delta\in(0,1)$.

We claim for any $\mathbf{x}\in P''$ (or in fact any $\mathbf{x}\in\mathbb{R}^d$) we have $\mathsf{RNP}(\mathbf{B},\mathbf{x})=\mathbf{x}+Y$ where $Y=(Y_i)_{i=1}^d$ is a random variable with components satisfying

$$\begin{aligned}&\mathbb{E}[Y_d]=0,\\ &\mathbb{E}[Y_i\mid Y_d,Y_{d-1},\ldots,Y_{i+1}]=0\quad 1\le i\le d-1.\end{aligned}$$

For unitriangular lattices we have $\mathsf{RNP}(\mathbf{B}, \mathbf{x}) = \mathsf{NearestPlane}_U(\mathbf{B}, \mathbf{x}+\mathbf{e})$ with $\mathbf{e} \leftarrow \mathsf{U}([-1/2, 1/2))^d$ by definition, and we use the description of $\mathsf{NearestPlane}$ for unitriangular lattices in Figure 2. Now

$$\mathsf{NearestPlane}_U(\mathbf{B}, \mathbf{x}+\mathbf{e})_d = a = \lfloor x_d + e_d \rceil \sim \lfloor x_d \rfloor + \mathrm{Ber}(x_d - \lfloor x_d \rfloor)$$

which gives the first statement. Further, exploiting the recursive structure,

$$\begin{aligned}
&\mathsf{NearestPlane}_U(\mathbf{B}, \mathbf{x}+\mathbf{e})_{[d-1]} \\
&= \mathsf{NearestPlane}_U(\mathbf{B}_{[d-1]}, (\mathbf{x} - a \cdot \mathbf{b}_d)_{[d-1]} + \mathbf{e}_{[d-1]}) + a \cdot \mathbf{b}_d \\
&= \mathsf{RNP}(\mathbf{B}_{[d-1]}, (\mathbf{x} - a \cdot \mathbf{b}_d)_{[d-1]}) + a \cdot \mathbf{b}_d
\end{aligned}$$

and so the remaining statements in the claim follow by induction on d.

Hence, $Z_j = \sum_{k=1}^{j} m_{i,d+1-k} Y_{d+1-k}$ forms a martingale – note the reversed ordering – and thus by a similar computation as in Lemma 12,

$$\mathop{\mathbb{P}}_{\mathbf{x} \leftarrow \mathsf{U}(P'')} [\mathsf{RNP}(\mathbf{B}, \mathbf{x}) \in P] = \mathbb{P}\left[\mathbf{M}^t Y \leq \mathbf{v} - \mathbf{v}''\right] \geq 1 - \varepsilon.$$

On the other hand, we have by a union bound

$$\mathop{\mathbb{P}}_{\mathbf{x} \leftarrow \mathsf{U}(P'')} [\mathsf{RNP}(\mathbf{B}, \mathbf{x}) \in P] \leq \frac{|P \cap \Lambda|}{\mathrm{vol}(P'')}$$

and thus chaining the two inequalities above,

$$|P \cap \Lambda| \geq (1 - \varepsilon) \cdot \mathrm{vol}(P'')\,. \tag{1}$$

We use this alongside Lemma 12 to bound the probability in the statement,

$$\begin{aligned}
\mathop{\mathbb{P}}_{\mathbf{x} \leftarrow \mathsf{U}(P')} [\mathsf{RNP}(\mathbf{B}, \mathbf{x}) \in P] &= \sum_{\mathbf{z} \in P \cap \Lambda} \mathop{\mathbb{P}}_{\mathbf{x} \leftarrow \mathsf{U}(P')} [\mathsf{RNP}(\mathbf{B}, \mathbf{x}) = \mathbf{z}] \\
&\geq \frac{|P \cap \Lambda|}{\mathrm{vol}(P')} \cdot (1 - \varepsilon) \qquad (2) \\
&\geq \frac{\mathrm{vol}(P'')}{\mathrm{vol}(P')} \cdot (1 - \varepsilon)^2.
\end{aligned}$$

We conclude with

$$\frac{\mathrm{vol}(P'')}{\mathrm{vol}(P')} = \left(\frac{1 - \delta}{1 + \delta}\right)^d$$

since P' and P'' are dilations of P by $1 + \delta$ and $1 - \delta$ respectively. □

Theorem 3. *Let $\varepsilon = f \cdot \exp(-\delta^2\, r(P; \mathbf{c})^2/2)$.*

1. *The statistical distance to uniform satisfies*

$$\Delta\Big[\mathsf{SampPoly}(\mathbf{B}, P, \mathbf{c}, \delta),\ \mathsf{U}(P \cap \Lambda)\Big] \leq \frac{1}{2} \frac{\varepsilon}{1 - \varepsilon}.$$

2. The expected number of iterations is at most

$$M = \frac{1}{(1-\varepsilon)^2}\left(\frac{1+\delta}{1-\delta}\right)^d.$$

Proof. For every $\mathbf{z} \in P \cap \Lambda$,

$$\begin{aligned}\mathbb{P}\left[\mathsf{SampPoly}(\mathbf{B}, P, \mathbf{c}, \delta) = \mathbf{z}\right] &= \mathop{\mathbb{P}}_{\mathbf{x}\leftarrow\mathsf{U}(P')}\left[\mathsf{RNP}(\mathbf{B}, \mathbf{x}) = \mathbf{z} \mid \mathsf{RNP}(\mathbf{B}, \mathbf{x}) \in P\right] \\ &= \frac{\mathbb{P}\left[\mathsf{RNP}(\mathbf{B}, \mathbf{x}) = \mathbf{z}\right]}{\mathbb{P}\left[\mathsf{RNP}(\mathbf{B}, \mathbf{x}) \in P\right]}\end{aligned}$$

where by Lemma 12 the numerator lies in $[1-\varepsilon, 1]\cdot\mathrm{vol}(P')^{-1}$ and the denominator lies in $[1-\varepsilon, 1]\cdot|P\cap\Lambda|\mathrm{vol}(P')^{-1}$. A routine calculation now gives the first statement. The second statement is immediate from Lemma 13. □

3.3 A point counting bound

The techniques used in our analysis give a bound on the size of the intersection between any (not necessarily effective) polytope and any unitriangular lattice. This is a generalisation of [KV97, Thm. 3] which may be of independent interest in convex and discrete geometry, and is used later to prove our preimage sampling has sufficient min entropy.

Theorem 4. *Let $\Lambda \subset \mathbb{R}^d$ be a unitriangular lattice. Let $P \subset \mathbb{R}^d$ be a polytope with f facets and inradius $r = r(P)$. For every $\delta \in (0,1)$ define*

$$\varepsilon = f \cdot \exp\left(-\frac{\delta^2 r^2}{2}\right), \quad M = \frac{1}{(1-\varepsilon)^2}\left(\frac{1+\delta}{1-\delta}\right)^d.$$

Then the volume of P and the number of lattice points in P are a multiplicative factor at most $M(1-\varepsilon)$ apart,

$$\frac{1}{M(1-\varepsilon)}|P\cap\Lambda| \le \mathrm{vol}(P) \le M(1-\varepsilon)|P\cap\Lambda|.$$

Proof. Resume the notation in the proof of Lemma 13, and note we are free to choose $\mathbf{c}$ such that $r(P;\mathbf{c})$ is maximised, i.e. is equal to the inradius. From (1)

$$|P\cap\Lambda| \ge (1-\varepsilon)\mathrm{vol}(P'') \ge (1-\varepsilon)\frac{\mathrm{vol}(P)\mathrm{vol}(P'')}{\mathrm{vol}(P')} \ge \frac{1}{M(1-\varepsilon)}\mathrm{vol}(P).$$

Furthermore, from (2) we have

$$1 \ge \mathop{\mathbb{P}}_{\mathbf{x}\leftarrow\mathsf{U}(P')}\left[\mathsf{RNP}(\mathbf{x}) \in P\right] \ge \frac{|P\cap\Lambda|}{\mathrm{vol}(P')}\cdot(1-\varepsilon)$$

and thus

$$|P\cap\Lambda| \le \frac{\mathrm{vol}(P')}{1-\varepsilon} \le \frac{1}{1-\varepsilon}\frac{\mathrm{vol}(P')\mathrm{vol}(P)}{\mathrm{vol}(P'')} \le M(1-\varepsilon)\mathrm{vol}(P).$$ □

Remark 1. We generalise Theorem 4 and the analysis of §3.2 from unitriangular lattices to arbitrary lattices $\Lambda = \Lambda(\mathbf{B})$ by first applying the transformation $\widetilde{\mathbf{B}}^{-1}$ to Λ, P and $\mathbf{c}$; see Theorem 5.

