On Rejection Sampling in Lyubashevsky's Signature Scheme

Julien Devevey[1], Omar Fawzi[1,2], Alain Passelègue[1,2], and Damien Stehlé[1,3]

[1] ENS de Lyon, Lyon, France
[2] Inria Lyon, Lyon, France
[3] Institut Universitaire de France, Paris, France

Abstract. Lyubashevsky's signatures are based on the Fiat-Shamir with aborts paradigm, whose central ingredient is the use of rejection sampling to transform secret-dependent signature samples into samples from (or close to) a secret-independent target distribution. Several choices for the underlying distributions and for the rejection sampling strategy can be considered. In this work, we study Lyubashevsky's signatures through the lens of rejection sampling, and aim to minimize signature size given signing runtime requirements. Several of our results concern rejection sampling itself and could have other applications.

We prove lower bounds for compactness of signatures given signing runtime requirements, and for expected runtime of perfect rejection sampling strategies. We also propose a Rényi-divergence-based analysis of Lyubashevsky's signatures which allows for larger deviations from the target distribution, and show hyperball uniforms to be a good choice of distributions: they asymptotically reach our compactness lower bounds and offer interesting features for practical deployment. Finally, we propose a different rejection sampling strategy which circumvents the expected runtime lower bound and provides a worst-case runtime guarantee.

1 Introduction

Lyubashevsky's signature scheme [Lyu09,Lyu12] may be viewed as a lattice variant of Schnorr's group-based signature scheme [Sch91], with a core conceptual difference being the use of rejection sampling and the associated introduction of aborts and repeats in the Fiat-Shamir heuristic [FS86]. The use of rejection sampling in Lyubashesvky's scheme is the focus of the present work. It is hard to overstate the importance of Lyubashevsky's signature scheme in lattice-based cryptography. Thanks to its elementary and flexible design, numerous variants and optimizations have been proposed (see [AFLT16,GLP15,DDLL13,BG14], or [Lyu16], for instance). Notably, it led to the TESLA [ABB$^+$17,AAB$^+$19] and Dilithium [DKL$^+$18,BDK$^+$20] candidates to the NIST standardization project on post-quantum cryptography. It also led to lattice-based zero-knowledge proofs (see [LNP22] and the references therein).

Lyubashevsky's scheme involves a publicly shared matrix $\mathbf{A} \in \mathbb{Z}_q^{n\times m}$ (note that other algebraic setups are possible, but this is not relevant to the present discussion). The signing key is a matrix $\mathbf{S} \in \mathbb{Z}^{m \times k}$. It is small in the sense that all its entries have absolute values significantly smaller than q. The verification key associated to $\mathbf{S}$ is $\mathbf{T} = \mathbf{AS}$. Given a message $\mu \in \{0,1\}^*$, the signer samples a small masking vector $\mathbf{y} \in \mathbb{Z}^m$ and computes a random-looking commitment $\mathrm{com} = \mathbf{Ay}$. By using a hash function H taking small values in $\mathbb{Z}^k$, it computes a challenge $\mathbf{c} = H(\mathrm{com}, \mu)$. Finally, if some (possibly probabilistic) test passes, it outputs a signature $\sigma = (\mathbf{z}, \mathbf{c})$ with $\mathbf{z} = \mathbf{y} + \mathbf{Sc}$, and else it restarts from scratch. Given a signature $\sigma = (\mathbf{z}, \mathbf{c})$ for a message μ, the verifier accepts if and only if $\mathbf{z}$ is small and $H(\mathbf{Az} - \mathbf{Tc}, \mu) = \mathbf{c}$. We refer the reader to Figure 2 for a formal description. As suggested by the choice of terminology, Lyubashevsky's signature can be viewed as an identification protocol made non-interactive by relying on the Fiat-Shamir heuristic, i.e., by replacing a truly random $\mathbf{c}$ by the output of a hash function. The security proof relies on the Random Oracle Model (ROM) as it models H as a function such that each image is distributed as $\mathbf{c}$ is supposed to be.

Compared to Schnorr's signature scheme, the signing key and mask do not belong to a finite set, preventing the use of a uniform mask $\mathbf{y}$ to hide the sensitive term $\mathbf{Sc}$.[4] One possibility (see, e.g., [DPSZ12]) is to sample $\mathbf{y}$ exponentially larger than $\mathbf{Sc}$ as a function of the security parameter, so that the distributions of $\mathbf{y}$ and $\mathbf{y} + \mathbf{Sc}$ have exponentially small statistical distance. As q must be larger than $\mathbf{y}$ and the smallness of $\mathbf{S}$ relative to q impacts security, this flooding approach leads to large parameters. Instead, Lyubashevsky [Lyu09,Lyu12] put forward the notion of Fiat-Shamir with aborts. This is the reason for the test concerning $\mathbf{z}$ in the signing algorithm: it is so that the output signature $(\mathbf{z}, \mathbf{c})$ follows a distribution that is independent of the sensitive term $\mathbf{Sc}$.

A classic application of rejection sampling (see, e.g., [Dev86, Chapter 2]) is to use a source distribution Q that is convenient to sample from, to create samples from a target distribution P. In Lyubashevsky's scheme, the purpose differs: we start from a pre-source distribution Q for $\mathbf{y}$; it is shifted by $\mathbf{Sc}$, leading to a distribution $Q_{+\mathbf{Sc}}$ for $\mathbf{y} + \mathbf{Sc}$; the latter is the source distribution; it is rejected to a target distribution P for $\mathbf{z}$ that does not depend on the signing key $\mathbf{S}$. The purpose of rejection sampling here is to hide the sensitive data $\mathbf{Sc}$. Diverse choices of pairs of distributions have been put forward in the literature: uniform in hypercubes [Lyu09], Gaussian with the same standard deviation while allowing for some small statistical inaccuracy in the target distribution [Lyu12], a bimodal Gaussian source distribution with a Gaussian target distribution in association with an accomodating arithmetic modification of the scheme [DDLL13] (the modification consists in replacing q with $2q$ and changing key generation to ensure that $\mathbf{T} = -\mathbf{T} = q\mathbf{I} \bmod 2q$). The pre-source distribution Q is shifted by $(-1)^b\mathbf{Sc}$ for a uniform bit b, leading to a source distribution $Q_{\pm\mathbf{Sc}}$. We refer to this as the bimodal setting. By opposition, we now refer to the former two cases where the source distribution is $Q_{+\mathbf{Sc}}$ as the unimodal setting. The first

[4] If we view $\mathbf{y}$ and $\mathbf{S}$ over $\mathbb{Z}_q$ rather than $\mathbb{Z}$, then they do belong to a finite set; but for security, the masking should preserve smallness relative to q, which the uniform distribution modulo q does not achieve.

choice (uniform distributions in hypercubes) leads to a simple design, whereas the latter two allow for more compact signatures. One may also want to add constraints on the number of loop iterations, notably to guarantee a signing runtime upper bound. In the extreme case of removing rejection altogether, it was recently shown in [ASY22] that a limited flooding suffices, compared to the exponential flooding discussed earlier. This leads us to the question we address in this work:

Given signing runtime requirements, which rejection sampling strategy leads to the most compact signatures?

In a signature $\sigma = (\mathbf{z}, \mathbf{c})$, the second component contributes to a small fraction of the bitsize: the main requirement on $\mathbf{c}$ is that it has sufficiently high min-entropy to make it hard to guess. On the other hand, the contribution of $\mathbf{z}$ towards signature length is mostly driven by $\|\mathbf{z}\|$, as this directly impacts security: for a given security level, the smaller $\|\mathbf{z}\|$, the more compact the signatures. For this reason, we simplify the overall objective to minimizing $\mathbb{E}_{\mathbf{x} \hookleftarrow P}(\|\mathbf{x}\|)$ under signing runtime requirements.

Contributions. Our main contributions concern the optimality of rejection sampling design choices towards optimizing signature sizes and signing runtime. We provide lower bounds, and study ways to reach and circumvent them.

Before describing the main results, we need to quantify the runtime of rejection sampling strategies. We note that for classic rejection sampling with target P and source Q, the expected number of samples needed is $R_\infty(P\|Q)$ where $R_\infty(D_1\|D_2) = \sup_x D_1(x)/D_2(x)$ refers to the Rényi divergence of infinite order. Indeed, for classic rejection sampling, one samples x from Q and accepts with probability $P(x)/(M \cdot Q(x))$, for $M = R_\infty(P\|Q)$. This justifies using $R_\infty(P\|Q)$ to quantify the runtime for rejecting Q to P.

We start with our lower bounds.

- Considering Lyubashevsky's scheme with perfect rejection sampling to the target distribution P (as in [Lyu09]), the relevant quantity measuring the signing runtime is then given by $M = \max_{\mathbf{S},\mathbf{c}} R_\infty(P\|Q_{+\mathbf{Sc}})$. We show (under a mild assumption discussed below) that for all P and Q such that M is finite, the expected norm $\mathbb{E}_{\mathbf{x} \hookleftarrow P}(\|\mathbf{x}\|)$ is $\Omega((m/\log M) \cdot \max_{\mathbf{S},\mathbf{c}} \|\mathbf{Sc}\|)$. Interestingly, this bound is a factor $\sqrt{m}$ lower than what is obtained for the typical choice of P and Q set as uniform distributions in hypercubes.
- In the case of perfect rejection with the accommodating arithmetic modification from [DDLL13], then the relevant quantity for measuring the signing runtime is $M = \max_{\mathbf{S},\mathbf{c}} R_\infty(P\|Q_{\pm\mathbf{Sc}})$, where $Q_{\pm\mathbf{Sc}}$ denotes the balanced mixture of $Q_{+\mathbf{Sc}}$ and $Q_{-\mathbf{Sc}}$. In this case, we show (under the same mild assumption) that for all P and Q such that M is finite, the expected norm $\mathbb{E}_{\mathbf{x} \hookleftarrow P}(\|\mathbf{x}\|)$ is $\Omega(\sqrt{m/\log M} \cdot \max_{\mathbf{S},\mathbf{c}} \|\mathbf{Sc}\|)$. This lower bound is actually reached (up to a constant factor) for P and Q Gaussian as in [DDLL13].
- We show that for any algorithm (terminating with probability 1) that selects one out of many samples from Q to get a sample from P, the expected number of required samples from Q is $\geq R_\infty(P\|Q)$. This lower bound is

reached by classic rejection sampling. In the case of Lyubashesvky's signatures with exact rejection sampling, this general result implies that classic rejection sampling is the appropriate strategy when it comes to minimize the expected runtime.

The lower bounds above seem to give little margin of improvement in the design choices of Lyubashevsky's signatures, except for the unimodal case, for which uniform distributions in hypercubes do not reach the lower bound. Our second set of results considers ways to reach or circumvent these lower bounds.

- Concerning the unimodal case, one way to circumvent the results above is to consider imperfect rejection sampling, by allowing for an approximation to P whose accuracy is parameterized by some $\varepsilon > 0$ (as introduced in [Lyu12]). Then the relevant quantity to bound the runtime becomes $\max_{\mathbf{S},\mathbf{c}} R_{\infty}^{\varepsilon}(P \| Q_{+\mathbf{Sc}})$, where R_{∞}^{ε} is a smoothed variant of R_{∞} that we define. In this case, we improve the signature security analysis from [Lyu12] by using the Rényi divergence instead of the statistical distance to quantify the closeness to P of the output distribution. This allows choosing ε larger than previously, leading to a (limited) signature compactness improvement.
- Gaussian distributions provide better signature compactness in the bimodal and imperfect unimodal regimes, than uniforms in hypercubes in the perfect unimodal regime. However, uniforms in hypercubes are sometimes preferred (see, e.g., Dilithium), because they lead to a simpler implementation, which in turn makes protection against timing attacks easier. We consider uniforms in hyperballs as a new alternative for the choice of source and target distributions. We show that this choice reaches the two lower bounds for $\mathbb{E}_{\mathbf{x} \hookleftarrow P}(\|\mathbf{x}\|)$ for perfect rejection sampling and is as good as Gaussians for imperfect rejection sampling (up to a constant factor). Interestingly, the rejection test for uniforms in hyperballs is very simple, similarly to uniforms in hypercubes. We not only study the choice of uniforms in hyperballs in the asymptotic regime, but also compare it to Dilithium.
- Finally, imperfect rejection from Q to P allows us to describe and analyze variants of rejection sampling where the maximum number of loop iterations is bounded. This provides trade-offs between maximum signing runtime and signature sizes. When instantiated to rejection-free sampling, we recover the scheme and analysis from [ASY22], whereas it quickly converges to Lyubashevsky's signature scheme when the signing runtime bound grows.

The results concerning signature compactness for unbounded (perfect and imperfect) rejection sampling are summarized in Table 1.

Technical overview. In Section 2, we provide the background necessary to this work, including rejection sampling and Lyubashevsky's signature scheme.

After identifying the notion of expected number of iterations during rejection sampling with the notion of smooth-Rényi divergence that we define, we start addressing our main question of understanding to which extent the expected norm of a signature can be small for target expected signing runtime constraints.

Lower Bounds. In Section 3, we prove lower bounds in the case of exact rejection sampling in both unimodal and bimodal settings. These lower bounds

	Unimodal ($\varepsilon = 0$)	Unimodal ($\varepsilon \geq 2^{-o(m)}$ and $\varepsilon = o(1/m)$)	Bimodal ($\varepsilon = 0$)
Hypercube	$\frac{tm^{3/2}}{\log M}$	$\frac{tm^{3/2}}{\log M}$	$\frac{tm^{3/2}}{\log M}$
Gaussian	∞	$\frac{t\sqrt{m}\sqrt{\log\frac{1}{\varepsilon}+\log M}}{\log M}$	$\frac{t\sqrt{m}}{\sqrt{\log M}}$
Hyperball	$\frac{tm}{\log M}$ (Lemma 6)	$\frac{t\sqrt{m}\sqrt{\log\frac{1}{\varepsilon}+\log M}}{\log M}$ (Lemma 6)	$\frac{t\sqrt{m}}{\sqrt{\log M}}$ (Lemma 7)
Lower bound	$\frac{tm}{\log M}$ (Corollary 1)	?	$\frac{t\sqrt{m}}{\sqrt{\log M}}$ (Corollary 2)

Table 1. This table expresses the compactness of the signature modeled as $\mathbb{E}_{\mathbf{x}\hookleftarrow P}(\|\mathbf{x}\|)$ given the signing runtime constraint for various choices of distributions P and Q. The column indicates the signing runtime constraint which is modeled in the unimodal case by $\max_{\mathbf{v}\in\mathcal{B}_m(t)} R_\infty^\varepsilon(P\|Q_{+\mathbf{v}}) \leq M$ where ε quantifies the accuracy of rejection sampling and in the bimodal case by $\max_{\mathbf{v}\in\mathcal{B}_m(t)} R_\infty(P\|Q_{\pm\mathbf{v}}) \leq M$. In the first row, P and Q are chosen to be uniform in m-dimensional hypercubes of appropriate side-lengths, in the second row, they are chosen to be m-dimensional Gaussians of appropriate variance. In the third row, they are chosen to be uniform in the m-dimensional hyperballs of appropriate radii. The last row gives a lower bound on the compactness for any choice of P and Q. Multiplicative constants are omitted in this table, and we make the assumption that $\log M \leq m$.

are obtained following a similar path. In what follows, we focus on the unimodal setting. To ease the analysis, we place ourselves in a slightly simplified setup where shifts belong to a hyperball $\mathcal{B}_m(t)$ of radius t instead of being defined as $\mathbf{Sc}$. Given that $\mathbf{S}$ is unknown, this simplification seems reasonable and allows avoiding significant complications in the proof. In this setting, we prove that for a given constraint $\max_{\mathbf{v}\in\mathcal{B}_\mathbf{m}(\mathbf{t})} R_\infty(P\|Q_{+\mathbf{v}}) \leq M$, we have $\mathbb{E}_{\mathbf{x}\hookleftarrow P}(\|\mathbf{x}\|) \geq (t/M^{1/(m-1)} - 1) - \sqrt{m}/2$.

Our lower bounds are obtained in three steps: (1) considering the same setting with continuous distributions, we first prove that we can restrict ourselves to the case of isotropic distributions over $\mathbb{R}^m$, where isotropic means that their densities only depend on the norm. Specifically, we prove that for any two densities f, g, there exist isotropic distributions f^*, g^* satisfying $\max_{\mathbf{v}\in\mathcal{B}_m(t)} R_\infty(f^*\|g^*_{+\mathbf{v}}) \leq M$ as well as $\mathbb{E}_{\mathbf{x}\hookleftarrow f^*}(\|\mathbf{x}\|) = \mathbb{E}_{\mathbf{x}\hookleftarrow f}(\|\mathbf{x}\|)$. The latter distributions are essentially obtained from f, g by averaging their respective densities on hyperspheres. (2) Starting with f and g isotropic, we show that $\mathbb{E}_{\mathbf{x}\hookleftarrow f}(\|\mathbf{x}\|) = \mu_m/\mu_{m-1}$ where $\mu_k = \int_0^\infty r^k f(r)\,\mathrm{d}r$. The main technicality consists in proving an intermediate lower bound $\mu_{m-1}/\mu_{m-2} \geq (t/M^{1/(m-1)} - 1)$ which results from the constraint $\max_{\mathbf{v}\in\mathcal{B}_m(t)} R_\infty(f\|g_{+\mathbf{v}}) \leq M$. Our lower bound is then obtained by applying the Cauchy-Schwarz inequality $|\mathbb{E}(XY)|^2 \leq \mathbb{E}(X^2)\mathbb{E}(Y^2)$ to random variables $X = \|\mathbf{x}\|^{m/2}$ and $Y = \|\mathbf{x}\|^{(m-2)/2}$, where $\mathbf{x} \hookleftarrow f$. Indeed, it immediately leads to inequality $\mu_m \cdot \mu_{m-2} \geq (\mu_{m-1})^2$, which results in $\mu_m/\mu_{m-1} \geq \mu_{m-1}/\mu_{m-2} \geq (t/M^{1/(m-1)} - 1)$. (3) A similar lower bound in the discrete setting is obtained by considering the continuous density $p(\mathbf{x}) = P(\lceil \mathbf{x} \rfloor)$ with P being a discrete probability. These lower bounds provide us with a target to reach, and we can compare them with the signature size obtained when instantiating the above scheme with various distributions.

On Alternative Rejection Sampling Strategies. In Section 3.3, we investigate the question of the relevance of rejection sampling strategies differing from the classic one. We consider the following setting. As above, the goal is to sample from a distribution P given access to a sampler from a distribution Q, and we consider a sequence of samples $(X_i)_{i\geq 1}$ from distribution Q. Any strategy is allowed as long as we output one of the X_i's. A strategy is given by a sequence of algorithms $(A_i)_{i\geq 1}$ that take samples $(X_j)_{j\leq i}$ as input and return either an index $j \in [i]$, which corresponds to halting with output X_j, or a special symbol r which corresponds to rejecting and moving to A_{i+1}. We restrict ourselves to the case of procedures that terminate with probability 1. Considering i^* the random variable denoting the number of samples observed in a strategy, our objective is then to measure how small $\mathbb{E}(i^*)$ can be. We prove that for any P, Q, we have $\mathbb{E}(i^*) \geq R_\infty(P \| Q)$. This result is obtained by proving that for any x, we have $P(x) \leq \mathbb{E}(i^*) \cdot Q(x)$, leading to the former inequality by definition of R_∞.

Rényi-Based Analysis for Imperfect Rejection Sampling. All lower bounds are for perfect rejection sampling, in the sense that one obtains a sample from (exactly) P. In [Lyu12], Lyubashevsky showed that one can consider imperfect rejection sampling, and shows that it is particularly beneficial in the case of Gaussians. We propose an analysis that replaces the use of the statistical distance as done in [Lyu12] by that of the smooth Rényi divergence, and allows loosening the constraints on imperfectness. We first recall that in [Lyu12], the statistical distance is used to bound the statistical distance between a (single) execution of the imperfect rejection sampling algorithm and the target distribution. Using imperfect rejection sampling in a signature scheme and given bound ε for the above statistical distance, one can then bound the distinguishing advantage of an adversary between the real security game and the ideal game (where signatures are simulated by sampling them from the target distribution) by $q_{\mathsf{sig}} \cdot \varepsilon$. Here q_{sig} is a bound on the number of signature queries an adversary can make. In Section 4, we prove that for P, Q such that $R_\infty^\varepsilon(P \| Q)$ is finite, the Rényi divergence of infinite order between a (single) execution of the imperfect rejection sampling algorithm and the target distribution is bounded by $1/(1-\varepsilon)$. Combining this result with the multiplicativity of the Rényi divergence, we can then bound the Rényi divergence of infinite order between the adversary's view in the real game and its view in the ideal game by $1/(1-\varepsilon)^{q_{\mathsf{sig}}}$ for the resulting signature. The probability preservation property of the Rényi divergence then allows completing the analysis. Our analysis leads to potential improvements as the former statistical bound $q_{\mathsf{sig}} \cdot \varepsilon$ imposes that $\varepsilon = 2^{-\Omega(\lambda)}$, while our bound can be used setting $\varepsilon = 1/q_{\mathsf{sig}}$. Since q_{sig} is a (possibly large) polynomial of the security parameter λ, this puts less constraint on the condition P and Q must satisfy, which results in compactness improvement.

Hyperball Uniforms. We show that (continuous) uniform distributions over hyperballs reach the signature compactness lower bound (up to constant factors) in both unimodal and bimodal settings, as shown in Section 5.1. We also show that they are as good as Gaussians for imperfect rejection sampling (up to a constant factor). These results reduce to Rényi divergence computations, which involve geometric properties of hyperballs. We emphasize that while Gaussian distributions also achieve similar signature size (up to a constant factor) in both unimodal and bimodal settings (but only in the case of imperfect rejection sampling with polynomial loss for the unimodal case), using uniform distributions over hyperballs makes the rejection test as simple as computing $\|\mathbf{z}\|$ since it consists only in checking that $\mathbf{z}$ is in the hyperball of the target distribution P. In order to use this distribution in a signature, we propose a generalization of Lyubashevsky's signature that allows for continuous source and target distributions, by adding a rounding step after accepting a sample. Its security relies on the same mechanisms as the discrete case. This strategy could also benefit to Gaussian distributions, by allowing to replace discrete Gaussian sampling with possibly simpler continuous Gaussian sampling. To assess the practicality of this new choice of distributions, we propose parameters for a variant of Dilithium with uniform distributions in hyperballs. If considering the sum of bitsizes of a verification key and a signature, the gains range from $\sim 15\%$ to $\sim 25\%$, depending on the security level.

Bounded Rejection Sampling. We conclude this work by proposing an original strategy to use rejection sampling while guaranteeing a (moderate) worst-case runtime. This could be beneficial in the context of real-time systems. A simple strategy could consist in fixing a (very large) bound i on the number of iterations such that it fails to produce a sample with negligible probability. While this guarantees a worst-case runtime, the change is mainly cosmetic since it has to be large enough for the sampling to succeed. In Section 6, we propose an alternative solution that leaves the choice of i open without ever failing: for a fixed bound i, it performs (up to) $i-1$ iterations of the classic rejection sampling and outputs a sample if it ever succeeds, otherwise, the last (i-th) iteration uses one-shot flooding techniques (as done in [ASY22]) to guarantee an output. The analysis makes heavy use of the smooth Rényi divergence and its properties. Different choices for the bound i offer various trade-offs, ranging from one-shot signatures ($i = 1$) as in [ASY22] to Lyubashevsky's expected polynomial-time signatures (i going to ∞).

Open problems. Our results suggest that instantiating the Fiat-Shamir with aborts using uniform distributions in hyperballs is a relevant choice, both in the unimodal and bimodal settings, as it provides more compact signatures than uniform distributions in hypercubes but also much simpler rejection test than Gaussians. We believe it is an interesting open question to investigate a constant-time implementation with this choice. Regarding further improvements of signatures, our results show that there is not much room for improvement if the goal is to minimize signature size or $\mathbb{E}(i^*)$. However, other quantities could be considered, such as the shape of the tail of the distribution of i^*.

2 Preliminaries

See full version for notations and some standard background.

We introduce a relaxed version of the Rényi divergence, termed the smooth Rényi divergence, where one is able to remove a few problematic points from the support, including those that may lie in $\mathrm{Supp}(p) \setminus \mathrm{Supp}(q)$. Doing so, we can compare a wider set of probability distributions. For instance, while the Rényi divergence of infinite order between $D_{\mathbb{Z}^m,\sigma}$ and $D_{\mathbb{Z}^m,\sigma,\mathbf{v}}$ is infinite when $\mathbf{v} \neq \mathbf{0}$, their smooth divergence is finite, as we show in the full version and is implicit in [Lyu12]. We could give this definition for any order $a \in [1, +\infty]$. However, only the case $a = +\infty$ is relevant for this work.

This definition is useful to link previous works on rejection sampling and the Rényi divergence. A similar quantity has been previously defined in the quantum information literature [Ren05,Dat09], though the specific notion of smoothing we consider here is slightly different.

Definition 1 (Smooth Rényi Divergence). *Let $\varepsilon \geq 0$. Let p, q be two probability densities such that $\int_{\mathrm{Supp}(q)} p(x)\,\mathrm{d}\mu(x) \geq 1 - \varepsilon$. Their ε-smooth Rényi divergence of infinite order is*

$$R_\infty^\varepsilon(p\|q) := \inf_{\substack{S \subseteq \mathrm{Supp}(q) \\ \int_S p(x)\,\mathrm{d}\mu(x) \geq 1-\varepsilon}} \operatorname*{ess\,sup}_{x \in S} \frac{p(x)}{q(x)}.$$

This definition is equivalent to

$$R_\infty^\varepsilon(p\|q) := \inf\{M > 0 \mid \Pr_{x \hookleftarrow p}(p(x) \leq Mq(x)) \geq 1 - \varepsilon\}.$$

By convention, if $\int_{\mathrm{Supp}(q)} p(x)\,\mathrm{d}\mu(x) < 1 - \varepsilon$, we define $R_\infty^\varepsilon(p\|q) = +\infty$.

In the full version, we prove that the two definitions are indeed equivalent and give useful properties of the smooth Rényi divergence.

2.1 Rejection Sampling

Given two close enough densities p_t and p_s, either both continuous or both discrete, rejection sampling is a way to generate samples from p_t given access to samples from p_s, as explained for instance in [Dev86]. It was used mainly to generate samples from complex distributions that were "close" to easier-to-sample distributions. However, in cryptography and particularly in the line of works started with [Lyu09], it found a peculiar use that diverged from its primary use. Given a family of densities $(p_s^{(v)})$, rejection sampling can be used to hide the parameter v given a density p_t that is close to every density in this family. It was later observed in [Lyu12] that an "imperfect" rejection procedure is sufficient for this use and leads to smaller parameters, notably standard deviation of p_s.

In the case of Lyubashevsky's signature scheme [Lyu09,Lyu12], a signature is a pair of vectors $(\mathbf{y}+\mathbf{Sc},\mathbf{c})$ where $\mathbf{y} \hookleftarrow P_{\mathbf{y}}$ and $\mathbf{c}$ would ideally be sampled from $P_{\mathbf{c}}$. Here $P_{\mathbf{c}} : \mathcal{C} \to \mathbb{R}$ and $P_{\mathbf{y}} : \mathbb{Z}^m \to \mathbb{R}$ are two discrete probability distributions, where $\mathcal{C} \subset \mathbb{Z}^k, m, k \geq 1$, and $\mathbf{S} \in \mathbb{Z}^{m \times k}$ is fixed (it is the signing key). The joint distribution of this pair corresponds to the source distribution $P_s^{(\mathbf{Sc})}$ above, which depends on the sensitive data $\mathbf{Sc}$. Rejection sampling is used to ensure that the output of the signing algorithm is of the form $(\mathbf{z},\mathbf{c})$ where $\mathbf{z} \hookleftarrow P_{\mathbf{z}}$ and $\mathbf{c} \hookleftarrow P_{\mathbf{c}}$ are statistically independent and $P_{\mathbf{z}}$ is well-chosen. Their joint distribution corresponds to the target distribution P_t above. The case of BLISS [DDLL13] is identical, except that signatures are of the form $(\mathbf{y}+(-1)^b\mathbf{Sc},\mathbf{c})$, where $b \hookleftarrow U(\{0,1\})$.