4 Trapdoor sampling for polytopes

We introduce trapdoor samplers for polytopes in the sense of sampling a point of an effective polytope P that lies within some arbitrary coset $\mathbf{v}+\Lambda$ of lattice Λ. For example, if $P = B_m^\infty(\beta)$ and $\Lambda = \Lambda_q^\perp(\mathbf{A})$ for some $\mathbf{A} \in \mathbb{Z}_q^{n\times m}$ then sampling from $P \cap (\mathbf{v}+\Lambda)$ yields a solution for $\mathbf{Ax} = \mathbf{Av}$ with $\|\mathbf{x}\|_\infty \leq \beta$, or an $\mathrm{ISIS}_{n,m,q,\beta}^\infty$ solution for target $\mathbf{t} = \mathbf{Av}$. How close to $\mathsf{U}(P \cap (\mathbf{v}+\Lambda))$ we are able to sample, and how efficiently, is detailed in Theorem 5 and Corollary 2, with the latter clarifying the required $r(P)$ with respect to $\|\widetilde{\mathbf{B}}\|$ of known Λ basis $\mathbf{B}$. Note that in both we assume we know a point $\mathbf{c}$ for which $r(P) = r(P;\mathbf{c})$ as this will often be the case in applications, however both statements can be adapted to any $\mathbf{c}$ for which a lower bound on $r(P;\mathbf{c})$ is known.

Theorem 5. *Let Λ be a lattice with basis $\mathbf{B}$ and let $\mathbf{v} \in \mathbb{R}^d$ define an arbitrary coset. Let $P \subset \mathbb{R}^d$ be an effective polytope with f facets and $\mathbf{c} \in P$ realise the inradius of P, $r(P) = r(P;\mathbf{c})$. Let $\delta \in (0,1)$ and define*

$$\varepsilon = f \cdot \exp\left(-\frac{1}{2}\,\frac{\delta^2 \cdot r(P)^2}{\|\widetilde{\mathbf{B}}\|^2}\right).$$

Let $\mathcal{D}$ be distributed according to

$$\mathbf{v} + \mathsf{SampPoly}(\mathbf{B}, P-\mathbf{v}, \mathbf{c}-\mathbf{v}, \delta)$$

which returns a point in $P \cap (\mathbf{v}+\Lambda)$.

1. *The statistical distance to uniform satisfies*

$$\Delta\big[\mathcal{D}, \mathsf{U}(P \cap (\mathbf{v}+\Lambda))\big] \leq \frac{1}{2}\frac{\varepsilon}{1-\varepsilon}.$$

2. *The expected number of iterations within* SampPoly *is at most*

$$M = \frac{1}{(1-\varepsilon)^2}\left(\frac{1+\delta}{1-\delta}\right)^d.$$

3. *The cardinality of* $\mathrm{supp}(\mathcal{D}) = P \cap (\mathbf{v}+\Lambda)$ *satisfies*

$$\frac{1}{M(1-\varepsilon)}|P \cap (\mathbf{v}+\Lambda)| \leq \frac{\mathrm{vol}(P)}{\mathrm{vol}(\Lambda)} \leq M(1-\varepsilon)|P \cap (\mathbf{v}+\Lambda)|.$$

Remark 2. In particular, if $\varepsilon \in \mathrm{negl}(d)$ then $\mathcal{D}$ is statistically close to uniform. If further $\delta \in O(\log d/d)$ then $M \in \mathrm{poly}(d)$ and if instead $\delta \in O(1/d)$ then $M \in O(1)$. In either case, by a standard argument, the algorithm terminates in at most $M \cdot \omega(\log d)$ iterations with overwhelming probability. For guaranteed termination in polynomial time, one can fix a maximum number of iterations $T_{\max} \in M \cdot \omega(\log d)$ and output $\perp$ or an arbitrary known point in $P \cap (\mathbf{v}+\Lambda)$ if no iteration succeeded. This introduces an additional negligible statistical distance of at most $(1-1/M)^{T_{\max}}$.

Remark 3. To make ε small we want small f and $\|\widetilde{\mathbf{B}}\|$, and large δ and $r(P)$. To make M small we want small ε and δ. Therefore, only our selection of $\delta \in (0,1)$ is in tension and to achieve small ε and M simultaneously requires P to have sufficiently good geometric properties, see Corollary 2 below.

Proof. One can check SampPoly and its subprocedure RNP are compatible with conjugation by any $\mathbf{S} \in \mathrm{GL}_d(\mathbb{R})$: applying $\mathbf{S}$ to all geometric inputs, running the algorithm, and multiplying the output by $\mathbf{S}^{-1}$ is equivalent to running the algorithm on the original inputs

$$\mathsf{SampPoly}(\mathbf{B}, P, \mathbf{c}, \delta) \sim \mathbf{S}^{-1} \cdot \mathsf{SampPoly}(\mathbf{SB}, \mathbf{S}P, \mathbf{Sc}, \delta).$$

Hence,

$$\begin{aligned}&\mathsf{SampPoly}(\mathbf{B}, P - \mathbf{v}, \mathbf{c} - \mathbf{v}, \delta)\\ &\qquad \sim \widetilde{\mathbf{B}} \cdot \mathsf{SampPoly}(\widetilde{\mathbf{B}}^{-1}\mathbf{B}, \widetilde{\mathbf{B}}^{-1}(P - \mathbf{v}), \widetilde{\mathbf{B}}^{-1}(\mathbf{c} - \mathbf{v}), \delta),\end{aligned}$$

recalling $\widetilde{\mathbf{B}}^{-1}\mathbf{B}$ is unitriangular. Thus to prove the first two statements we apply the result of Theorem 3 with

$$\begin{aligned}&r(\widetilde{\mathbf{B}}^{-1}(P - \mathbf{v}); \widetilde{\mathbf{B}}^{-1}(\mathbf{c} - \mathbf{v}))\\ &= r(\widetilde{\mathbf{B}}^{-1}P; \widetilde{\mathbf{B}}^{-1}\mathbf{c}) &(3)\\ &\geq r(P; \mathbf{c})/s_1(\widetilde{\mathbf{B}}) &(4)\\ &= r(P)/\|\widetilde{\mathbf{B}}\| &(5)\end{aligned}$$

where (3) holds by translation, (4) by Lemma 4, and (5) since $r(P) = r(P; \mathbf{c})$ by hypothesis and $s_1(\widetilde{\mathbf{B}}) = \|\widetilde{\mathbf{B}}\|$ because $\widetilde{\mathbf{B}}$ has orthogonal columns.

The third statement follows from Theorem 4 and noting $P' = \widetilde{\mathbf{B}}^{-1}(P - \mathbf{v})$ has $\mathrm{vol}(P') = \det(\widetilde{\mathbf{B}}^{-1})\mathrm{vol}(P) = \mathrm{vol}(P)/\mathrm{vol}(\Lambda)$ and $r(P') \geq r(P)/\|\widetilde{\mathbf{B}}\|$ by the above. □

Corollary 2. *Let Λ be a lattice with basis $\mathbf{B}$ and let $\mathbf{v} \in \mathbb{R}^d$ define an arbitrary coset. Let $P \subset \mathbb{R}^d$ be an effective polytope with f facets and $\mathbf{c} \in P$ realise the inradius of P, $r(P) = r(P; \mathbf{c})$.*

There exists a $\delta \in (0,1)$ such that Theorem 5 yields a polynomial time algorithm for sampling negligibly statistically close to $\mathsf{U}(P \cap (\mathbf{v} + \Lambda))$ if either

1. *$r(P) \in \|\widetilde{\mathbf{B}}\| \cdot \omega(d/\sqrt{\log d})$ and $f \in \mathrm{poly}(d)$, or*
2. *$r(P) \in \|\widetilde{\mathbf{B}}\| \cdot \Omega(d\sqrt{\log f}/\log d)$ and $f \notin \mathrm{poly}(d)$.*

Furthermore, the expected number of iterations M within SampPoly *is constant if either*

1. *$r(P) \in \|\widetilde{\mathbf{B}}\| \cdot \omega(d\sqrt{\log d})$ and $f \in \mathrm{poly}(d)$, or*
2. *$r(P) \in \|\widetilde{\mathbf{B}}\| \cdot \Omega(d\sqrt{\log f})$ and $f \notin \mathrm{poly}(d)$.*

Proof. Let $r = r(P)$ throughout. For the first part, using Remark 2, it suffices to choose $\delta \in O(\log d/d)$ while ensuring $\varepsilon \in \mathrm{negl}(d)$.