We consider the algorithms from Figure 1, which take some $M \geq 1$ as a parameter. Algorithm $\mathcal{A}^{\mathsf{ideal}}$ corresponds to what we would like to have, whereas $\mathcal{A}^{\mathsf{real}}$ is the algorithm corresponding to the real distribution. We are typically interested in calling these algorithms until they output something, which is what $\mathcal{B}_\infty^{\mathsf{real}}$ and $\mathcal{B}_\infty^{\mathsf{ideal}}$ do. It remains to understand when the outputs of these algorithms are statistically close. The lemma below is proved in the full version.

```
Algorithm $\mathcal{A}^{\mathsf{real}}$:
  $x \hookleftarrow p_s$
  with probability $\min\left(\frac{p_t(x)}{M \cdot p_s(x)}, 1\right)$, return $x$
  return $\perp$
```

```
Algorithm $\mathcal{A}^{\mathsf{ideal}}$:
  $x \hookleftarrow p_t$
  with probability $\frac{1}{M}$, return $x$
  return $\perp$
```

```
Algorithm $\mathcal{B}_\infty^{\mathsf{real}}$:
  $z \leftarrow \perp$
  while $z = \perp$ do
      $z \leftarrow \mathcal{A}^{\mathsf{real}}$
  end while
  return $z$
```

```
Algorithm $\mathcal{B}_\infty^{\mathsf{ideal}}$:
  $z \leftarrow \perp$
  while $z = \perp$ do
      $z \leftarrow \mathcal{A}^{\mathsf{ideal}}$
  end while
  return $z$
```

Fig. 1. Rejection sampling algorithms.

Lemma 1 (Adapted from [Lyu12, Lemma 4.7]). *Assume that $M \geq 1$ and $\varepsilon \in [0, 1/2]$ are such that*

$$\Pr_{z \hookleftarrow p_t}\left(p_t(z) \leq M \cdot p_s(z)\right) \geq 1 - \varepsilon,$$

which can be rewritten in terms of smooth Rényi divergence as $R_\infty^\varepsilon(p_t \| p_s) \leq M$. Then the probability $\mathcal{A}^{\mathsf{real}}(\perp)$ that $\mathcal{A}^{\mathsf{real}}$ aborts is such that

$$\frac{M-1}{M} \leq \mathcal{A}^{\mathsf{real}}(\perp) \leq \frac{M-1+\varepsilon}{M}.$$

Moreover, we have

$$\Delta(\mathcal{A}^{\mathsf{real}}, \mathcal{A}^{\mathsf{ideal}}) \leq \varepsilon/M \quad \text{and} \quad \Delta(\mathcal{B}_\infty^{\mathsf{real}}, \mathcal{B}_\infty^{\mathsf{ideal}}) \leq \varepsilon.$$

2.2 Lyubashevsky's signature scheme

All the following parameters are functions of a security parameter λ. We let $k, m, n \geq 1$ and $q \geq 2$ specify matrix spaces over $\mathbb{Z}_q$, with $m > n$. The distribution $P_{\mathbf{S}}$ over $\mathbb{Z}^{m \times k}$ is for signing keys and has support $\mathcal{S} = \mathrm{Supp}(P_{\mathbf{S}})$. Let $\mathcal{M}$ be the message space. Let $\mathcal{C} \subset \mathbb{Z}^k$ finite and $H : \mathbb{Z}_q^n \times \mathcal{M} \to \mathcal{C}$ a hash function, which is modeled as a random oracle in the signature scheme analysis. The parameter $\gamma > 0$ is used in the verification algorithm to quantify the smallness of vectors corresponding to valid signatures. To obtain a 2^λ security against known attacks, one typically sets $m, n, k = \Omega(\lambda)$ and $\gamma, q = \mathsf{poly}(\lambda)$.

Let $\varepsilon \geq 0$ and $M \geq 1$ be parameters related to rejection sampling, for a source distribution Q and a target distribution P over $\mathbb{Z}^m$. Most works directly instantiate these distributions. For example, uniform distributions in well-chosen hypercubes are used in [Lyu09] and $P = Q$ Gaussian are used in [Lyu12]. We assume that the support of Q is contained in $(-q/2, q/2]^m$.

We consider the scheme presented in Figure 2, borrowed from [Lyu12] with the aforementioned rejection sampling generalization. For simplicity, we assume that the verification key $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$ is in Hermite normal form, i.e., we have $\mathbf{A} = (\mathbf{I}_n | \mathbf{B})$ for some matrix $\mathbf{B}$ and with $\mathbf{I}_n \in \mathbb{Z}_q^{n \times n}$ denoting the identity matrix. Up to mild conditions on k, n, m, q, this is without loss of generality.

```
KeyGen($1^\lambda$) :
1: $\mathbf{B} \hookleftarrow \mathbb{Z}_q^{n \times (m-n)}$ and $\mathbf{S} \hookleftarrow P_{\mathbf{S}}$
2: $\mathbf{A} \leftarrow (\mathbf{I}_n | \mathbf{B})$
3: $\mathbf{T} \leftarrow \mathbf{AS}$
4: return $\mathsf{vk} = (\mathbf{A}, \mathbf{T})$ and $\mathsf{sk} = (\mathbf{A}, \mathbf{S})$
```

```
Sign($\mu, \mathbf{A}, \mathbf{S}$) :
1: $\mathbf{y} \hookleftarrow Q$
2: $\mathbf{c} \leftarrow H(\mathbf{Ay}, \mu)$
3: $\mathbf{z} \leftarrow \mathbf{y} + \mathbf{Sc}$
4: $u \hookleftarrow U([0,1])$
5: if $u \leq \min\left(\frac{P(\mathbf{z})}{M \cdot Q(\mathbf{y})}, 1\right)$ then
6:     return $(\mathbf{z}, \mathbf{c})$
7: else
8:     go to Step 1
9: end if
```

```
Verify($\mu, \mathbf{z}, \mathbf{c}, \mathbf{A}, \mathbf{T} = \mathbf{AS}$) :
1: if $\|\mathbf{z}\| \leq \gamma$ and $\mathbf{c} = H(\mathbf{Az} - \mathbf{Tc}, \mu)$
   then
2:     return 1
3: else
4:     return 0
5: end if
```

Fig. 2. Lyubashevsky's signature scheme.

Runtime and correctness follow from the two lemmas below.

Lemma 2 (Sign Runtime). *Let $\varepsilon \geq 0$, $M \geq 1$ and $B = \lceil \lambda / \log \frac{M}{M-1+\varepsilon} \rceil$. Assume that P and Q satisfy $\max_{(\mathbf{S},\mathbf{c}) \in \mathcal{S} \times \mathcal{C}} R_\infty^\varepsilon(P \| Q_{+\mathbf{Sc}}) \leq M$. Let $(\mathbf{y}_0^\top | \mathbf{y}_1^\top)^\top \hookleftarrow$*

Q, where $\mathbf{y}_0$ takes values in $\mathbb{Z}^n$. In the ROM, the number of loop iterations i^* of a Sign *execution satisfies*

$$\forall i : \Pr(i^* \geq i) \leq \left(1 - \frac{1-\varepsilon}{M}\right)^i + B^2 \cdot 2^{-H_\infty(\mathbf{y}_0|\mathbf{y}_1)_Q} + 2^{-\lambda}.$$

Note that when $M \leq \mathsf{poly}(\lambda)$, $\varepsilon \leq 1 - 1/\mathsf{poly}(\lambda)$ and $2^{-H_\infty(\mathbf{y}_0|\mathbf{y}_1)_Q} \leq \mathsf{negl}(\lambda)$, we have that $B^2 \cdot 2^{-H_\infty(\mathbf{y}_0|\mathbf{y}_1)_Q} + 2^{-\lambda} \leq \mathsf{negl}(\lambda)$.

Lemma 3 (Correctness). *Let $\varepsilon \geq 0$ and $M \geq 1$. Let P and Q satisfying $\max_{(\mathbf{S},\mathbf{c})\in\mathcal{S}\times\mathcal{C}} R^\varepsilon_\infty(P\|Q_{+\mathbf{Sc}}) \leq M$. Let $(\mathbf{y}_0^\top|\mathbf{y}_1^\top)^\top \hookleftarrow Q$, where $\mathbf{y}_0$ takes value in $\mathbb{Z}^n$. Further assume that $2^{-H_\infty(\mathbf{y}_0|\mathbf{y}_1)_Q} \leq \mathsf{negl}(\lambda)$, $\varepsilon \leq \mathsf{negl}(\lambda)$ and the probability that* Sign *terminates is $\geq 1 - \mathsf{negl}(\lambda)$. Then, in the ROM, the scheme is correct if $\gamma \geq \gamma_P$ with γ_P such that $\Pr_{\mathbf{z}\hookleftarrow P}(\|\mathbf{z}\| \geq \gamma_P) \leq \mathsf{negl}(\lambda)$.*

We only highlight components of typical security proofs that are relevant to our work, and refer to prior works for more details [Lyu09,Lyu12,AFLT16]. The security proofs of Lyubashevsky's signature scheme all proceed by sequences of games and argue that the adversary's advantages in successive games differ by small amounts and that no efficient adversary can solve the last game with a significant advantage.

An early step in the sequence of games is to replace the calls to H at Step 2 of the Sign algorithm by truly uniform and independent samples $\mathbf{c} \leftarrow U(\mathcal{C})$. To ensure that the adversary cannot notice the difference in the ROM, this requires that a given input $(\mathbf{Ay}, \mu)$ to H cannot occur twice. This is obtained by having the conditional min-entropy $H_\infty(\mathbf{y}_0|\mathbf{y}_1)_Q$ satisfy:

$$H_\infty(\mathbf{y}_0|\mathbf{y}_1)_Q = \Omega(\lambda).$$

An important other game hop consists in making Steps 1 to 6 of the Sign algorithm signing-key independent. Concretely, this means arguing that the distributions of the pair $(\mathbf{z}, \mathbf{c})$ in the experiments from Figure 3 are statistically close, by using Lemma 1. (Note that this also requires programming H consistently with all appearing $\mathbf{c}$'s.)

To complete the security proof, Lyubashevsky [Lyu12] reduces the SIS problem to the sEU-CMA security of a signing-key independent simulation of the Sign algorithm, by relying on the forking lemma. At this stage of the security proof, rejection sampling does not play a role anymore. We only note that the SIS instance has parameters q, m, n and $\beta = 2(\gamma + \gamma')$, with γ as in the Verify algorithm and $\gamma' = \max_{(\mathbf{S},\mathbf{c})\in\mathcal{S}\times\mathcal{C}}\|\mathbf{Sc}\|$. Note that γ is always significantly larger than γ'. We stress that there is a tension in setting γ: it should be sufficiently high to provide correctness (see Lemma 3 above) and as small as possible to provide higher security and hence allow more compact instantations.

3 Lower Bounds in the Case of Perfect Rejection Sampling

We start by studying *the case of perfect rejection sampling*, which corresponds to the setting of [Lyu09,DDLL13]. That is, we set $\varepsilon = 0$ in the formalism of

```
$\mathbf{y} \hookleftarrow Q$
$\mathbf{c} \leftarrow U(\mathcal{C})$
$\mathbf{z} \leftarrow \mathbf{y} + \mathbf{Sc}$
$u \hookleftarrow U([0,1])$
if $u \leq \min\left(\frac{P(\mathbf{z})}{M \cdot Q(\mathbf{y})}, 1\right)$ then
    return $(\mathbf{z}, \mathbf{c})$
else
    return $(\perp, \perp)$
end if
```

```
$\mathbf{c} \leftarrow U(\mathcal{C})$
$\mathbf{z} \leftarrow P$
$u \hookleftarrow U([0,1])$
if $u \leq \frac{1}{M}$ then
    return $(\mathbf{z}, \mathbf{c})$
else
    return $(\perp, \perp)$
end if
```

Fig. 3. Simulating signatures.

Section 2.2. We prove two lower bounds: (1) regarding signature size in both unimodal and bimodal settings (Sections 3.1 and 3.2), and (2) regarding the expected number of iterations of the rejection step (Section 3.3).

First, we analyze to which extent the expected norm of a distribution P can be decreased, under the constraint that we can reject to it using shifted samples from Q, where the Euclidean norm of the shift is bounded from above. This gives lower bounds on the norm of the signature vector $\mathbf{z}$ in Lyubashevsky's signature scheme, as recalled in Section 2.2. We start by studying the easier case of continuous distributions, and then provide a way to discretize the results.

Second, we prove than the classical rejection sampling strategy described above is optimal if one aims to minimize the expected number of iterations of the rejection step in the case of perfect rejection sampling from P to Q. Specifically, the expected number of iterations of any strategy is at least $R_\infty(P\|Q)$, which is reached by classical rejection sampling.

3.1 Optimal Compactness in the Unimodal Setting

The main result of this subsection is the following.

Theorem 1. *Let $m \geq 1, t > 0$, $V = \mathcal{B}_m(t)$ and $M > 1$. Let $f, g : \mathbb{R}^m \to [0,1]$ be two probability densities over $\mathbb{R}^m$ such that $\sup_{\mathbf{v}\in V} R_\infty(f\|g_{+\mathbf{v}}) \leq M$. Then we have:*

$$\mathbb{E}_{\mathbf{x}\hookleftarrow f}(\|\mathbf{x}\|) \geq \frac{t}{M^{1/(m-1)} - 1}.$$

Note that we place ourselves in a setup where shifts belong to a hyperball. In the context of Lyubashesvky's signature scheme, the shift is $\mathbf{Sc}$, where $\mathbf{S}$ is the signing key and $\mathbf{c}$ is the challenge (which is part of the signature). Given that $\mathbf{S}$ is unknown, replacing the set of $\mathbf{Sc}$'s by a hyperball seems to be a reasonable approach. Refining this approximation would lead to significant difficulties in the proof, with unlikely gains.

We now discuss the parameters M and m. As exhibited in Lemma 2, the variable M is related to the rejection probability. The smaller M, the faster we expect signing to be. To obtain a signing algorithm that terminates in polynomial time with overwhelming probability, we are interested in $M \leq \mathsf{poly}(\lambda)$. Recall that we have $m = \Omega(\lambda)$. In this parameter regime, we have $t/(M^{1/(m-1)} - 1) \approx t(m-1)/\log M$.

The role of distribution g in Theorem 1 may seem puzzling, as it does not appear in the result. It acts as a control of the discrepancy of f: distribution f must be sufficiently wide to hide (in the Rényi divergence sense) a version of V that is blurred by g. This forces $\mathbb{E}_{\mathbf{x} \hookleftarrow f}(\|\mathbf{x}\|)$ to be rather large. The proof proceeds in two steps. The first one consists in showing that there is no point favoring any direction and that we can restrict the study to isotropic distributions, i.e., distributions whose density is a function of the norm of the vector. The proof, which may be found in the full version, proceeds by averaging on shells. Theorem 1 is then obtained by integrating the local constraint $\sup_{\mathbf{v} \in V} R_\infty(f \| g_{+\mathbf{v}}) \leq M$ over the whole support, with appropriate scaling.

Lemma 4. *Let $m \geq 1, t > 0$ and $V = \mathcal{B}_m(t)$. Let $f, g : \mathbb{R}^m \to [0,1]$ be two probability densities over $\mathbb{R}^m$ and define $M = \sup_{\mathbf{v} \in V} R_\infty(f \| g_{+\mathbf{v}})$. Then there exist two probability densities f^*, g^* that satisfy*

- $\sup_{\mathbf{v} \in V} R_\infty(f^* \| g^*_{+\mathbf{v}}) \leq M$,
- $\|\mathbf{x}\| = \|\mathbf{y}\| \implies g^*(\mathbf{x}) = g^*(\mathbf{y})$ *and* $f^*(\mathbf{x}) = f^*(\mathbf{y})$,
- $\mathbb{E}_{\mathbf{z} \hookleftarrow f}(\|\mathbf{z}\|) = \mathbb{E}_{\mathbf{z} \hookleftarrow f^*}(\|\mathbf{z}\|)$.

Proof (Theorem 1). Thanks to Lemma 4, we can, without loss of generality, assume that both f and g are isotropic. For $k \geq 0$, we define $\mu_k = \int_0^\infty r^k f(r)\,\mathrm{d}r$, which is the k-th order moment of f. In particular, we have $\mu_{m-1} = 1/S_m$ and $\mu_m = \mathbb{E}_{\mathbf{x} \hookleftarrow f}(\|\mathbf{x}\|)/S_m$. Indeed, using a hyperspherical variable change, we see that, for any $\beta \in \{0,1\}$:

$$\begin{aligned}
\mathbb{E}_{\mathbf{x} \hookleftarrow f}(\|\mathbf{x}\|^\beta) &= \int_{\mathbb{R}^m} \|\mathbf{x}\|^\beta f(\mathbf{x})\,\mathrm{d}\mathbf{x} \\
&= \int_0^\infty \rho^{m-1+\beta} f(\rho) \int_{[0,\pi]^{m-2} \times [0,2\pi]} D(\vec{\theta})\,\mathrm{d}\vec{\theta}\,\mathrm{d}\rho \\
&= S_m \cdot \mu_{m-1+\beta}.
\end{aligned}$$

The above implies that $\mathbb{E}_{\mathbf{x} \hookleftarrow f}(\|\mathbf{x}\|) = \mu_m/\mu_{m-1}$.

For any $x \geq 0$ and $u \in [-t,t]$, it holds that $f(x) \leq M \cdot g(|x-u|)$. In particular, for $x \geq t$, we have $f(x-t) \leq M \cdot g(x)$. Let us multiply both sides by x^{m-1} and integrate over $[t, +\infty)$. With a change of variable on the left-hand side, this gives

$$\begin{aligned}
\int_0^\infty (x+t)^{m-1} f(x)\,\mathrm{d}x &\leq M \cdot \int_t^\infty x^{m-1} g(x)\,\mathrm{d}x \\
&\leq M \cdot \int_0^\infty x^{m-1} g(x)\,\mathrm{d}x \\
&= M \cdot \int_0^\infty x^{m-1} f(x)\,\mathrm{d}x,
\end{aligned}$$

by recognizing that the right-hand side is $M \cdot \mu_{m-1}$ (which is the same for f and g). Grouping everything on the same side, we have

$$0 \leq \int_0^\infty \left(Mx^{m-1} - (x+t)^{m-1} \right) f(x) \, \mathrm{d}x. \tag{1}$$

Let $C = t/(M^{1/(m-1)} - 1)$. For $m > 2$, we rewrite the integrand as

$$\begin{aligned} Mx^{m-1} - (x+t)^{m-1} &= \left(M^{\frac{1}{m-1}} x - (x+t) \right) \cdot \sum_{k=0}^{m-2} \left(x M^{\frac{1}{m-1}} \right)^k (x+t)^{m-2-k} \\ &= \left(M^{\frac{1}{m-1}} - 1 \right) (x - C) \cdot \sum_{k=0}^{m-2} \left(x M^{\frac{1}{m-1}} \right)^k (x+t)^{m-2-k}. \end{aligned}$$

For $m = 2$, the above holds by replacing the sum by 1. Now, note that the inequality $xM^{1/(m-1)} \geq x + t$ holds if and only if $x \geq C$. Hence the following upper bound holds for any $x \geq 0$, if $m > 2$:

$$(x - C) \cdot \sum_{k=0}^{m-2} (x M^{\frac{1}{m-1}})^k (x+t)^{m-2-k} \leq (x - C)(m-1) M^{\frac{m-2}{m-1}} x^{m-2}.$$

When $m > 2$, we can divide by $(M^{1/(m-1)} - 1) M^{(m-2)/(m-1)} (m-1) > 0$ in Equation (1), and obtain:

$$C \cdot \int_0^\infty x^{m-2} f(x) \, \mathrm{d}x \leq \int_0^\infty x^{m-1} f(x) \, \mathrm{d}x.$$

Note that it also holds for $m = 2$. This can be rewritten as $\mu_{m-1}/\mu_{m-2} \geq C$.

Now, observe that $\mu_m \cdot \mu_{m-2} \geq (\mu_{m-1})^2$. Indeed, the Cauchy-Schwarz inequality states that for any real random variables X, Y, it holds that $|\mathbb{E}(XY)|^2 \leq \mathbb{E}(X^2)\mathbb{E}(Y^2)$. We instantiate it with the (non-independent) random variables $X = \|\mathbf{x}\|^{m/2}$ and $Y = \|\mathbf{x}\|^{(m-2)/2}$, where $\mathbf{x} \hookleftarrow f$. Then $XY = \|\mathbf{x}\|^{\frac{m}{2} + \frac{m-2}{2}} = \|\mathbf{x}\|^{m-1}$. To conclude, note $\mu_m \cdot \mu_{m-2} \geq (\mu_{m-1})^2$ implies that $\mu_m/\mu_{m-1} \geq \mu_{m-1}/\mu_{m-2} \geq C$. This completes the proof. □

For the discrete case, we observe that given a discrete distribution P, letting $f : \mathbf{x} \mapsto P(\lceil \mathbf{x} \rfloor)$ be a probability density over $\mathbb{R}^m$, we have, by the triangle inequality

$$\mathbb{E}_{\mathbf{x} \hookleftarrow f}(\|\mathbf{x}\|) \leq \mathbb{E}_{\mathbf{x} \hookleftarrow P}(\|\mathbf{x}\|) + \frac{\sqrt{m}}{2}.$$

Theorem 1 can then be adapted to the discrete case, up to subtracting $\frac{\sqrt{m}}{2}$ from the lower bound. In all setups considered in this work, this term is significantly smaller than $t/(M^{1/(m-1)} - 1)$.

Corollary 1. *Let $m \geq 1$, $t > 0$, $V = \mathcal{B}_m(t) \cap \mathbb{Z}^m$ and $M > 1$. Let P and Q be two discrete probability distributions over $\mathbb{Z}^m$ such that $\sup_{\mathbf{v} \in V} R_\infty(P \| Q_{+\mathbf{v}}) \leq M$. Then we have:*

$$\mathbb{E}_{\mathbf{x} \hookleftarrow P}(\|\mathbf{x}\|) \geq \frac{t}{M^{1/(m-1)} - 1} - \frac{\sqrt{m}}{2}.$$

3.2 Optimal Compactness in the Bimodal Setting

We obtain the following result in the bimodal setting. The proof, which is similar to the one of Theorem 1, it provided in the full version.

Theorem 2. *Let $m \geq 3, t > 0$, $V = \mathcal{B}_m(t)$ and $M > 1$. Let $f, g : \mathbb{R}^m \to [0,1]$ be two probability densities over $\mathbb{R}^m$ such that $\sup_{\mathbf{v}\in V} R_\infty(f\|g_{\pm\mathbf{v}}) \leq M$, where $g_{\pm\mathbf{v}}$ is the density $\mathbf{x} \mapsto \frac{1}{2}(g(\mathbf{x}-\mathbf{v}) + g(\mathbf{x}+\mathbf{v}))$.Then the following holds:*

$$\mathbb{E}_{\mathbf{x}\hookleftarrow f}(\|\mathbf{x}\|) \geq \frac{t}{\sqrt{M^{\frac{2}{m-2}} - 1}}.$$

For $M \leq \mathsf{poly}(\lambda)$ and $m = \Omega(\lambda)$ as in the discussion following Theorem 1, we have $t/(M^{2/(m-2)} - 1)^{1/2} \approx t\sqrt{(m-2)/(2\log M)}$. Similarly to the unimodal case, the lower bound can be adapted to integer distributions with limited loss (for all setups considered in this work).

Corollary 2. *Let $m \geq 3, t > 0, V = \mathcal{B}_m(t) \cap \mathbb{Z}^m$ and $M > 1$. Let P and Q be two discrete probability distributions over $\mathbb{Z}^m$ such that $\sup_{\mathbf{v}\in V} R_\infty(P\|Q_{\pm\mathbf{v}}) \leq M$, where $Q_{\pm\mathbf{v}}$ is as in Theorem 2. Then the following holds:*

$$\mathbb{E}_{\mathbf{x}\hookleftarrow P}(\|\mathbf{x}\|) \geq \frac{t}{\sqrt{M^{\frac{2}{m-2}} - 1}} - \frac{\sqrt{m}}{2}.$$

3.3 Optimality of the Expected Number of Iterations

We now analyze to which extent the expected number of iterations of the rejection step could be reduced in the case of exact rejection sampling from P to Q, and prove the classical strategy to be optimal. This question arises from the variety of rejection sampling techniques that have been studied in other fields.

There exist multiple variants of rejection sampling. For instance, a procedure described in [HJMR07], and recalled in the full version, takes a greedy approach to rejection sampling and differs from the one we presented up until now. We are in the setting where we have access to a sampler from distribution Q. These samples are denoted by $(X_i)_{i\geq 1}$ with $X_i \in \mathcal{X}$ for some set $\mathcal{X}$ and we are required to output a sample from the distribution P over $\mathcal{X}$. Any design of procedure is allowed, as long as the output is one of the observed samples X_i. Let i^* be the random variable denoting the number of samples observed by an algorithm and we wish to determine how small $\mathbb{E}(i^*)$ can be. We note that the work of [HJMR07], establishes that there exists a rejection sampling algorithm achieving $\mathbb{E}(\log i^*) = \log R_1(P\|Q)$ up to lower order terms in $R_1(P\|Q)$, and that this is optimal. Here, we show that the minimum value for $\mathbb{E}(i^*)$ is $R_\infty(P\|Q)$.

We model a rejection sampling algorithm by a family of randomized functions $A_i : \mathcal{X}^i \to \{1, \ldots, i\} \cup \{\mathrm{r}\}$. At step i, it sees the new sample X_i and based on $X_1, \ldots, X_i$ it computes $A_i(X_1, \ldots, X_i)$. If it is equal to r, the algorithm asks for one more sample and otherwise if $A_i(X_1, \ldots, X_i) \in \{1, \ldots, i\}$, the algorithm terminates and outputs the sample $X_{A_i(X_1,\ldots,X_i)}$. Note that the running time of the algorithm is defined by $i^* = \inf\{i \geq 1 : A_i(X_1, \ldots, X_i) \neq \mathrm{r}\}$. We only consider algorithms for which $i^* < \infty$ almost surely. We define the random variable $J = A_{i^*}(X_1, \ldots, X_{i^*}) \in \mathbb{N}_+$, note that $J \leq i^*$ and the output of the algorithm is X_J (i.e., the output sample may not be the last one that was generated).

Theorem 3. *Let P, Q be two discrete probability distributions. Any rejection sampling algorithm $(A_i)_{i \geq 1}$ sampling from P satisfies $\mathbb{E}(i^*) \geq R_\infty(P \| Q)$.*

Proof. We have by assumption for any $x \in \mathcal{X}$,

$$P(x) = \Pr[X_J = x] = \sum_{j=1}^{\infty} \Pr[J = j, X_j = x] \leq \sum_{j=1}^{\infty} \Pr[i^* \geq j, X_j = x],$$

where we used the fact that the event $[J = j]$ is contained in $[i^* \geq j]$. Now, observe that the event $[i^* < j]$ only depends on $X_1, \ldots, X_{j-1}$ and as such it is independent of the event $[X_j = x]$. This implies that $[i^* \geq j]$ is independent of $[X_j = x]$. As a result, we have

$$P(x) \leq \sum_{j=1}^{\infty} \Pr[i^* \geq j] \Pr[X_j = x] = \mathbb{E}(i^*) Q(x) ,$$

which proves the desired result.