If the first condition is met any $\delta = C \log d/d$ for $C > 0$ has $\delta \in O(\log d/d)$ and gives $\varepsilon = f \cdot e^{-\delta^2 r^2/2\|\widetilde{\mathbf{B}}\|^2} \in \mathrm{poly}(d) \cdot e^{-\omega(\log d)} = \mathrm{negl}(d)$.

If instead the second condition is met, write $r \in \|\widetilde{\mathbf{B}}\| \cdot \Omega_1(d) \cdot d\sqrt{\log f}/\log d$ for some $\Omega_1 \in \Omega(1)$. Let $\nu > 0$ and $\delta = O_1(d) \cdot \log d/d$ for any $O_1(d) \in O(1)$ with $O_1(d) \geq \sqrt{2}/\Omega_1(d) + \nu$. Noting that $f \notin \mathrm{poly}(d)$ implies $\log f \in \omega(\log d)$ we have

$$\frac{\delta^2 r^2}{2\|\widetilde{\mathbf{B}}\|^2} - \log f \in \Omega(\log f) \subseteq \omega(\log d) \implies \varepsilon \in \mathrm{negl}(d).$$

The second part follows by a similar analysis with $\delta \in O(1/d)$. □

4.1 Comparisons and examples

Plançon–Prest. The trapdoor sampling framework of [PP21] may be applied to sample from $\mathsf{U}(P \cap \Lambda)$ for any polytope P and lattice $\Lambda = \Lambda(\mathbf{B})$. The algorithm consists of taking a dilation $P' \supset P$, sampling $\mathbf{x} \leftarrow \mathsf{U}(P')$, and then applying a *deterministic* rounding procedure (either RoundOff or NearestPlane)[5] and iterating until the result lies in P.

SampPoly is essentially a randomised version of their NearestPlane approach, replacing NearestPlane with RNP.[6] This yields two main advantages (at the cost of introducing a negligible statistical distance).

1. Our analysis of the expected number of iterations (Lemma 13) is non-heuristic whereas the analysis of [PP21, §3.2] approximates the probability of success by $|P \cap \Lambda|/|P' \cap \Lambda| \approx \mathrm{vol}(P)/\mathrm{vol}(P')$.
2. Our inradius requirements in Corollary 2 are a factor $\sqrt{\log f}/\sqrt{d}$ smaller than [PP21]: if $f \in e^{o(d)}$, e.g. $f \in \mathrm{poly}(d)$, then we may sample from smaller P, see Example 1.

The second advantage follows from the approach of [PP21] requiring for each $\mathbf{z} \in P \cap \Lambda$ the set $\mathbf{z} + \mathcal{P}_{1/2}(\widetilde{\mathbf{B}})$ to be fully contained within P', thus the dilation must shift the facets of P' a distance $\xi \geq (\sqrt{d}/2)\|\widetilde{\mathbf{B}}\|$ from the corresponding facets of P. On the other hand in Corollary 2

$$\xi = \delta \cdot r(P) \in \begin{cases} \|\widetilde{\mathbf{B}}\| \cdot \omega(\sqrt{\log d}) & f \in \mathrm{poly}(d), \\ \|\widetilde{\mathbf{B}}\| \cdot \Omega(\sqrt{\log f}) & f \notin \mathrm{poly}(d). \end{cases}$$

Smaller dilations imply smaller lower bounds on $r(P)$ for any acceptance rate.

[5] They also consider an (in general inefficient) ExactCVP oracle.

[6] RoundOff is analogously improved by sampling from $\mathbf{B}^{-1}P \cap \mathbb{Z}^d$ with SampPoly and multiplying by $\mathbf{B}$. On the integer lattice RNP simplifies to the randomised rounding of [KV97]. The resulting trapdoor sampler requires inradii as in Corollary 2 but with $s_1(\widetilde{\mathbf{B}}) = \|\widetilde{\mathbf{B}}\|$ replaced with $s_1(\mathbf{B})$, in general a factor $\sqrt{d}$ longer.

Lyubashevsky–Wichs. In [LW15] the authors use the lattice trapdoors of [MP12] to sample distributions $\mathbf{z} = (\mathbf{z}_1, \mathbf{z}_2)$ where $\mathbf{z}_1$ follows a prescribed distribution, e.g. uniform, which in turn determines the distribution of $\mathbf{z}_2$. The restricted form of $\mathbf{z}_2$ disallows sampling uniformly, i.e. with maximum entropy, from many standard polytopes. We nevertheless compare to the hypercube case in Example 1.

Gentry–Peikert–Vaikuntanathan. [GPV08, Thm. 4.1] shows one can efficiently sample from a discrete Gaussian with parameter $s \geq \|\widetilde{\mathbf{B}}\| \cdot \omega(\sqrt{\log d})$, which is concentrated within $B_d^2(r)$ with $r = s\sqrt{d} = \|\widetilde{\mathbf{B}}\| \cdot \omega(\sqrt{d \log d})$. Since $B_d^2(r)$ is not a polytope and the discrete Gaussian is not uniform, we cannot give a direct comparison, but r is $\sqrt{d/\log d}$ shorter than our minimal $r(P)$.

Example 1 (Hypercubes). Let $P = B_d^\infty(w)$ be the hypercube of width w, which is effective. We have $r(P) = w$, $\mathbf{c} = \mathbf{0}$ and $f = 2d \in \text{poly}(d)$.

By Corollary 2 we can sample $\mathsf{U}(P \cap \Lambda)$ at widths

$$w \geq \|\widetilde{\mathbf{B}}\| \cdot \omega(d\sqrt{\log d}) \quad \text{with } M \in O(1),$$
$$w \geq \|\widetilde{\mathbf{B}}\| \cdot \omega(d/\sqrt{\log d}) \quad \text{with } M \in \text{poly}(d).$$

which are a factor $\sqrt{d/\log d}$ shorter than in [PP21]. When instantiated over random q-ary lattices with trapdoor quality $\|\widetilde{\mathbf{B}}\| \in O(\sqrt{d})$ (recall Theorem 1), the first matches the width attained in [LW15],[7] but their algorithm does not sample *uniformly* from $P \cap \Lambda$.

Example 2 (Orthoplexes). Let $P = B_d^1(w)$ be the orthoplex of width w, which is effective (see §2.1). We have $r(P) = w/\sqrt{d}$, $\mathbf{c} = \mathbf{0}$ and $f = 2^d \notin \text{poly}(d)$.

By Corollary 2 we can sample $\mathsf{U}(P \cap \Lambda)$ at widths

$$w \geq \|\widetilde{\mathbf{B}}\| \cdot \Omega(d^2) \quad \text{with } M \in O(1),$$
$$w \geq \|\widetilde{\mathbf{B}}\| \cdot \Omega(d^2/\sqrt{\log d}) \quad \text{with } M \in \text{poly}(d).$$

matching the widths of [PP21] (though with provable acceptance rates).

5 New collision resistant preimage sampleable function families

In this section we provide a generic construction of a CR-PSF over q-ary lattices given an effective polytope P. In §5.1 we derive a leftover hash lemma result we require. In §5.2 we give a construction and set of conditions that allow us to realise it.

[7] In [LW15, §3.2, Example 1] the distribution has support in the hypercube of width $w = \Theta(d^{3/2}\sqrt{\lambda})$ for $\lambda \in \omega(\log d)$. There exist $\mathbf{z}$ in the support with $\|\mathbf{z}\|_\infty = w$.