In the context of Lyubashevsky's signature schemes with source distribution Q', target distribution P', challenge set $\mathcal{C}$ and signing key $\mathbf{S}$, we would have $P = P' \otimes U(\mathcal{C})$ and Q would be the distribution of the pair $(\mathbf{z}, \mathbf{c})$ obtained by sampling $\mathbf{y}$ from Q', $\mathbf{c}$ from $U(\mathcal{C})$ and defining $\mathbf{z} = \mathbf{y} + \mathbf{Sc}$.

The above proof can be adapted in the setting where P and Q are continuous distributions by considering a sequence of balls converging to $\{x\}$ instead of x.

4 Improved Analysis via the Rényi Divergence

For the rest of the paper, we flip our focus and prove positive results (upper bounds). In this section, we propose an improved analysis of Lyubashevsky's signatures that relies on the Rényi divergence rather than the statistical distance, allowing larger sampling errors in the case of imperfect rejection sampling. Then, in Section 5 propose a new choice of distributions that (asymptotically) reaches our lower bounds. Finally, in Section 6, we propose a way to circumvent the lower bound for the expected number of iterations by providing an alternate strategy which allows to fix a-priori a maximal number of loop iterations.

Our lower bounds apply to perfect rejection sampling, but rejecting to an inaccurate approximation to the target distribution also allows to instantiate Lyubashevsky's signature, as done in [Lyu12] and already mentioned in Section 2.2 (when $\varepsilon > 0$ in Lemma 1). In particular, imperfect rejection sampling is used when instantiating the signature scheme with Gaussian distributions.

In this section, we study the case of imperfect rejection sampling and describe a way to improve the analysis of the digital signature from Section 2.2, by replacing the statistical distance (in Lemma 1) with the Rényi divergence to quantify the closeness between ideal and real rejection sampling algorithms. As already observed in prior works (see in particular the discussion in [BLR$^+$18]), the Rényi divergence is well-suited for improving the analyses of digital signatures, as the security game is of a search type. While the analysis based on the statistical distance imposes $\varepsilon = 2^{\Omega(\lambda)}$, as it requires the statistical distance to be negligible, our analysis allows larger sampling errors as it only imposes $\varepsilon \approx 1/q_{\mathsf{sig}}$ where q_{sig} is the number of signing queries (which is $\mathsf{poly}(\lambda) \ll 2^{\Omega(\lambda)}$).

4.1 Rényi Divergence Bounds for Imperfect Rejection Sampling

Let p_t and p_s be two probability densities, both continuous or both discrete. We consider algorithms $\mathcal{A}^{\mathsf{real}}$, $\mathcal{A}^{\mathsf{ideal}}$, $\mathcal{B}^{\mathsf{real}}$ and $\mathcal{B}^{\mathsf{ideal}}$ from Figure 1.

Lemma 5. *Assume that $M > 1$ and $\varepsilon < 1$ are such that $R_\infty^\varepsilon(p_t \| p_s) \leq M$. Then for any $a \in (1, +\infty)$ we have:*

$$R_a(\mathcal{A}^{\mathsf{real}} \| \mathcal{A}^{\mathsf{ideal}}) \leq \left(\frac{1}{M} + \frac{M-1+\varepsilon}{M} \cdot \left(1 + \frac{\varepsilon}{M-1}\right)^{a-1} \right)^{\frac{1}{a-1}},$$

$$R_a(\mathcal{B}_\infty^{\mathsf{real}} \| \mathcal{B}_\infty^{\mathsf{ideal}}) \leq \frac{1}{(1-\varepsilon)^{a/(a-1)}}.$$

Moreover, for $a = \infty$, we have:

$$R_\infty(\mathcal{A}^{\mathsf{real}} \| \mathcal{A}^{\mathsf{ideal}}) \leq 1 + \frac{\varepsilon}{M-1} \quad \text{and} \quad R_\infty(\mathcal{B}_\infty^{\mathsf{real}} \| \mathcal{B}_\infty^{\mathsf{ideal}}) \leq \frac{1}{1-\varepsilon}.$$

Note that for $\varepsilon = 0$, we recover the distributional equalities $\mathcal{A}^{\mathsf{real}} = \mathcal{A}^{\mathsf{ideal}}$ and $\mathcal{B}_\infty^{\mathsf{real}} = \mathcal{B}_\infty^{\mathsf{ideal}}$ of Lemma 1. We are interested in the case $\varepsilon > 0$.

Proof. Let $\mathcal{A}^{\mathsf{real}}(\perp)$ and $\mathcal{A}^{\mathsf{ideal}}(\perp)$ denote the probabilities that $\mathcal{A}^{\mathsf{real}}$ or $\mathcal{A}^{\mathsf{ideal}}$ output nothing. We have, using results from Lemma 1:

$$\begin{aligned}
R_a(\mathcal{A}^{\mathsf{real}} \| \mathcal{A}^{\mathsf{ideal}})^{a-1} &= \left[\int_{\mathrm{Supp}(p_s)} \frac{\left(p_s(x) \min\left(\frac{p_t(x)}{M \cdot p_s(x)}, 1\right)\right)^a}{(p_t(x)/M)^{a-1}} \, \mathrm{d}x \right] + \frac{(\mathcal{A}^{\mathsf{real}}(\perp))^a}{(\mathcal{A}^{\mathsf{ideal}}(\perp))^{a-1}} \\
&\leq \int_{\mathrm{Supp}(p_s)} \frac{\left(p_s(x) \frac{p_t(x)}{M \cdot p_s(x)}\right)^a}{(p_t(x)/M)^{a-1}} \, \mathrm{d}x + \frac{(1-(1-\varepsilon)/M)^a}{(1-1/M)^{a-1}} \\
&= \int_{\mathrm{Supp}(p_s)} \frac{p_t(x)}{M} \, \mathrm{d}x + \frac{M-1+\varepsilon}{M} \cdot \left(\frac{M-1+\varepsilon}{M-1} \right)^{a-1} \\
&\leq \frac{1}{M} + \frac{M-1+\varepsilon}{M} \cdot \left(1 + \frac{\varepsilon}{M-1}\right)^{a-1}.
\end{aligned}$$

We move on to bounding the second divergence. For any $x \in \mathrm{Supp}(p_s)$:

$$\mathcal{B}_\infty^{\mathsf{real}}(x) = \frac{\mathcal{A}^{\mathsf{real}}(x)}{1 - \mathcal{A}^{\mathsf{real}}(\perp)}.$$

This also holds for $\mathcal{B}_\infty^{\mathsf{ideal}}$ with $\mathcal{A}^{\mathsf{ideal}}$ instead of $\mathcal{A}^{\mathsf{real}}$. We obtain:

$$\begin{aligned} R_a(\mathcal{B}_\infty^{\mathsf{real}} \| \mathcal{B}_\infty^{\mathsf{ideal}})^{a-1} &= \int_{\mathrm{Supp}(p_s)} \frac{1}{M^{a-1}} \cdot \frac{\left(p_s(x) \min\left(\frac{p_t(x)}{M \cdot p_s(x)}, 1\right)\right)^a}{(\mathcal{A}^{\mathsf{real}}(\perp))^a (p_t(x)/M)^{a-1}} \\ &\leq \frac{M}{(1-\varepsilon)^a} \int_{\mathrm{Supp}(p_s)} \frac{\left(p_s(x) \min\left(\frac{p_t(x)}{M \cdot p_s(x)}, 1\right)\right)^a}{(p_t(x)/M)^{a-1}}. \end{aligned}$$

This sum was already computed just above and is at most $1/M$.

The continuity of $a \mapsto R_a(P_t \| P_s)$ at $a = +\infty$ gives the last bounds. □

4.2 Improved Analysis of Lyubashevsky's Scheme

We now go back to the scheme described in Section 2.2 with imperfect rejection sampling, and show that the analysis above allows setting $\varepsilon \approx 1/q_{\mathsf{sig}}$ instead of $\varepsilon = 2^{-\Omega(\lambda)}$. Here q_{sig} refers to the number of signing queries that an adversary can make. As a signing query requires an interaction with the signer, it is typically considered to be a large polynomial in λ, which is much smaller than $2^{\Omega(\lambda)}$. As a result, this refined analysis puts less stress on the condition that P_s and P_t must satisfy and hence to reach smaller values for $\mathbb{E}_{\mathbf{z} \hookleftarrow P_t}(\|\mathbf{x}\|)$: this is beneficial to security and then allows for smaller parameter sets.

To achieve this improvement, we replace the statistical distance with the Rényi divergence in the scheme analysis, when simulating signature queries (see Figure 3). By Lemma 5 and the Rényi divergence data processing inequality, replacing $\mathcal{A}^{\mathsf{real}}$ by $\mathcal{A}^{\mathsf{ideal}}$ once in the security proof (i.e., in one loop iteration of one signature query) leads to a multiplicative loss of a factor $\leq 1 + \varepsilon/(M-1)$ in the adversary's advantage. Now, note that the probability that at least one among the q_{sig} sign queries requires more than $B = (\lambda + \log q_{\mathsf{sig}}) / \log(M/(M-1+\varepsilon))$ loop iterations is exponentially small. Assuming this is not the case, we can bound the number of times $\mathcal{A}^{\mathsf{real}}$ is replaced by $\mathcal{A}^{\mathsf{ideal}}$ in the security proof by $B \cdot q_{\mathsf{sig}}$. By the Rényi divergence multiplicativity property, this induces a multiplicative loss of a factor $\leq (1 + \varepsilon/(M-1))^{B \cdot q_{\mathsf{sig}}}$ in the adversary's advantage.

5 Reaching the Lower Bounds with Hyperballs

In this section, we show that continuous uniform distributions in hyperballs reach the lower bounds in both the unimodal and bimodal perfect rejection sampling settings. We also consider the imperfect unimodal setting and find parameters that are asymptotically at least as good as the ones obtained for the Gaussian distribution (using our analysis described in Section 4). As continuous hyperball uniform distributions are easier both to study and implement than their discrete counterpart, we consider the case of continuous distributions.Further, we show that a slight modification of Lyubashevsky's signature allows for the target and source distributions to be continuous.

We also compare this choice of distributions with the uniform distributions in hypercubes and with Gaussians, both asymptotically and with concrete parameters.

5.1 Uniform Distributions in Hyperballs

The first step is to compute the divergence in the three settings: unimodal, either perfect or imperfect rejection sampling and bimodal perfect rejection sampling. The first case can actually be seen as a particular case of the second one, and we summarize both in the following lemma. We use the notation $I_x(a,b) = B(x;a,b)/B(a,b)$ for $x \in [0,1]$ and $a,b > 0$, where $B(a,b)$ is the Beta function and $B(x;a,b)$ is the regularized incomplete Beta function.

Lemma 6 (Smooth Divergence). *Let $m \geq 1$ and $\mathbf{v} \in \mathbb{R}^m$. Let $\varepsilon \in [0,1/2)$ and $\eta \geq 1$ be such that $2\varepsilon = I_{1-1/\eta^2}(\frac{m+1}{2}, \frac{1}{2})$. Let $r, r' > 0$ such that $r'^2 \geq r^2 + \|\mathbf{v}\|^2 + 2r\|\mathbf{v}\|/\eta$. Then it holds that:*

$$R_\infty^\varepsilon\Big(U(\mathcal{B}_m(r)) \| U(\mathcal{B}_m(r', \mathbf{v}))\Big) = \Big(\frac{r'}{r}\Big)^m.$$

Let $M > 1$. The above is $\leq M$ if $r \geq \|\mathbf{v}\| \cdot \frac{\frac{1}{\eta} + \sqrt{\frac{1}{\eta^2} + M^{2/m} - 1}}{M^{2/m} - 1}$ and $r' = M^{1/m} r$.

Note that when $\varepsilon = 0$, we have $\eta = 1$. In that case, we can set $r = \|\mathbf{v}\|/(M^{1/m} - 1)$, which almost matches the lower bound from Theorem 1. For $\varepsilon = 2^{-c \cdot m}$ with a constant $c > 0$, we have that $1/\eta^2$ tends to $1 - 2^{-c}$ when m goes to infinity; for ε satisfying $\varepsilon \geq 2^{-o(m)}$ and $\varepsilon = o(1/m)$ with m going to infinity, we have that $1/\eta^2 \sim -\log(\varepsilon)/m$ (see full version).

Proof. Assume that there exists some cut $\mathcal{C}$ with $\mathrm{vol}(\mathcal{C})/V_m(r) \leq \varepsilon$ such that the divergence is defined, i.e., with $\mathcal{B}_m(r) \setminus \mathcal{C} \subseteq \mathcal{B}_m(r', \mathbf{v})$. Then the divergence is $(r'/r)^m$, as the ratio of densities is constant and equal to $(r'/r)^m$ over $\mathcal{B}_m(r) \setminus \mathcal{C}$. To prove the first claim, it hence suffices to show that such a cut $\mathcal{C}$ exists.

We introduce the cut $\mathcal{C}_\eta := \{\mathbf{x} \in \mathcal{B}_m(r) | \langle \mathbf{x}, \mathbf{v} \rangle \geq -\|\mathbf{v}\| r/\eta\}$. This is the intersection of a ball with an affine half-space, i.e., an m-dimensional hyperspherical cap. Its volume is $\frac{V_m(r)}{2} \cdot I_{1-1/\eta^2}(\frac{m+1}{2}, \frac{1}{2})$ (see full version). The definition of η ensures that $\mathrm{vol}(\mathcal{C}_\eta)/V_m(r) = \varepsilon$. We now check that $\mathcal{B}_m(r) \setminus \mathcal{C}_\eta \subseteq \mathcal{B}_m(r', \mathbf{v})$. Let $\mathbf{x} \in \mathcal{B}_m(r) \setminus \mathcal{C}_\eta$. We have

$$\|\mathbf{x} - \mathbf{v}\| \leq \sqrt{r^2 + \|\mathbf{v}\|^2 + 2r\|\mathbf{v}\|/\eta}.$$

By assumption, the latter is no larger than r', implying that $\mathbf{x} \in \mathcal{B}_m(r', \mathbf{v})$. This completes the proof of the first claim.

If we combine the condition on r and r' and the equality $r' = M^{1/m} r$, we get

$$r^2 + \|\mathbf{v}\|^2 + 2\frac{r\|\mathbf{v}\|}{\eta} \leq M^{2/m} r^2,$$

which is a degree-2 inequality on r. Solving it completes the proof. □

Lemma 7 (Divergence in the Bimodal Setting). *Let $m \geq 1$ and $\mathbf{v} \in \mathbb{R}^m$. Let $r, r' > 0$ such that $r'^2 \geq r^2 + \|\mathbf{v}\|^2$. Let $U(\mathcal{B}_m(r'), \pm\mathbf{v})$ denote the continuous probability distribution which samples $b \hookleftarrow U(\{0,1\})$ and returns $\mathbf{z} \hookleftarrow U(\mathcal{B}_m(r', (-1)^b \mathbf{v}))$. Then it holds that:*

$$R_\infty\Big(U(\mathcal{B}_m(r)) \| U(\mathcal{B}_m(r'), \pm\mathbf{v})\Big) = \big(1 + \chi_{<r+\|\mathbf{v}\|}(r')\big) \cdot \Big(\frac{r'}{r}\Big)^m,$$

where $\chi_{<r+\|\mathbf{v}\|}$ denotes the indicator function of reals smaller than $r + \|\mathbf{v}\|$. Let $M > 1$. The above is $\leq M$ if $r \geq \|\mathbf{v}\|/\sqrt{(M/2)^{2/m} - 1}$ and $r' = (M/2)^{1/m} r$.

Note that the choice of r almost matches the lower bound from Theorem 2.

Proof. The support of $U(\mathcal{B}_m(r'), \pm\mathbf{v})$ is exactly $\mathcal{B}_m(r', \mathbf{v}) \cup \mathcal{B}_m(r', -\mathbf{v})$ and its density is $\mathbf{z} \mapsto (\chi_{\mathcal{B}_m(r',\mathbf{v})}(\mathbf{z}) + \chi_{\mathcal{B}_m(r',-\mathbf{v})}(\mathbf{z}))/(2V_m(r'))$. The divergence is finite when $\mathcal{B}_m(r) \subseteq \mathcal{B}_m(r', \mathbf{v}) \cup \mathcal{B}_m(r', -\mathbf{v})$. This is the case if any $\mathbf{x}$ with $\|\mathbf{x}\| \leq r$ satisfies $\|\mathbf{x} - \mathbf{v}\| \leq r'$ or $\|\mathbf{x} + \mathbf{v}\| \leq r'$. Let us assume, w.l.o.g., that $\|\mathbf{x} - \mathbf{v}\| \leq \|\mathbf{x} + \mathbf{v}\|$. Then we write

$$\|\mathbf{x} - \mathbf{v}\| = \sqrt{\|\mathbf{x}\|^2 + \|\mathbf{v}\|^2 - 2\langle \mathbf{x}, \mathbf{v}\rangle} \leq \sqrt{\|\mathbf{x}\|^2 + \|\mathbf{v}\|^2}.$$

Thanks to the assumption on r and r', we conclude that the divergence is finite.

Now, the ratio of the densities only takes three values. If $\mathbf{x} \notin \mathcal{B}_m(r)$ then the ratio is 0. If $\mathbf{x} \in \mathcal{B}_m(r) \cap \mathcal{B}_m(r', \mathbf{v}) \cap \mathcal{B}_m(r', -\mathbf{v})$ then the ratio is $(r'/r)^m$. Finally, if $\mathbf{x}$ belongs to $\mathcal{B}_m(r) \cap \mathcal{B}_m(r', \mathbf{v})$ but not to $\mathcal{B}_m(r', -\mathbf{v})$, then the ratio is $2(r'/r)^m$. This last case only occurs if $\mathcal{B}_m(r) \not\subseteq \mathcal{B}_m(r', -\mathbf{v})$. This is the case only if $r' < r + \|\mathbf{v}\|$. This completes the proof of the first claim.

For the second claim, note that the assumption on r and r' is satisfied, and that the divergence bound is indeed $\leq M$. □

Finally, in order to use the uniform distribution in a hyperball, we verify that there is sufficient min-entropy in the first n coordinates given the remaining $m-n$ coordinates. The proof of the following lemma can be found in the full version.

Lemma 8. *Let $m \geq 6, n \geq 1$ and $r \geq 2\sqrt{m}$. Let $\mathbf{x} = (\mathbf{x}_0^\top | \mathbf{x}_1^\top)^\top$ be a random variable over $\mathbb{R}^m$ whose distribution is $U(\mathcal{B}_m(r))$, where $\mathbf{x}_0$ has dimension n. It holds that*

$$H_\infty\big(\lceil \mathbf{x}_0 \rfloor | \lceil \mathbf{x}_1 \rfloor\big)_{U(\mathcal{B}_m(r))} \geq \Big(\log_2 \frac{1}{0.85}\Big) \cdot n\ .$$

5.2 Lyubashevsky's Signature with Continuous Distributions

We consider continuous distributions over hyperballs, which are not directly compatible with Lyubashevsky's signature scheme, as recalled in Section 2. Switching to uniform distributions over the integer points inside hyperballs leads to several difficulties: sampling from such a distribution seems delicate, in particular if the radius of the ball is moderate. Similarly, adapting Lemmas 6 and 7 seems difficult. Rather, we argue that it is possible to extend Lyubashevsky's signature scheme to the case of continuous distributions, and that this comes with very limited complications (in the case of Gaussians, it could be simpler to use continuous Gaussians with this modified scheme, than using discrete Gaussians with the original scheme).

In order to adapt Lyubashevsky's signature scheme to continuous distributions, a rounding step is added after acceptance of a sample, as well as during hashing. Concretely, the changes compared to the construction described in Figure 2 are as follows: (i) $\mathbf{y}$ is now sampled from a continuous distribution with density g, (ii) $\mathbf{c}$ is now computed as $H(\mathbf{A}\lceil\mathbf{y}\rfloor, \mu)$, (iii) with $\mathbf{z}$ still being defined as $\mathbf{y} + \mathbf{Sc}$, if the test passes, and the returned signature is now $(\lceil\mathbf{z}\rfloor, \mathbf{c})$. This adaptation is discussed in more details in the full version. We note that this leads to the requirement that the min-entropy of $\lceil\mathbf{x}_0\rfloor || \lceil\mathbf{x}_1\rfloor$ is large, where $\mathbf{x} = (\mathbf{x}_0^\top | \mathbf{x}_1^\top)^\top$ is a random variable over $\mathbb{R}^m$ whose distribution is g and $\mathbf{x}_0$ has dimension n. In the case of the uniform distribution in a hyperball, this is provided by Lemma 8.

We further remark that this applies to the analysis relying on the statistical distance as well as our improved analysis which relies on the Rényi divergence. Also, we note that the modified scheme involves computations over real numbers. These can be securely replaced by finite precision computations, using standard techniques such as described in [Pre17].

5.3 Comparison with other Distributions

Let $t = \max_{\mathbf{S},\mathbf{c}} \|\mathbf{Sc}\|$. In Table 2, we summarize the expected norm of signatures (up to a constant factor) for diverse distributions P and Q, and for a target expected number of iterations M. We consider three specific pairs of distributions, two of them being previously considered distributions (Gaussians and uniforms in hypercubes), and the last one being uniform distributions in hyperballs, introduced above. We also consider three different scenarios:

- unimodal distributions and perfect rejection sampling, corresponding to the column $\varepsilon = 0$;
- unimodal distributions and imperfect rejection sampling – we use approximations specific to the choice of $\varepsilon \geq 2^{-o(m)}$ and $\varepsilon = o(1/m)$;
- bimodal source distribution and perfect rejection sampling, corresponding to column "Bimodal".

Note that the second scenario relies on our improved analysis relying on the Rényi divergence for the imperfect case (see Section 4). This parameter range for ε is not appropriate when using the analysis relying on the statistical distance.

In the last column, we also emphasize if the test that decides to keep or reject a sample is simple or not. For hyperballs, it simply consists in comparing the norm of the sample with the radius of the target hyperball.

The entries in the table are approximations for $m \to \infty$, $t = \omega(1)$ and $M = 2^{o(m)}$, and for a given choice of P, we optimize the parametrization of Q (e.g., the radius in case of a hyperball) to minimize the signature norm.

Choices for P and Q	$\varepsilon = 0$	$\varepsilon \geq 2^{-o(m)}$ and $\varepsilon = o(1/m)$	Bimodal	Rejection Test
Hypercubes	$\frac{tm^{3/2}}{\log M}$	$\frac{tm^{3/2}}{\log M}$	$\frac{tm^{3/2}}{\log M}$	Simple
Gaussians	∞	$\frac{t\sqrt{m}\sqrt{\log \frac{1}{\varepsilon} + \log M}}{\log M}$	$\frac{t\sqrt{m}}{\sqrt{\log M}}$	Complex
Hyperballs	$\frac{tm}{\log M}$	$\frac{t\sqrt{m}\sqrt{\log \frac{1}{\varepsilon} + \log M}}{\log M}$	$\frac{t\sqrt{m}}{\sqrt{\log M}}$	Simple

Table 2. Expected norm of signatures depending on the choice of distributions and (im)perfectness of rejection sampling.

The values of the table are obtained by computing the parameters for the underlying distributions (radii r, r' of the hypercubes or hyperballs and standard deviation σ of Gaussians) for our constraints M and t. This is done by computing their (smooth) Rényi Divergence, as done in Lemmas 6 and 7 for hyperballs. Proofs for hypercubes and Gaussians can be found in the full version. Given these parameters, the expected norm immediately follows ($r\sqrt{m}$ for a hypercube of radius r, σr for a Gaussian of standard deviation σ, and r for a hyperball of radius r). To conclude this section, we emphasize the following points:

- Gaussians and Hyperballs are asymptotically equivalent and reach the lower bounds in the bimodal setting; Hyperballs further reach our lower bound in the exact unimodal setting as well;
- Hyperballs benefits from a significantly simpler rejection test compared to Gaussians;
- The bimodal setting (in both Gaussian and Hyperballs cases) leads to the most compact signatures.

5.4 Concrete Parameters

To study the concrete impact of the choice of distributions on signature size, we consider Dilithium. The left side of Table 3 shows the parameters for three security levels of the round-3 documentation of the CRYSTALS-Dilithium submission to the NIST post-quantum project [BDK+20]. The right side of Table 3 gives updated parameters for Dilithium-G, a modification of Dilithium using Gaussian distributions whose description is available in the first version of the eprint version of [DKL+18]. For this updated version, we set the value of M to 4.

	Hypercube-Uniform			Previous Gaussian		
	Medium	Recommended	Very High	Medium	Recommended	Very High
Ring dimension ℓ	256	256	256	256	256	256
q	8380417	8380417	8380417	254977	254977	254977
$(n, m-n)$	$(4,4)$	$(6,5)$	$(8,7)$	$(4,3)$	$(5,4)$	$(7,6)$
η	2	4	2	2	3	2
S	N/A	N/A	N/A	91	134	111
τ	39	49	60	39	49	60
$t = S \cdot \sqrt{\tau}$	N/A	N/A	N/A	568	938	860
B	N/A	N/A	N/A	864K	535K	664K
γ_2	$\frac{q-1}{88}$	$\frac{q-1}{32}$	$\frac{q-1}{32}$	$\frac{q-1}{48}$	$\frac{q-1}{24}$	$\frac{q-1}{32}$
d	13	13	13	11	11	11
M	4.25	5.1	3.85	4	4	4
BKZ block-size b to break SIS	423 (417)	638 (603)	909 (868)	450 (390)	677 (588)	1018 (891)
Best known classical bit-cost	123 (121)	186 (176)	265 (253)	131 (114)	198 (171)	297 (260)
Best known quantum bit-cost	112 (110)	169 (159)	241 (230)	119 (103)	179 (155)	270 (236)
BKZ block-size b to break LWE	422	622	860	403	623	1018
Best known classical bit-cost	123	181	251	117	182	297
Best known quantum bit-cost	111	164	228	1076	165	170
Expected signature size	2420	3293	4595	1737	2372	3478
Expected public key size	1312	1952	2592	672	1312	1600

Table 3. Parameters for Dilithium and updated Dilithium-G.

In these schemes, the verification key is a module-LWE sample $\mathbf{B}\mathbf{s}_1 + \mathbf{s}_2$ where $\mathbf{s}_1$ and $\mathbf{s}_2$ have ℓ_∞-norms $\leq \eta$. For each coordinate, the lowest d bits are dropped. A parameter τ is used to control the ℓ_1-norm of any hashed value $\mathbf{c}$, so that $\mathbf{c}$ has sufficient min-entropy. In Dilithium-G, the bound t is $S\sqrt{\tau}$, where S is the median over the key generation randomness of the largest singular value of $(\mathsf{rot}(\mathbf{s}_1)^\top, \mathsf{rot}(\mathbf{s}_2)^\top)^\top$. A rejection step is added in KeyGen to check that the key satisfies this bound. The value of the SIS bound corresponding to unforgeability is computed using [BDK$^+$20, Equation (6)]. The strong unforgeability bound is obtained by multiplying this bound by 2. The security is estimated using block-size optimized BKZ to break the module-SIS or module-LWE instances.[5]

For Dilithium, i.e., the hypercube version, we take $t_\infty = \tau\eta$ as a bound on the ℓ_∞-norm of the secret key, which drives the radius of the hypercube and subsequently the unforgeability SIS bound (in ℓ_∞-norm).