5.1 Leftover hash lemma

We make use of a form of the leftover hash lemma for the multiplication by $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$ map when its argument is sampled uniformly from a finite subset of $\mathbb{Z}^m$. This subset of $\mathbb{Z}^m$ must contain sufficiently many points that are not equivalent modulo q, as any $\mathbf{x}, \mathbf{x}' \in \mathbb{Z}^m$ such that $\mathbf{x} = \mathbf{x}' \bmod q$ have $\mathbf{A}\mathbf{x} = \mathbf{A}\mathbf{x}'$. To prove our leftover hash lemma we use the following result.

Lemma 14 ([BL26, Lem. 3i]). *Let X be a random variable over a finite set S. Then $2\Delta(X, \mathsf{U}(S)) \leq \sqrt{|S| \cdot \mathrm{Coll}(X) - 1}$. Here $\mathrm{Coll}(X) = \sum_{x \in S} \mathbb{P}\left[X = x\right]^2$.*

Lemma 15. *Let n, m, q be positive integers with $m \geq n$ and q prime, $T \subset \mathbb{Z}^m$ be a finite set of points which are pairwise inequivalent modulo q, and $\mathbf{A} \sim \mathsf{U}(\mathbb{Z}_q^{n \times m})$ and $\mathbf{x} \sim \mathsf{U}(T)$. Then*

$$\Delta((\mathbf{A}, \mathbf{A}\mathbf{x}), (\mathbf{A}, \mathsf{U}(\mathbb{Z}_q^n))) \leq \frac{1}{2}\sqrt{\frac{q^n - 1}{|T|}} \ .$$

Proof. Let $S = \mathbb{Z}_q^{n \times m} \times \mathbb{Z}_q^n$ and X be the random variable of $(\mathbf{A}, \mathbf{A}\mathbf{x})$. By Lemma 14 we bound the statistical distance by $\frac{1}{2}\sqrt{q^{nm} q^n \,\mathrm{Coll}(X) - 1}$ where

$$\mathrm{Coll}(X) = \frac{1}{q^{nm}} \mathop{\mathbb{P}}_{\substack{\mathbf{A} \leftarrow \mathsf{U}(\mathbb{Z}_q^{n \times m}) \\ \mathbf{x}, \mathbf{x}' \leftarrow \mathsf{U}(T)}} \left[\mathbf{x} - \mathbf{x}' \in \Lambda_q^{\perp}(\mathbf{A})\right]$$

in this setting.[8] We consider two cases. If $\mathbf{x} = \mathbf{x}'$, which occurs with probability $1/|T|$, then $\mathbf{x} - \mathbf{x}' = \mathbf{0}$ is in the lattice with certainty. Otherwise, if $\mathbf{x} - \mathbf{x}' \neq 0$, then it lies in the lattice with probability q^{-n} by Corollary 1. Thus

$$\mathrm{Coll}(X) = \frac{1}{q^{nm}} \left(\frac{1}{|T|} + \left(1 - \frac{1}{|T|}\right)\frac{1}{q^n}\right)$$

and the result follows. □

Remark 4. Note the equivalent formulation in terms of the expected statistical distance over the randomness of $\mathbf{A}$

$$\mathop{\mathbb{E}}_{\mathbf{A} \sim \mathsf{U}(\mathbb{Z}_q^{n \times m})} \left[\Delta(\mathbf{A}\mathbf{x}, \mathsf{U}(\mathbb{Z}_q^n))\right] \leq \frac{1}{2}\sqrt{\frac{q^n - 1}{|T|}}.$$

That is, let Y be a random variable mapping $\mathbf{A}$ to $\Delta(\mathbf{A} \cdot \mathsf{U}(T), \mathsf{U}(\mathbb{Z}_q^n))$ then we have $2\mathbb{E}[Y] \leq \sqrt{(q^n - 1)/|T|}$. If $\mathbb{E}[Y] \in \mathrm{negl}(n)$ then by Markov's inequality there exist $0 < \mathrm{negl}_1, \mathrm{negl}_2 \in \mathrm{negl}(n)$ such that

$$\mathbb{P}\left[Y \geq \mathrm{negl}_1(n)\right] \leq \mathbb{E}[Y]/\,\mathrm{negl}_1(n) = \mathrm{negl}_2(n).$$

In particular, one may take $\mathrm{negl}_1 = \mathrm{negl}_2 = \sqrt{\mathbb{E}[Y]}$.

[8] See e.g. the proof of [BL26, Thm. 4].

5.2 Generic CR-PSF construction

Definition 15 (Generic CR-PSF). *Let $n \in \mathbb{N}$ parameterise $m, q \in \mathbb{N}$ with $m \geq n$ and q prime. Let $P \subset \mathbb{R}^m$ be an effective polytope, $\mathbf{c} \in P$ realise the inradius of P, $r(P) = r(P; \mathbf{c})$, and $\delta \in (0, 1)$, all implicitly parameterised by m.*

Define CR-PSF $=$ (TrapGen, SampDom, SampPre) *as follows.*

- $\mathsf{TrapGen}(1^n)$: *return* $(\mathbf{A}, \mathbf{T})$. *Here* $\mathbf{T}$ *is a trapdoor basis for* $\Lambda_q^\perp(\mathbf{A})$, *and* $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$ *describes the function*

$$f_{\mathbf{A}} \colon (P \cap \mathbb{Z}^m) \to \mathbb{Z}_q^n$$
$$\mathbf{x} \mapsto \mathbf{A}\mathbf{x}$$

- $\mathsf{SampDom}(1^n)$: *return* $\mathbf{z} \leftarrow \mathsf{SampPoly}(\mathbf{I}_d, P, \mathbf{c}, \delta)$. *Alternatively, use a bespoke sampling procedure to return a sample from* $\mathsf{U}(P \cap \mathbb{Z}^m)$.
- $\mathsf{SampPre}(1^n, \mathbf{A}, \mathbf{T}, \mathbf{u})$: *solve for any* $\mathbf{x} \in \mathbb{Z}^m$ *such that* $\mathbf{A}\mathbf{x} = \mathbf{u}$. *If no such* $\mathbf{x}$ *exists, return* $\perp$. *Else return* $\mathbf{z} \leftarrow \mathbf{x} + \mathsf{SampPoly}(\mathbf{T}, P - \mathbf{x}, \mathbf{c} - \mathbf{x}, \delta)$.

Theorem 6. *Inherit the notation and conditions of Definition 15 where also* $6n \log_2 q \leq m \in O(n \log q)$, $q \leq 2^{\mathrm{poly}(n)}$, *and* $P \subset \mathbb{R}^m$ *has* f *facets and inradius* $r(P)$ *such that*

- $P \subseteq (-q/2, q/2)^m$,
- $\dfrac{\mathrm{vol}(P)}{q^n} \in 2^{\omega(\log n)}$,
- $r(P) \in \omega(m^{3/2}/\sqrt{\log m}) \cap \Omega(m^{3/2}\sqrt{\log f}/\log m)$.

Then there exists a choice of δ *such that the generic construction yields a* CR-PSF, *under the assumption that finding collisions*

$$\mathbf{x}, \mathbf{x}' \in P \cap \mathbb{Z}^m, \quad \mathbf{x} \neq \mathbf{x}', \quad \mathbf{A}\mathbf{x} = \mathbf{A}\mathbf{x}'$$

for uniform $\mathbf{A} \leftarrow \mathsf{U}(\mathbb{Z}_q^{n \times m})$ *is hard.*

Given the stronger condition $r(P) \in \omega(m^{3/2}\sqrt{\log m}) \cap \Omega(m^{3/2}\sqrt{\log f})$, *the expected number of iterations required by* SampPoly *(in both* SampDom *and* SampPre*) is* $O(1)$.

Proof. Theorem 1 gives $\mathbf{A}$ close to uniform and basis $\mathbf{T}$ for $\Lambda = \Lambda_q^\perp(\mathbf{A})$ with $\|\widetilde{\mathbf{T}}\| \in O(m^{1/2})$. Thus $r(P)/\|\widetilde{\mathbf{T}}\| \in \omega(m/\sqrt{\log m}) \cap \Omega(m\sqrt{\log f}/\log m)$ and so Corollary 2 gives a choice of $\delta \in O(m/\log m)$ ensuring that

$$\varepsilon = f \cdot \exp\left(-\frac{1}{2}\frac{\delta^2 \cdot r(P)^2}{\|\widetilde{\mathbf{T}}\|^2}\right), \quad M = \frac{1}{(1-\varepsilon)^2}\left(\frac{1+\delta}{1-\delta}\right)^d$$

satisfy $\varepsilon \in \mathrm{negl}(n)$ and $M \in \mathrm{poly}(n)$, noting $m \in \mathrm{poly}(n)$.

Let $\varepsilon' = f \cdot \exp(-\delta^2 r(P)^2/2)$ and $M' = (1/(1-\varepsilon')^2) \cdot ((1+\delta)/(1-\delta))^d$ correspond to sampling from $P \cap \mathbb{Z}^m$. We have $\|\widetilde{\mathbf{T}}\| \geq 1$, as such $\varepsilon' \leq \varepsilon \in \mathrm{negl}(n)$ and $M' \leq M \in \mathrm{poly}(n)$. We check the properties of a CR-PSF per Definition 12.

1. *Trapdoor generation.* Satisfied via Theorem 1.
2. *Domain sampling with uniform output.* Since $\mathbf{I}_d$ is a unitriangular basis, we have by Theorem 3 that $\mathsf{SampPoly}(\mathbf{I}_d, P, \mathbf{c}, \delta)$ is efficient and samples uniformly on $P \cap \mathbb{Z}^m$ up to a statistical distance at most $\varepsilon'/(1-\varepsilon') \in \mathrm{negl}(n)$. Next, we show the output looks uniform under $f_{\mathbf{A}}$. Using Theorem 4,

$$|P \cap \mathbb{Z}^m| \geq \frac{\mathrm{vol}(P)}{M'(1-\varepsilon')} \in \frac{\mathrm{vol}(P)}{\mathrm{poly}(n)}$$

and thus

$$\frac{q^n}{|P \cap \mathbb{Z}^m|} \leq \mathrm{poly}(n) \cdot \frac{q^n}{\mathrm{vol}(P)} \in \mathrm{negl}(n)$$

since $\mathrm{vol}(P)/q^n \in 2^{\omega(\log n)}$. Furthemore, as $P \subseteq (-q/2, q/2)^m$ all points in $P \cap \mathbb{Z}^m$ are inequivalent modulo q. Using Lemma 15,

$$\mathop{\mathbb{E}}_{\mathbf{A} \leftarrow \mathsf{U}(\mathbb{Z}_q^{m})} \left[\Delta(\mathbf{A} \cdot \mathsf{U}(P \cap \mathbb{Z}^m), \mathsf{U}(\mathbb{Z}_q^n))\right] \leq \frac{1}{2}\sqrt{\frac{q^n}{|P \cap \mathbb{Z}^m|}} \in \mathrm{negl}(n).$$

Finally, by Remark 4 and $\mathsf{TrapGen}$ giving $\mathbf{A}$ statistically close to uniform, we have $\Delta(\mathbf{A} \cdot \mathsf{U}(P \cap \mathbb{Z}^m), \mathsf{U}(\mathbb{Z}_q^n))$ is negligible with overwhelming probability over $(\mathbf{A}, \cdot) \leftarrow \mathsf{TrapGen}(1^n)$.
3. *Preimage sampling with trapdoor.* Recall the definition of $X_n(a, y)$, written $X_n(\mathbf{A}, \mathbf{u})$ here, from §2.6. If no $\mathbf{x}$ exists such that $\mathbf{A}\mathbf{x} = \mathbf{u}$ then $\mathsf{SampPre}$ outputs $\perp$ and is distributed exactly as $X_n(\mathbf{A}, \mathbf{u})$. Else, by Theorem 5 and the choice of δ above, $\mathbf{z} \leftarrow \mathbf{x} + \mathsf{SampPoly}(\mathbf{T}, P - \mathbf{x}, \mathbf{c} - \mathbf{x}, \delta)$ is efficient and produces a sample from $\mathsf{U}(P \cap (\mathbf{x} + \Lambda))$ up to a negligible statistical distance $\varepsilon/(1-\varepsilon) \in \mathrm{negl}(n)$. Abusing notation in the second expression, we have

$$\Delta(\mathsf{U}(P \cap \mathbb{Z}^m)|_{\mathbf{x}+\Lambda}, \mathsf{U}(P \cap (\mathbf{x} + \Lambda))) = 0$$

and

$$\Delta(\mathsf{U}(P \cap (\mathbf{x} + \Lambda)), \mathbf{z}) \in \mathrm{negl}(n).$$

It remains to show that $\mathsf{SampDom}$ restricted to $\mathbf{x} + \Lambda$, written $X_n(\mathbf{A}, \mathbf{u})$, is close to the uniform distribution on $P \cap \mathbb{Z}^m$ restricted to the same coset. Let $N_{\mathbf{u}} = |P \cap (\mathbf{x} + \Lambda)|$. By the proof of Theorem 3 each $\mathbf{z} \in P \cap \mathbb{Z}^m$ has mass in $[1-\varepsilon', 1/(1-\varepsilon')]/|P \cap \mathbb{Z}^m|$. Therefore

$$\mathbb{P}\left[X_n(\mathbf{A}, \mathbf{u}) = \mathbf{z}\right] \in \frac{[(1-\varepsilon')^2, (1/(1-\varepsilon'))^2]}{N_{\mathbf{u}}}, \ \mathbf{z} \in P \cap (\Lambda + \mathbf{x})$$

since the normalisation factor is in $[1-\varepsilon', 1/(1-\varepsilon')] \cdot N_{\mathbf{u}}/|P \cap \mathbb{Z}^m|$. A simple calculation shows

$$\Delta(X_n(\mathbf{A}, \mathbf{u}), \mathsf{U}(P \cap \mathbb{Z}^m)|_{\mathbf{x}+\Lambda}) \leq \left((1/(1-\varepsilon'))^2 - 1\right)/2 \in \mathrm{negl}(n).$$

4. *Preimage min-entropy.* Note $|P \cap (\mathbf{x} + \Lambda)| \in 2^{\omega(\log n)}$ for all $\mathbf{x} \in \mathbb{Z}^m$ and $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$. Indeed, the third item of Theorem 5 gives

$$|P \cap (\mathbf{x} + \Lambda)| \geq \frac{1}{M(1-\varepsilon)} \frac{\mathrm{vol}(P)}{\mathrm{vol}(\Lambda)} \in 2^{\omega(\log n)}$$

as $\varepsilon \in \mathrm{negl}(n)$, $M \in \mathrm{poly}(m)$, $\mathrm{vol}(\Lambda) \leq q^n$ and $\mathrm{vol}(P)/q^n \in 2^{\omega(\log n)}$. This implies $H_\infty(\mathsf{U}(P \cap (\mathbf{x}+\Lambda))) \in \omega(\log n)$. By 3. above, $\mathsf{SampDom}(1^n)$ restricted to $\mathbf{x}+\Lambda$ is statistically close to uniform and thus also has min-entropy $\omega(\log n)$. Given $\mathbb{P}\left[\mathsf{U}(\mathbb{Z}_q^{n\times m}) \cdot \mathbb{Z}^m = \mathbb{Z}_q^n\right] \geq 1 - \mathrm{negl}(n)$, $\mathbf{A}(\mathbf{x}+\Lambda) = \{\mathbf{u}\}$ for $\mathbf{A}\mathbf{x} = \mathbf{u}$ and $(\mathbf{A}, \cdot) \leftarrow \mathsf{TrapGen}(1^n)$ returns $\mathbf{A}$ statistically close to uniform, we conclude.

5. *Collision resistance without trapdoor.* Satisfied by assumption and $(\mathbf{A}, \cdot) \leftarrow \mathsf{TrapGen}(1^n)$ being statistically close to uniform.

Repeating the proof with $r(P) \in \omega(m^{3/2}\sqrt{\log m}) \cap \Omega(m^{3/2}\sqrt{\log f})$ allows a choice $\delta \in O(1/m)$ such that the expected number of iterations in $\mathsf{SampPoly}$ within $\mathsf{SampPre}$ and $\mathsf{SampDom}$ are $M, M' \in O(1)$ respectively. □

6 Instantiations for various polytopes

In this section, we instantiate the generic construction of §5.2 for three different families of polytopes. In each case we choose parameters such that the conditions of Theorem 6 are satisfied, and hence attain a $\mathsf{CR\text{-}PSF}$ under the assumption that the corresponding collision resistance problem is hard.