It was argued in [DKL$^+$18] that it seems difficult for BKZ to solve SIS with ℓ_∞-norm bound close to q, i.e., ℓ_2-norm above q. To analyze the runtime of BKZ in the case of an ℓ_2-norm bound $B \geq q$, one can remove the trivial vectors of the input basis (i.e., the vectors with coordinates in $q\mathbb{Z}$) by some randomizing step. This approach was however not considered for Dilithium-G and q was chosen such that $B < q$, leading to bigger parameters overall. Our updated parameters allow for $B \geq q$, for a fairer comparison to Dilithium. We note that for $B > q\sqrt{n}/2$, linear algebra modulo q allows to solve SIS efficiently – our choice of B is always significantly lower than this threshold.

Finally, the computation of the verification key and signature sizes (in bytes) is performed as in [BDK$^+$20] and [DKL$^+$18], respectively. We note that the

[5] We use the scripts from `https://github.com/pq-crystals/security-estimates`.

updated Dilithium-G has signature sizes $\sim 25\%$ smaller than those of Dilithium. To compute signature sizes for Gaussian and Hypercube versions, we rely on a strategy explained in [ETWY22, Section 5].

	Hyperball-Uniform			Improved Gaussian		
	Medium	Recommended	Very High	Medium	Recommended	Very High
Ring dimension ℓ	256	256	256	256	256	256
q	254977	254977	254977	254977	254977	254977
(m, n)	$(4, 3)$	$(6, 4)$	$(8, 6)$	$(4, 3)$	$(5, 4)$	$(7, 6)$
η	2	3	2	2	3	2
S	91	140	115	91	134	111
τ	39	49	60	39	49	60
$t = S \cdot \sqrt{\tau}$	568	980	890	568	938	860
B	741K	1894K	2330K	836K	413K	760K
γ_2	$\frac{q-1}{16}$	$\frac{q-1}{8}$	$\frac{q-1}{8}$	$\frac{q-1}{64}$	$\frac{q-1}{48}$	$\frac{q-1}{48}$
d	10	13	13	12	11	10
M	4	4	4	4	4	4
BKZ block-size b to break SIS	464 (402)	677 (595)	958 (848)	453 (393)	715 (620)	991 (868)
Best known classical bit-cost	135 (117)	198 (174)	280 (248)	132 (114)	209 (181)	289 (253)
Best known quantum bit-cost	123 (106)	179 (157)	254 (224)	120 (104)	189 (164)	262 (230)
BKZ block-size b to break LWE	403	623	953	403	623	1018
Best known classical bit-cost	117	182	278	117	182	297
Best known quantum bit-cost	106	165	252	106	165	170
Expected signature size	1900	2710	3989	1672	2284	3347
Expected public key size	1056	1184	1824	672	1152	1376

Table 4. Parameters for hyperball-uniform and improved Dilithium-G.

We apply to Dilithium-G two modifications introduced in this work. In Table 4 (right side), we show the improvements we obtain when the standard deviation σ is computed using our refined bound (available in the full version) on the smooth Rényi divergence between two Gaussians and instantiated with $\varepsilon = 2^{-64}$ instead of $\varepsilon = 2^{-128}$, as allowed by the use of Rényi divergence (as discussed in Section 4). Keeping $M = 4$, the standard deviation σ drops from $11t$ to $6.85t$ and leads to an additional saving of $\sim 5\%$ on the signature size. When compared to Dilithium, we obtain up to $\sim 30\%$ of signature size savings.

Finally, we explore the use of the continuous uniform distributions in hyperballs. We take the algorithms from Dilithium-G, which are adapted to radial distributions and replace the Gaussians with the continuous uniform distributions in hyperballs, adding coefficient-wise rounding to integers when computing commitments. We also emphasize that the rejection step is deterministic. To set parameters, the bound B is computed using the radius of the hyperball instead of the probabilistic upper bound on the norm of a Gaussian vector. In Table 4 (left side), we provide the instantiations that we obtained. We note that the signature sizes are larger than those obtained with Gaussians. The growth of the signature size comes from two factors: first, the bound B is larger than the Gaussian case, likely because of constant factors hidden in the Rényi divergence computations of this section. Second, in order to encode a signature, we use a coordinate-wise Huffman coding of the rounded vector, which is less efficient than in the Gaussian case, as the Gaussian distribution minimizes entropy across distributions with a fixed standard deviation. When compared to Dilithium, the signature size still drops by $\sim 10\%$ to $\sim 20\%$, which underlines the trade-off offered by the uniform distributions in hyperballs, between the efficiency of Gaussians and the ease of implementation provided by the uniform distributions in hypercubes.

All figures of Tables 3 and 4 can be reproduced using scripts available at `https://github.com/jdevevey/security-estimates`.

6 Circumventing the Second Lower Bound via Bounded Rejection Sampling

We conclude this work by investigating an alternative way to perform rejection sampling which circumvents our lower bound on the expected number of loop iterations from Section 3.3. Notably, this approach makes the resulting signature scheme run within a given amount of time, which may be required in some practical applications (e.g., in real-time systems).

A first solution could be to set a bound on the maximal number of iterations, based on the run-time analysis from Lemma 2. However, this leads to a very large bound, of the order of $\omega(\log\lambda + \log q_{\mathsf{sig}})/\log(M/(M-1+\varepsilon))$, to ensure that with probability $1-\lambda^{-\omega(1)}$, no signature among q_{sig} requires more iterations.

In the following, we propose a rejection sampling strategy that lets us fix an arbitrary bound $i \geq 1$ on the number of iterations while still guaranteeing an output is produced at the end of the process. This strategy consists in first running $i-1$ iterations of the rejection sampling procedure. If something was output, then we are done, but if all iterations failed, we have to sample something that is related to the target distribution, in one-shot. For this last step, we use some sort of flooding. Note that, setting $i = 1$, one obtains one-shot signatures based on flooding, as in [ASY22]. Hence, this strategy can be seen as a generalization of both rejection sampling and flooding techniques.

6.1 Bounded Rejection Sampling Lemma

Let $i \geq 1$ be an arbitrary bound for the number of loop iterations. Instead of simply having one distribution P_s to sample from, we now use two distributions P_f and P_s, where P_s is used for the rejection sampling part (the first $i-1$ iterations) and P_f is used in case of $i-1$ successive failures. If the divergences $R_\infty(P_f\|P_s)$ and $R_\infty(P_s\|P_t)$ are small, this strategy works. Moreover, the resulting distribution has a divergence with P_s and is a weighted mean of the classical rejection sampling-resulting distribution and the flooding distribution. This is what we prove in the following lemma.

Lemma 9 (Bounded Rejection Sampling). *Let p_f, p_t, p_s be probability densities, either all continuous or all discrete, and $\varepsilon_0, \varepsilon_1 \geq 0, M_0, M_1 \geq 1$ with*

$$R_\infty^{\varepsilon_0}(p_f\|p_t) \leq M_0 \quad \textit{and} \quad R_\infty^{\varepsilon_1}(p_t\|p_s) \leq M_1.$$

Then

$$R_\infty^{\frac{M}{M_0}\varepsilon_0}(\mathcal{B}_i^{\mathsf{real}} \| \mathcal{B}_i^{\mathsf{ideal}}) \leq M,$$

where

$$M = \left(1 - \left(1 - \frac{1}{M_1}\right)^{i-1}\right)\frac{1}{1-\varepsilon_1} + \left(1 - \frac{1+\varepsilon_1}{M_1}\right)^{i-1} \cdot M_0,$$

and $\mathcal{B}_i^{\mathsf{real}}$ *and* $\mathcal{B}_i^{\mathsf{ideal}}$ *are defined in Figure 4.*

Note that in the case where $i = 1$, distribution p_s is useless, as $\mathcal{B}_1^{\mathsf{real}}$ samples $z \hookleftarrow p_f$ and returns it: this is flooding. Our lemma captures this situation, as $M = M_0$ in that case. It is then not only a generalization of rejection sampling but also of flooding techniques.

Algorithms $\mathcal{B}_i^{\mathsf{ideal}}$ and $\mathcal{B}_i^{\mathsf{ideal}'}$ produce the same distribution for variable z, and hence Lemma 9 also holds when replacing $\mathcal{B}_i^{\mathsf{ideal}}$ by $\mathcal{B}_i^{\mathsf{ideal}'}$. Algorithm $\mathcal{B}_i^{\mathsf{ideal}'}$ is more convenient when analyzing the adapted Lyubashevsky signature scheme.

```
Algorithm $\mathcal{B}_i^{\mathsf{real}}$:
$\ell \leftarrow 1$
while $\ell \leq i-1$ do
    $z \hookleftarrow p_s$
    with probability $\min(\frac{p_t(z)}{M_1 \cdot p_s(z)}, 1)$, return $z$
    $\ell \leftarrow \ell + 1$
end while
return $z \hookleftarrow p_f$
```

```
Algorithm $\mathcal{B}_i^{\mathsf{ideal}}$:
return $z \hookleftarrow p_t$
```

```
Algorithm $\mathcal{B}_i^{\mathsf{ideal}'}$:
$\ell \leftarrow 1$
while $\ell \leq i-1$ do
    $z \hookleftarrow p_t$
    with probability $\frac{1}{M_1}$, return $z$
    $\ell \leftarrow \ell + 1$
end while
return $z \hookleftarrow p_t$
```

Fig. 4. Bounded rejection sampling algorithms.

Proof. With p_t and p_s, for $\mathtt{t} \in \{\mathsf{real}, \mathsf{ideal}\}$, we can view $\mathcal{B}_i^{\mathtt{t}}$ as calling $i-1$ times $\mathcal{A}^{\mathtt{t}}$ from Figure 1, returning the value of the first call that does not abort, and if all calls failed, returning some independent sample $z \hookleftarrow p_f$ (or p_t). Using probability bounds from Lemma 1 and letting $\mathcal{A}^{\mathsf{real}}(\perp)$ denote the probability that $\mathcal{A}^{\mathsf{real}}$ aborts, we know that

$$\begin{aligned}
\mathcal{B}_i^{\mathsf{real}}(x) &= \left[\sum_{0 \leq j \leq i-2} (\mathcal{A}^{\mathsf{real}}(\perp))^j \cdot \min\left(\frac{p_t(x)}{M_1}, p_s(x)\right)\right] + (\mathcal{A}^{\mathsf{real}}(\perp))^{i-1} \cdot p_f(x) \\
&= \frac{1 - (\mathcal{A}^{\mathsf{real}}(\perp))^{i-1}}{1 - \mathcal{A}^{\mathsf{real}}(\perp)} \cdot \min\left(\frac{p_t(x)}{M_1}, p_s(x)\right) + (\mathcal{A}^{\mathsf{real}}(\perp))^{i-1} \cdot p_f(x) \\
&\leq \frac{1 - \left(1 - \frac{1}{M_1}\right)^{i-1}}{\frac{1-\varepsilon_1}{M_1}} \cdot \frac{p_t(x)}{M_1} + \left(\frac{M_1 - 1 + \varepsilon_1}{M_1}\right)^{i-1} \cdot p_f(x).
\end{aligned}$$

Let us define

$$M = \left(1 - \left(1 - \frac{1}{M_1}\right)^{i-1}\right) \cdot \frac{1}{1-\varepsilon_1} + \left(\frac{M_1 - 1 + \varepsilon_1}{M_1}\right)^{i-1} \cdot M_0.$$

For this to be an upper bound on $R_\infty^{\frac{M}{M_0}\varepsilon_0}(\mathcal{B}_i^{\mathsf{real}}\|\mathcal{B}_i^{\mathsf{ideal}})$, it suffices that

$$\Pr_{x \hookleftarrow \mathcal{B}_i^{\mathsf{real}}}[\mathcal{B}_i^{\mathsf{real}}(x) > M \cdot p_t(x)] \le \frac{M}{M_0}\varepsilon_0.$$

For any output x such that $\mathcal{B}_i^{\mathsf{real}}(x) > Mp_t(x)$, it holds $p_f(x) > M_0 p_t(x)$ according to the above upper bound on $\mathcal{B}_i^{\mathsf{real}}(x)$. This yields, by definition of M_0:

$$\Pr_{x \hookleftarrow p_f}\left[\mathcal{B}_i^{\mathsf{real}}(x) > M \cdot p_t(x)\right] \le \varepsilon_0.$$

The probability is however not taken over the desired distribution for x. Note that if we combine $p_f(x) > M_0 \cdot p_t(x)$ with the above bound on the distribution of the output of $\mathcal{B}_i^{\mathsf{real}}$, we get

$$\mathcal{B}_i^{\mathsf{real}}(x) < \frac{M}{M_0} \cdot p_f(x).$$

Then $\Pr_{x \hookleftarrow \mathcal{B}_i^{\mathsf{real}}}[\mathcal{B}_i^{\mathsf{real}}(x) > Mp_t(x)] < \frac{M}{M_0}\varepsilon_0$. □

6.2 Lyubashevsky's Signature with Bounded Rejection

In this section, we present a way to modify Lyubashevsky's signature scheme by relying on bounded rejection sampling, as decribed above. This can be seen as a hybrid version between one-shot signatures which use flooding, as in [ASY22], and Lyubashevsky's unbounded signature.