The first two constructions are based on B_m^∞ and B_m^1, and thus collision resistance naturally reduces from SIS^∞ and SIS^1 respectively. The third is based on a polytope previously considered in [BBRS24], for which we give two reductions, from SIS^1 and from SIS^2. We note that in each case P has inradius $r(P) = r(P; \mathbf{0})$, so we set $\mathbf{c} = \mathbf{0}$ in all constructions.

6.1 CR-PSF based on SIS^∞

We let $P = B_m^\infty(w) \subset \mathbb{R}^m$ be the ℓ_∞ ball of radius $w = w(m) > 0$.

Corollary 3. *Instantiating Theorem 6 with*

- $P = B_m^\infty(w)$,
- $6n \log_2 q \leq m \in O(n \log q)$,
- $q/2 > w \geq \omega(m^{3/2}/\sqrt{\log m})$, *for prime* $q \leq 2^{\mathrm{poly}(n)}$,

gives a $\mathsf{CR\text{-}PSF}$ *assuming* $\mathrm{SIS}^\infty_{n,m,q,2w}$ *is hard.*

The stronger condition $q/2 > w \geq \omega(m^{3/2}\sqrt{\log m})$ *guarantees the expected number of iterations in* $\mathsf{SampPoly}$ *is constant.*

Remark 5. For $\mathsf{SampDom}$, rather than using $\mathsf{SampPoly}$, it is trivial to efficiently sample $\mathbf{x} \leftarrow \mathsf{U}(\{-\lfloor w \rfloor, \ldots, \lfloor w \rfloor\}^m)$.

Proof. P is effective since checking if $\mathbf{x} \in P$ requires only the m conditions $|x_i| \leq w$, and $\mathsf{U}(P) = 2w \cdot (U - 1/2)^m$ with $U \sim \mathsf{U}([0,1])$ is efficiently sampleable. Note the following properties of P,

$$\begin{aligned} \mathrm{vol}(P) &= (2w)^m, \\ r(P) &= w, \\ f &= 2m. \end{aligned}$$

The conditions of Theorem 6 are straightforward to verify using the properties above and the hypotheses.

Finally, by the triangle inequality for ℓ_∞, finding distinct $\mathbf{x}, \mathbf{x}' \in B_m^\infty(w)$ with $\mathbf{Ax} = \mathbf{Ax}'$ gives $\mathbf{x} - \mathbf{x}' \in B_m^\infty(2w) \cap (\Lambda_q^\perp(\mathbf{A}) \setminus \{\mathbf{0}\})$, i.e. solves $\mathrm{SIS}^\infty_{n,m,q,2w}$. □

6.2 CR-PSF based on SIS^1

Let $P = B_m^1(w) \subset \mathbb{R}^m$ be the ℓ_1 ball of radius $w = w(m)$.

Corollary 4. *Instantiating Theorem 6 with*

- $P = B_m^1(w)$,
- $6n \log_q q \leq m \in O(n \log q)$,
- $q/2 > w \geq \Omega(m^{5/2}/\log m)$, *for prime* $q \leq 2^{\mathrm{poly}(n)}$,

gives a CR-PSF *assuming* $\mathrm{SIS}^1_{n,m,q,2w}$ *is hard.*

The stronger condition $q/2 > w \geq \Omega(m^{5/2})$ *guarantees the expected number of iterations in* SampPoly *is constant.*

Proof. P is effective since one checks if $\mathbf{x} \in P$ by checking $\sum_i^m |x_i| \leq w$, and efficiently sampling uniformly from P is detailed in §2.1. We have, noting $w > m$,

$$\begin{aligned} \mathrm{vol}(P) &= \frac{(2w)^m}{m!} \geq 2^m, \\ r(P) &= \frac{w}{\sqrt{m}}, \\ f &= 2^m. \end{aligned}$$

and the conditions of Theorem 6 are easy to check. By the triangle inequality for ℓ_1, finding distinct $\mathbf{x}, \mathbf{x}' \in B_m^1(w)$ with $\mathbf{Ax} = \mathbf{Ax}'$ gives $\mathbf{x} - \mathbf{x}' \in B_m^1(2w) \cap (\Lambda_q^\perp(\mathbf{A}) \setminus \{\mathbf{0}\})$, i.e. solves $\mathrm{SIS}^1_{n,m,q,2w}$. □

6.3 CR-PSF based on SIS^2 or SIS^1

In [BBRS24] the following polytope, which is the intersection of appropriately scaled ℓ_1 and ℓ_∞ balls, is introduced. This polytope attempts to approximate a sphere well in the sense that it has inessentially more facets than B_m^1 and its inradius and outradius are fairly close, i.e. it is 'close' to a ball.

$$\mathcal{H}_m(w) = B_m^1(w\sqrt{m}) \cap B_m^\infty(w). \tag{6}$$

As the intersection of two convex polytopes, $\mathcal{H}_m$ is also convex. Indeed, it is described as the region bounded by $f = 2^m + 2m$ hyperplanes:

$$\mathcal{H}_m(w) = \left\{ \mathbf{x} \in \mathbb{R}^m : \bigwedge_{\mathbf{m} \in \{-1,1\}^m} (\mathbf{m}^t \mathbf{x} \leq w\sqrt{m}) \wedge \bigwedge_{\mathbf{m} = \pm \mathbf{e}_i} (\mathbf{m}^t \mathbf{x} \leq w) \right\}$$

from which we can compute the inradius $r(\mathcal{H}_m(w)) = w$. Furthermore, [BBRS24, Prop. 5] implies $\mathcal{H}_m(w)$ is contained in a ball of radius not much larger than w,

$$\mathcal{H}_m(w) \subseteq B_m^2(w\sqrt[4]{m}).$$

To instantiate the generic construction we must justify that it is an effective polytope. Determining whether $\mathbf{x} \in \mathcal{H}_m(w)$ is achieved by checking the m conditions $|x_i| \leq w$ and that $\sum_i^m |x_i| \leq w\sqrt{m}$. To achieve uniform sampling from $\mathcal{H}_m(w)$ we need the following two lemmas.

Lemma 16 ([BBRS24, Lem. 7]). *Define α_m as the probability that $\mathbf{x}$ sampled uniformly from $B_1(w\sqrt{m})$ belongs to $\mathcal{H}_m(w)$. Then*

$$\alpha_m = \sum_{i=0}^{\lfloor\sqrt{m}\rfloor} (-1)^i \binom{m}{i} \left(1 - \frac{i}{\sqrt{m}}\right)^{m+1}$$

Note that α_m is scale invariant, i.e. has no dependence on w.

Lemma 17. *For α_m as defined in Lemma 16,*

$$1 - \alpha_m \in \mathrm{negl}(m)$$

Proof. Recall $1 \leq \binom{m}{i} \leq m^i$ and for all $x \in \mathbb{R}$ we have $1 - x \leq e^{-x}$. As $\alpha_m \in [0, 1]$ we have $|\alpha_m - 1| = 1 - \alpha_m$. Let $s(m) = \lfloor\sqrt{m}\rfloor$ and note $1 - s(m)/\sqrt{m} \geq 0$. Then

$$1 - \alpha_m \leq \sum_{i=1}^{s(m)} m^i e^{-i(m+1)/\sqrt{m}} \leq \sum_{i=1}^{s(m)} m^i e^{-i\sqrt{m}}.$$

Writing the ith term as $\exp(-i\,(\sqrt{m} - \log m))$, we see the $i = 1$ term is largest. Since $s(m) \leq \sqrt{m}$ we conclude that

$$1 - \alpha_m \leq m^{3/2} e^{-\sqrt{m}} \in \mathrm{negl}(m). \qquad \square$$

Hence, combining Lemmas 16 and 17, sampling $\mathsf{U}(\mathcal{H}_m(w))$ is realised by repeatedly sampling $\mathbf{x} \leftarrow \mathsf{U}(B_m^1)$ until $\mathbf{x} \in \mathcal{H}_m(w)$. We have $\alpha_m \geq 1 - e^{-k\sqrt{m}}$ for some $k > 0$ and large enough m, therefore the expected number of iterations[9] is $1/\alpha_m \in 1 + O(e^{-k\sqrt{m}})$.