Let $k, n, m, q \ge 1$ specify matrix spaces with $m > n$. Let $\mathcal{M}$ be the message space. Let H be a hash function modeled as a random oracle with domain $\mathbb{Z}_q^n \times \mathcal{M}$ and range some finite set $\mathcal{C} \subseteq \mathbb{Z}^k$. Let $\gamma > 0$. Let $\varepsilon_0, \varepsilon_1 \ge 0, M_0, M_1 \ge 1, i \ge 1$ be parameters related to bounded rejection sampling. Let $\mathcal{S} \subseteq \mathbb{Z}^{m \times k}$. Let P_0, P_1 and P_2 be three probability distributions over $\mathbb{Z}^m$ satisfying

$$\max_{(\mathbf{S},\mathbf{c}) \in \mathcal{S}\times\mathcal{C}} R_\infty^{\varepsilon_0}((P_0)_{+\mathbf{Sc}}\|P_1) \le M_0 \quad \text{and} \quad \max_{(\mathbf{S},\mathbf{c}) \in \mathcal{S}\times\mathcal{C}} R_\infty^{\varepsilon_1}(P_1\|(P_2)_{+\mathbf{Sc}}) \le M_1.$$

Let $(\mathbf{x}_0^\top|\mathbf{x}_1^\top)^\top \hookleftarrow P_0$ and $(\mathbf{y}_0^\top|\mathbf{y}_1^\top)^\top \hookleftarrow P_2$, where $\mathbf{y}_0$ and $\mathbf{x}_0$ take values in $\mathbb{Z}^n$. We present the modified scheme in Figure 5. The key generation algorithm is unchanged from Figure 2.

Before moving to the scheme analysis, let us define

$$M = \left(1 - \left(1 - \frac{1}{M_1}\right)^{i-1}\right)\frac{1}{1-\varepsilon_1} + \left(1 - \frac{1+\varepsilon_1}{M_1}\right)^{i-1} \cdot M_0.$$

The runtime of Sign is deterministically bounded, by at most i loop iterations. The correctness statement from Lemma 3 can be adapted as follows.

Lemma 10 (Correctness). *Let $\varepsilon_0, \varepsilon_1 \ge 0$ and $M_0, M_1 \ge 1$. Let P_0, P_1, P_2 satisfy $\max_{(\mathbf{S},\mathbf{c})\in\mathcal{S}\times\mathcal{C}} R_\infty^{\varepsilon_b}(P_b\|P_{b+1,+\mathbf{Sc}}) \le M_b$ for $b \in \{0,1\}$. Let $(\mathbf{x}_0^\top|\mathbf{x}_1^\top)^\top \hookleftarrow P_0$ and $(\mathbf{y}_0^\top|\mathbf{y}_1^\top)^\top \hookleftarrow P_2$, where $\mathbf{x}_0$ and $\mathbf{y}_0$ take values in $\mathbb{Z}^n$. Assume that $\varepsilon_0 \le \mathsf{negl}(\lambda)$, $M \le \mathsf{poly}(\lambda)$ and $2^{-H_\infty(\mathbf{x}_0|\mathbf{x}_1)_{P_0}}, 2^{-H_\infty(\mathbf{y}_0|\mathbf{y}_1)_{P_2}} \le \mathsf{negl}(\lambda)$. Then, in the ROM, the scheme is correct if $\gamma \ge \gamma_{P_1}$ with γ_{P_1} such that $\Pr_{\mathbf{z}\hookleftarrow P_1}(\|\mathbf{z}\| \ge \gamma_{P_1}) \le \mathsf{negl}(\lambda)$.*

```
$\mathsf{Sign}'(\mu, \mathbf{A}, \mathbf{S})$ :
 1: $\ell \leftarrow 1$
 2: if $\ell \leq i - 1$ then
 3:     $\mathbf{y} \hookleftarrow P_2$
 4: else
 5:     $\mathbf{y} \hookleftarrow P_0$
 6: end if
 7: $\mathbf{c} \leftarrow H(\mathbf{Ay}, \mu)$
 8: $\mathbf{z} \leftarrow \mathbf{y} + \mathbf{Sc}$
 9: $u \hookleftarrow U([0,1])$
10: if $u \leq \frac{P_1(\mathbf{z})}{M_1 P_2(\mathbf{y})}$ or $\ell = i$ then
11:     return $(\mathbf{z}, \mathbf{c})$
12: else
13:     $\ell \leftarrow \ell + 1$
14:     go to Step 2
15: end if
```

```
$\mathsf{Verify}(\mu, \mathbf{z}, \mathbf{c}, \mathbf{A}, \mathbf{T} = \mathbf{AS})$ :
 1: if $\|\mathbf{z}\| \leq \gamma$ and $\mathbf{c} = H(\mathbf{Az} - \mathbf{Tc}, \mu)$
    then
 2:     return 1
 3: else
 4:     return 0
 5: end if
```

Fig. 5. Lyubashevsky's signature scheme with bounded rejection.

The main modification towards analyzing the security of the scheme from Figure 5, compared to the one from Figure 2, resides in the observation that the distributions of the pair $(\mathbf{z}, \mathbf{c})$ obtained by the two processes from Figure 6 have $\frac{M}{M_0}\varepsilon_0$-smooth Rényi divergence of infinite order bounded by M. This is obtained by applying Lemma 9. Note that the hash function H needs to be consistently programmed for every $\mathbf{c}$ that is produced, which is why we use the formalism of Algorithm $\mathcal{B}_i^{\mathsf{ideal}'}$ rather than Algorithm $\mathcal{B}_i^{\mathsf{ideal}}$.

By the multiplicativity of the smooth Rényi divergence (see full version), we obtain that the $(q_{\mathsf{sig}} \cdot M\varepsilon_0/M_0)$-smooth Rényi divergence between the adversary's views in games where the changes from Figure 6 have been applied to all signature queries, is bounded by $M^{q_{\mathsf{sig}}}$. The probability preservation property can then be used meaningfully if $q_{\mathsf{sig}} \cdot M\varepsilon_0/M_0 = 2^{-\Omega(\lambda)}$ and $M^{q_{\mathsf{sig}}} \leq \mathsf{poly}(\lambda)$.

Once the signature queries are simulated without the signing key, the security proof can be completed as in prior works (see [Lyu09,Lyu12,AFLT16]).

Asymptotic trade-off. We now discuss the choices of the distributions P_0, P_1 and P_2. We require that $M^{q_{\mathsf{sig}}} = \mathsf{poly}(\lambda)$ and $q_{\mathsf{sig}} \cdot M\varepsilon_0/M_0 = 2^{-\Omega(\lambda)}$, with $\varepsilon_0, \varepsilon_1$, M_0, M_1 and M as in Lemma 9. We are aiming at not too large divergence bounds M_0, M_1, M as signatures typically become less efficient when they increase. For this reason, we set $\varepsilon_0 = 2^{-\Omega(\lambda)}$. As the condition $M^{q_{\mathsf{sig}}} = \mathsf{poly}(\lambda)$ forces M to be close to 1, the condition $q_{\mathsf{sig}} \cdot M\varepsilon_0/M_0 = 2^{-\Omega(\lambda)}$ is already satisfied. We now focus on ε_1, M_0 and M_1.

When i tends to infinity, we have $M \approx 1/(1-\varepsilon_1)$, so that we can set $\varepsilon_1 \approx 1/q_{\mathsf{sig}}$ as in Section 4. For $i = 1$, we have $M = M_0$, and we fall in the regime of [ASY22, Section 4]. Let us now consider the small i case, which is probably the most interesting one for applications requiring a bounded signature time. As $M \geq 1/(1-\varepsilon_1)$ and we must ensure that $M^{q_{\mathsf{sig}}} = \mathsf{poly}(\lambda)$, we set ε_1 at

```
1: $\ell \leftarrow 1$
2: if $\ell \leq i-1$ then
3:     $\mathbf{y} \hookleftarrow P_2$
4: else
5:     $\mathbf{y} \hookleftarrow P_0$
6: end if
7: $\mathbf{c} \leftarrow U(\mathcal{C})$
8: $\mathbf{z} \leftarrow \mathbf{y} + \mathbf{Sc}$
9: $u \hookleftarrow U([0,1])$
10: if $u \leq \frac{P_1(\mathbf{z})}{M_1 P_2(\mathbf{y})}$ or $\ell = i$ then
11:     return $(\mathbf{z}, \mathbf{c})$
12: else
13:     $\ell \leftarrow \ell + 1$
14:     go to Step 2
15: end if
```

```
1: $\ell \leftarrow 1$
2: $\mathbf{y} \hookleftarrow P_1$
3: $\mathbf{c} \leftarrow U(\mathcal{C})$
4: $\mathbf{z} \leftarrow \mathbf{y} + \mathbf{Sc}$
5: $u \hookleftarrow U([0,1])$
6: if $u \leq \frac{1}{M_1}$ or $\ell = i$ then
7:     return $(\mathbf{z}, \mathbf{c})$
8: else
9:     $\ell \leftarrow \ell + 1$
10:     go to Step 2
11: end if
```

Fig. 6. Simulating signatures.

most of the order of $1/q_{\mathsf{sig}}$. This implies that $M \approx 1 + M_0 \cdot (1 - 1/M_1)^{i-1}$, and hence we set $(M_0 - 1) \cdot (1 - 1/M_1)^{i-1} = O(1/q_{\mathsf{sig}})$. For Gaussian and hyperball-uniform instanciations, this leads to a standard deviation (resp. radius) growing polynomially in $q_{\mathsf{sig}}/(1 - 1/M_1)^{i-1}$.

We argue now that the trade-off above (for small i) seems essentially optimal. For $i = 1$, it was showed in [ASY22, Appendix C.2] that the folklore statistical attack against the Gaussian and rejection-free version of Lyubashevsky's signature scheme runs in subexponential time when $M_0 = q_{\mathsf{sig}}^{o(1)}$. Now, for larger i and sufficiently distinct target and flooding distributions, an adversary could consider the signatures for which all loop iterations failed (i.e., the output sample corresponds to the flooding distribution), and run the statistical attack described in [ASY22] for those samples. As the probability of rejecting all samples is essentially $(1 - 1/M_1)^{i-1}$, this attack matches with the trade-off above.

Acknowledgments. The authors thank Wonhee Cho for helpful discussions. The authors were supported by the AMIRAL ANR grant (ANR-21-ASTR-0016), the European Union Horizon 2020 Research and Innovation Program Grant 780701, the PEPR quantique France 2030 programme (ANR-22-PETQ-0008) and the PEPR Cyber France 2030 programme (ANR-22-PECY-0003).

References

AAB$^+$19. S. Akleylek, E. Alkim, P. S. L. M. Barreto, N. Bindel, J. Buchmann, E. Eaton, G. Gutoski, J. Krämer, P. Longa, H. Polat, J. E. Ricardini, and G. Zanon. qTESLA round-3 candidate to the NIST post-quantum cryptography standardisation project, 2019. Available at `https://qtesla.org/`.

ABB^{+}17. E. Alkim, N. Bindel, J. Buchmann, Ö. Dagdelen, E. Eaton, G. Gutoski, J. Krämer, and F. Pawlega. Revisiting TESLA in the quantum random oracle model. In *PQCrypto*, 2017.

AFLT16. M. Abdalla, P.-A. Fouque, V. Lyubashevsky, and M. Tibouchi. Tightly secure signatures from lossy identification schemes. *J. Cryptol.*, 2016.

ASY22. S. Agrawal, D. Stehlé, and A. Yadav. Round-optimal lattice-based threshold signatures, revisited. In *ICALP*, 2022.

BDK^{+}20. S. Bai, L. Ducas, E. Kiltz, T. Lepoint, V. Lyubashevsky, P. Schwabe, G. Seiler, and D. Stehlé. CRYSTALS-DILITHIUM round-3 candidate to the NIST post-quantum cryptography standardisation project, 2020. Avalailable at `https://pq-crystals.org/dilithium/`.

BG14. S. Bai and S. D. Galbraith. An improved compression technique for signatures based on learning with errors. In *CT-RSA*, 2014.

BLR^{+}18. S. Bai, T. Lepoint, A. Roux-Langlois, A. Sakzad, D. Stehlé, and R. Steinfeld. Improved security proofs in lattice-based cryptography: Using the Rényi divergence rather than the statistical distance. *J. Cryptol.*, 2018.

Dat09. N. Datta. Min-and max-relative entropies and a new entanglement monotone. *T. Inform. Theory*, 2009.

DDLL13. L. Ducas, A. Durmus, T. Lepoint, and V. Lyubashevsky. Lattice signatures and bimodal gaussians. In *CRYPTO*, 2013.

Dev86. L. Devroye. *Non-Uniform random variate generation.* 1986.

DKL^{+}18. L. Ducas, E. Kiltz, T. Lepoint, V. Lyubashevsky, P. Schwabe, G. Seiler, and D. Stehlé. CRYSTALS-DILITHIUM: A lattice-based digital signature scheme. *TCHES*, 2018.

DPSZ12. I. Damgård, V. Pastro, N. P. Smart, and S. Zakarias. Multiparty computation from somewhat homomorphic encryption. In *CRYPTO*, 2012.

ETWY22. T. Espitau, M. Tibouchi, A. Wallet, and Y. Yu. Shorter hash-and-sign lattice-based signatures. In *CRYPTO*, 2022.

FS86. A. Fiat and A. Shamir. How to prove yourself: Practical solutions to identification and signature problems. In *CRYPTO*, 1986.

GLP15. T. Güneysu, V. Lyubashevsky, and T. Pöppelmann. Lattice-based signatures: Optimization and implementation on reconfigurable hardware. *T. Comput.*, 2015.

HJMR07. P. Harsha, R. Jain, D. McAllester, and J. Radhakrishnan. The communication complexity of correlation. In *CCC*, 2007.

LNP22. V. Lyubashevsky, N. K. Nguyen, and M. Plançon. Lattice-based zero-knowledge proofs and applications: Shorter, simpler, and more general. In *CRYPTO*, 2022.

Lyu09. V. Lyubashevsky. Fiat-Shamir with aborts: Applications to lattice and factoring-based signatures. In *ASIACRYPT*, 2009.

Lyu12. V. Lyubashevsky. Lattice signatures without trapdoors. In *EUROCRYPT*, 2012.

Lyu16. V. Lyubashevsky. Digital signatures based on the hardness of ideal lattice problems in all rings. In *ASIACRYPT*, 2016.

Pre17. T. Prest. Sharper bounds in lattice-based cryptography using the Rényi divergence. In *ASIACRYPT*, 2017.

Ren05. R. Renner. *Security of quantum key distribution.* PhD thesis, ETH Zurich, 2005.

Sch91. C.-P. Schnorr. Efficient signature generation by smart cards. *J. Cryptol.*, 1991.