Corollary 5. *Instantiating Theorem 6 with*

- $P = \mathcal{H}_m(w)$,
- $6 \log_2 q \leq m \in O(n \log q)$,
- $q/2 > w \geq \Omega(m^2/\log m)$, *for prime* $q \leq 2^{\mathrm{poly}(n)}$,

gives a CR-PSF *assuming either*

[9] As for SampPoly, one may guarantee polynomial time by defining a maximum number of iterations $T_{\max} \in \omega(\log n)$ and return e.g. $\mathbf{0}$ if all fail.

- $\mathrm{SIS}^1_{n,m,q,\beta}$ *for* $\beta = 2w\sqrt{m}$ *is hard, or*
- $\mathrm{SIS}^2_{n,m,q,\beta}$ *for* $\beta = 2w\sqrt[4]{m}$ *is hard.*

The stronger condition $w \geq \Omega(m^2)$ *guarantees the expected number of iterations in* SampPoly *is constant.*

Proof. We proved P is effective above. We check the conditions of Theorem 6. The first is clear. For the second, $r(P) = w$ implies $\mathrm{vol}(P) \geq \mathrm{vol}(B^2_m(w))$. Asymptotically $w > m$ and $w^{1-1/m} > 2/\pi e$ and thus

$$\mathrm{vol}(B^2_m(w)) = \frac{\pi^{m/2}}{\Gamma(\frac{m}{2}+1)} w^m \sim \frac{1}{\sqrt{m\pi}} \left(\frac{2\pi e}{m}\right)^{m/2} w^m \geq \frac{2^m}{\sqrt{\pi}}$$
$$\implies \frac{\mathrm{vol}(P)}{q^n} \geq \frac{2^m}{q^n\sqrt{\pi}} \geq \frac{2^n}{\sqrt{\pi}} \in 2^{\omega(\log n)}.$$

For the third, $r(P) = w \geq \Omega(m^2/\log m) = \Omega(m^{3/2}\sqrt{\log f}/\log m)$, since we have $f \in (2^m, 2^{m+1})$.

Finally, since $P \subset B^2_m(w\sqrt[4]{m})$ and by the triangle inequality for ℓ_2, finding distinct $\mathbf{x}, \mathbf{x}' \in \mathcal{H}_m(w)$ with $\mathbf{Ax} = \mathbf{Ax}'$ gives $\mathbf{x} - \mathbf{x}' \in B^2_m(2w\sqrt[4]{m}) \cap (\Lambda^\perp_q(\mathbf{A}) \setminus \{\mathbf{0}\})$. The other hardness assumption follows as $\mathcal{H}_m(w) \subset B^1_m(w\sqrt{m})$. □

7 Comparisons

In this section we compare the instantiations of §6 to each other. Throughout we assume that m and q satisfy the requirements of constructions §§ 6.1 to 6.3. Furthermore, we assume that q is large enough to satisfy the worst-case to average-case reduction hypothesis of Theorem 2, such a $q \in \mathrm{poly}(n)$ always exists.

7.1 Comparison between polytopes

In Table 2 we list the norm bounds β_p for the SIS^p instances we reduce from in the CR-PSF constructions of §§ 6.1 to 6.3. We include both cases of the expected number of iterations M within SampPoly: constant and polynomial. In Figure 4 we write (P, SIS^p) when the construction with P admits a reduction from SIS^p. This corresponds to the row of Table 2 with polytope P and ℓ_p norm.

Assume for the sake of comparison that m and q are equal in each case, for example by taking the maximum such across cases that satisfy some criteria such as their constructions and Theorem 2. Then, in either regime of M, we have the following hardness relations.

- (B^1_m, SIS^1) *is not harder than* $(\mathcal{H}_m, \mathrm{SIS}^1)$. The norm bounds are identical and so the SIS problems are equal.
- $(\mathcal{H}_m, \mathrm{SIS}^2)$ *is not harder than* $(B^\infty_m, \mathrm{SIS}^\infty)$. Write β_2 and β_∞ for the respective norm bounds and note $\beta_2/\beta_\infty \in o(m^{3/4}/\sqrt{\log m})$. Therefore one may choose $\beta_2/\beta_\infty \geq \sqrt{m}$ and conclude by Lemma 10.

P	ℓ_p	SIS^p norm bound β_p for CR-PSF	
		$M \in \mathrm{poly}(m)$	$M \in O(1)$
B_m^1	$p = 1$	$\Omega(m^{5/2}/\log m)$	$\Omega(m^{5/2})$
$\mathcal{H}_m$	$p = 1$	$\Omega(m^{5/2}/\log m)$	$\Omega(m^{5/2})$
	$p = 2$	$\Omega(m^{9/4}/\log m)$	$\Omega(m^{9/4})$
B_m^∞	$p = \infty$	$\omega(m^{3/2}/\sqrt{\log m})$	$\omega(m^{3/2}\sqrt{\log m})$

Table 2: Polytopes P used in our CR-PSF constructions, and corresponding norm bounds β_p for the SIS^p problems used as hardness assumptions.

$$\begin{array}{ccc} (\mathcal{H}_m, \mathrm{SIS}^1) & \longleftarrow & (B_m^1, \mathrm{SIS}^1) \\ & & \overset{\text{heuristic}}{\uparrow} \\ & & (B_m^\infty, \mathrm{SIS}^\infty) \end{array} \longleftarrow (\mathcal{H}_m, \mathrm{SIS}^2)$$

Fig. 4: Hardness relations between the SIS^p problems in Table 2, assuming m and q are equal in all cases. An arrow $A \to B$ indicates A is not harder than B.

Some instantiations are not comparable via Lemma 10, so we perform a heuristic analysis. The following heuristic relies on the intuition that the hardness of SIS^p is determined by the number of points of $[0, q-1]^m$ that are solutions to a given instance. The heuristic states that the number of such points is roughly equal to the volume of $B_m^p(\beta_p)$ divided by the covolume of $\Lambda_q^\perp(\mathbf{A})$. For our parameters this covolume is q^n with overwhelming probability. If we therefore fix m and q we may compare the volume of the appropriate p-balls directly.

Heuristic 1 $\mathrm{SIS}^{p_1}_{n,m,q,\beta_{p_1}}$ *is not harder than* $\mathrm{SIS}^{p_2}_{n,m,q,\beta_{p_2}}$ *when*

$$\mathrm{vol}(B_m^{p_1}(\beta_{p_1})) \geq \mathrm{vol}(B_m^{p_2}(\beta_{p_2}))$$

i.e. the volume of the accepting region for p_1 *is larger.*

Lemma 18. *Let* r_1, r_∞ *be radii such that* $B_m^1(r_1) = B_m^\infty(r_\infty)$*. Then*

$$r_\infty = \frac{1}{(m!)^{1/m}} \sim_m \frac{e}{m}\, r_1.$$

We now compare $(B_m^\infty, \mathrm{SIS}^\infty)$ and (B_m^1, SIS^1). Writing β_∞ and β_1 for the corresponding norm bounds, we find that for m sufficiently large

$$\beta_\infty \geq \frac{\sqrt{\log m}}{m}\beta_1 > \frac{e}{m}\beta_1$$

and thus by Lemma 18, $\mathrm{vol}(B_m^\infty(\beta_\infty)) > \mathrm{vol}(B_m^1(\beta_1))$ (superexpontentially so), and we conclude that $(B_m^\infty, \mathrm{SIS}^\infty)$ *is heuristically not harder than* (B_m^1, SIS^1). We summarise the hardness relations in Figure 4.

7.2 Potential advantages

We conclude with a summary of potential advantages of our polytope CR-PSFs.

- *Flexibility of norm and maximum entropy.* Our constructions allow preimage sampling for ISIS^1 and ISIS^∞, and moreover sample uniformly from the solution space, maximising the min-entropy of preimages.
- *Simplicity (of implementation).* We avoid sampling from discrete Gaussian distributions. Sampling from the basic polytopes B_d^1, B_d^∞ is straightforward (see §2.1), and sampling from $\mathcal{H}_m$ requires an expected constant number of iterations that tends sharply to one. Checking membership is simple in each case. All polytopes appearing in our instantiations are affine linear transformations of these basic polytopes. The sampling of lattice points via SampPoly is efficient, requiring only the above operations and NearestPlane, and we can guarantee an expected constant number of iterations (at the expense of a growth in the norm bound by a factor of $\log m$).

Finally, we remark that if any of our open questions in §1.2 are answered positively, this immediately implies improved constructions.

References

Ajt96. Miklós Ajtai, *Generating hard instances of lattice problems (extended abstract)*, 28th ACM STOC, ACM Press, May 1996, pp. 99–108.

Ajt99. ———, *Generating hard instances of the short basis problem*, ICALP 99 (Jirí Wiedermann, Peter van Emde Boas, and Mogens Nielsen, eds.), LNCS, vol. 1644, Springer, Berlin, Heidelberg, July 1999, pp. 1–9.

AP11. Joël Alwen and Chris Peikert, *Generating shorter bases for hard random lattices*, Theory of Computing Systems **48** (2011), no. 3, 535–553.

Azu67. Kazuoki Azuma, *Weighted sums of certain dependent random variables*, Tohoku Mathematical Journal **19** (1967), no. 3, 357–367.

Bab86. L. Babai, *On lovász' lattice reduction and the nearest lattice point problem*, Combinatorica **6** (1986), no. 1, 1–13.

BBRS24. Henry Bambury, Hugo Beguinet, Thomas Ricosset, and Éric Sageloli, *Polytopes in the fiat-shamir with aborts paradigm*, CRYPTO 2024, Part I (Leonid Reyzin and Douglas Stebila, eds.), LNCS, vol. 14920, Springer, Cham, August 2024, pp. 339–372.

BGMN05. Franck Barthe, Olivier Guédon, Shahar Mendelson, and Assaf Naor, *A probabilistic approach to the geometry of the ℓ_p n-ball*, The Annals of Probability **33** (2005), no. 2, 480 – 513.

BL26. Katharina Boudgoust and Oleksandra Lapiha, *Leftover hash lemma(s) over cyclotomic rings*, Advances in Cryptology – EUROCRYPT 2026 (Cham) (Joan Daemen and Emmanuel Thomé, eds.), Springer Nature Switzerland, 2026, pp. 62–91.

CCD+24. Jung Hee Cheon, Hyeongmin Choe, Julien Devevey, Tim Güneysu, Dongyeon Hong, Markus Krausz, Georg Land, Marc Möller, Damien Stehlé, and MinJune Yi, *HAETAE: Shorter lattice-based fiat-shamir signatures*, IACR TCHES **2024** (2024), no. 3, 25–75.

Dev86. Luc Devroye, *Non-uniform random variate generation*, Springer-Verlag, 1986.

DFPS22. Julien Devevey, Omar Fawzi, Alain Passelègue, and Damien Stehlé, *On rejection sampling in Lyubashevsky's signature scheme*, ASIACRYPT 2022, Part IV (Shweta Agrawal and Dongdai Lin, eds.), LNCS, vol. 13794, Springer, Cham, December 2022, pp. 34–64.

FGdG+26. Pierre-Alain Fouque, Phillip Gajland, Hubert de Groote, Jonas Janneck, and Eike Kiltz, *A closer look at falcon*, Advances in Cryptology – EUROCRYPT 2026 (Cham) (Joan Daemen and Emmanuel Thomé, eds.), Springer Nature Switzerland, 2026, pp. 3–31.

GMRR22. Morgane Guerreau, Ange Martinelli, Thomas Ricosset, and Mélissa Rossi, *The hidden parallelepiped is back again: Power analysis attacks on falcon*, IACR TCHES **2022** (2022), no. 3, 141–164.

GPV07. Craig Gentry, Chris Peikert, and Vinod Vaikuntanathan, *Trapdoors for hard lattices and new cryptographic constructions*, Cryptology ePrint Archive, Report 2007/432, 2007.

GPV08. ———, *Trapdoors for hard lattices and new cryptographic constructions*, 40th ACM STOC (Richard E. Ladner and Cynthia Dwork, eds.), ACM Press, May 2008, pp. 197–206.

GR24. Morgane Guerreau and Mélissa Rossi, *A not so discrete sampler: Power analysis attacks on HAWK signature scheme*, IACR TCHES **2024** (2024), no. 4, 156–178.

Gru07. Peter M. Gruber, *Convex and discrete geometry*, no. 978-3-540-71132-2, Springer-Verlag Berlin Heidelberg, May 2007.

Haj42. Georg Hajós, *Über einfache und mehrfache bedeckung des n-dimensionalen raumes mit einem würfelgitter*, Mathematische Zeitschrift **47** (1942), no. 1, 427–467.

HvW24. Felicitas Hörmann and Wessel P. J. van Woerden, *FuLeakage: Breaking FuLeeca by learning attacks*, CRYPTO 2024, Part VI (Leonid Reyzin and Douglas Stebila, eds.), LNCS, vol. 14925, Springer, Cham, August 2024, pp. 253–286.

Kle00. Philip N. Klein, *Finding the closest lattice vector when it's unusually close*, 11th SODA (David B. Shmoys, ed.), ACM-SIAM, January 2000, pp. 937–941.

KV97. Ravi Kannan and Santosh Vempala, *Sampling lattice points*, Proceedings of the Twenty-Ninth Annual ACM Symposium on Theory of Computing (New York, NY, USA), STOC '97, Association for Computing Machinery, 1997, pp. 696–700.

Loe97. H. A. Loeliger, *Averaging bounds for lattices and linear codes*, IEEE Trans. Inf. Theor. **43** (1997), no. 6, 1767–1773.

LW15. Vadim Lyubashevsky and Daniel Wichs, *Simple lattice trapdoor sampling from a broad class of distributions*, PKC 2015 (Jonathan Katz, ed.), LNCS, vol. 9020, Springer, Berlin, Heidelberg, March / April 2015, pp. 716–730.

MP12. Daniele Micciancio and Chris Peikert, *Trapdoors for lattices: Simpler, tighter, faster, smaller*, EUROCRYPT 2012 (David Pointcheval and Thomas Johansson, eds.), LNCS, vol. 7237, Springer, Berlin, Heidelberg, April 2012, pp. 700–718.

MP13. ———, *Hardness of SIS and LWE with small parameters*, CRYPTO 2013, Part I (Ran Canetti and Juan A. Garay, eds.), LNCS, vol. 8042, Springer, Berlin, Heidelberg, August 2013, pp. 21–39.

MR07. Daniele Micciancio and Oded Regev, *Worst-case to average-case reductions based on gaussian measures*, SIAM Journal on Computing **37** (2007), no. 1, 267–302.

Pei08. Chris Peikert, *Limits on the hardness of lattice problems in ℓ_p norms*, computational complexity **17** (2008), no. 2, 300–351.

PFH+22. Thomas Prest, Pierre-Alain Fouque, Jeffrey Hoffstein, Paul Kirchner, Vadim Lyubashevsky, Thomas Pornin, Thomas Ricosset, Gregor Seiler, William Whyte, and Zhenfei Zhang, *FALCON*, Tech. report, National Institute of Standards and Technology, 2022, available at `https://csrc.nist.gov/Projects/post-quantum-cryptography/selected-algorithms-2022`.

PP21. Maxime Plançon and Thomas Prest, *Exact lattice sampling from non-gaussian distributions*, PKC 2021, Part I (Juan Garay, ed.), LNCS, vol. 12710, Springer, Cham, May 2021, pp. 573–595.

PS24. Amaury Pouly and Yixin Shen, *Provable dual attacks on learning with errors*, EUROCRYPT 2024, Part VII (Marc Joye and Gregor Leander, eds.), LNCS, vol. 14657, Springer, Cham, May 2024, pp. 256–285.

RBK+23. Stefan Ritterhoff, Sebastian Bitzer, Patrick Karl, Georg Maringer, Thomas Schamberger, Jonas Schupp, Georg Sigl, Antonia Wachter-Zeh, and Violetta Weger, *FuLeeca*, Tech. report, National Institute of Standards and Technology, 2023, available at `https://csrc.nist.gov/Projects/pqc-dig-sig/round-1-additional-signatures`.

RT87. Prabhakar Raghavan and Clark D. Thompson, *Randomized rounding: A technique for provably good algorithms and algorithmic proofs*, Combinatorica **7** (1987), no. 4, 365–374.

SC89. Carl Ludwig Siegel and Komaravolu Chandrasekharan, *Lectures on the geometry of numbers*, no. 978-3-642-08076-0, Springer Berlin, Heidelberg, December 1989.