

Isochronous Gaussian Sampling: From Inception to Implementation

With Applications to the Falcon Signature Scheme

James Howe¹, Thomas Prest¹, Thomas Ricosset², and Mélissa Rossi^{2,3,4,5}

¹ PQShield

`james.howe@pqshield.com` `thomas.prest@pqshield.com`

² Thales

`thomas.ricosset@thalesgroup.com`

³ ANSSI

⁴ École normale supérieure, CNRS, PSL University

⁵ Inria

`melissa.rossi@ens.fr`

Abstract Gaussian sampling over the integers is a crucial tool in lattice-based cryptography, but has proven over the recent years to be surprisingly challenging to perform in a generic, efficient and provably secure manner. In this work, we present a modular framework for generating discrete Gaussians with arbitrary center and standard deviation. Our framework is extremely simple, and it is precisely this simplicity that allowed us to make it easy to implement, provably secure, portable, efficient, and provably resistant against timing attacks. Our sampler is a good candidate for any trapdoor sampling and it is actually the one that has been recently implemented in the Falcon signature scheme. Our second contribution aims at systematizing the detection of implementation errors in Gaussian samplers. We provide a statistical testing suite for discrete Gaussians called SAGA (Statistically Acceptable GAussian). In a nutshell, our two contributions take a step towards trustable and robust Gaussian sampling real-world implementations.

Keywords.

Lattice based cryptography, Gaussian Sampling, Isochrony, Statistical verification tools

1 Introduction

Gaussian sampling over the integers is a central building block of lattice-based cryptography, in theory as well as in practice. It is also notoriously difficult to perform efficiently and securely, as illustrated by numerous side-channel attacks exploiting the Gaussian sampler in BLISS [7,20,49,56]. For this reason, some schemes choose to limit or discard the use of Gaussians [5,35]. However, in some situations, Gaussians are unavoidable. The most prominent example is trapdoor sampling [26,48,40]: performing trapdoor sampling using other distributions is an open question, except in one limited cases [36] which entail a growth $O(\sqrt{n})$ to $O(n)$ of the output, resulting in a dwindling security level. Given the countless applications of trapdoor sampling (full-domain hash signatures [26,53], identity-based encryption (or IBE) [26,17], hierarchical IBE [9,1], etc.), it is important to come up with Gaussian samplers over the integers which are not only efficient, but also provably secure, resistant to timing attacks, and more generally easy to deploy.

Our first contribution is to propose a Gaussian sampler over the integers with all the properties which are expected of a sampler for widespread deployment. It is simple and modular, making analysis and subsequent improvements easy. It is efficient and portable, making it amenable to a variety of scenarii. Finally, we formally prove its security and resistance against timing attacks. We detail below different aspects of our sampler:

- **Simplicity and Modularity.** At a high level, our framework only requires two ingredients (a base sampler and a rejection sampler) and combines them in a simple and black-box way. Not only does it make the description of our sampler modular (as one can replace either or both of the ingredients depending on the context), this simplicity and modularity also infuses all aspects of its analysis.
- **Genericity.** Our sampler is fully generic as it works with arbitrary center μ and standard deviation σ . In addition, it does not incur hidden precomputation costs: given a fixed base sampler of parameter σ_{\max} , our framework allows to sample from $D_{\mathbb{Z},\sigma,\mu}$ for any $\eta_\epsilon(\mathbb{Z}^n) \leq \sigma \leq \sigma_{\max}$. In comparison, [42] implicitly requires a different base sampler for each different value of σ ; this limits its applicability for use cases such as Falcon [53], which has up to 2048 different σ 's, all computed at key generation.
- **Efficiency and Portability.** Our sampler is instantiated with competitive parameters which make it very efficient in time and memory usage. For $\sigma_{\max} = 1.8205$ and SHAKE256 used as PRNG, our sampler uses only 512 bytes of memory and achieved 1,848,428 samples per second on an Intel i7-6500U clocked at 2.5 GHz. Moreover, our sampler can be instantiated in such way that it uses only integer operations which make it highly portable.
- **Provable Security.** A security analysis based on the statistical distance would either provide very weak security guarantees or require to increase the running time by an order of magnitude. We instead rely on the Rényi divergence, a tool which in the recent years has allowed dramatic practical gains for lattice-based schemes [3,52]. We carefully selected our parameters as to make them as amenable to a Rényi divergence-based analysis, while retaining as much efficiency as possible.
- **Isochrony.** We formally show that our sampler is isochronous: its running time is independent of the inputs σ, μ and of the output z . Isochrony is weaker than being constant-time, but it nevertheless suffices to argue security against timing attacks. Interestingly, our proof of isochrony relies on techniques and notions which are common in lattice-based cryptography: the smoothing parameter, the Rényi divergence, etc. In particular, the isochrony of our sampler is implied by parameters dictated by the current state of the art for *black-box* security of lattice-based schemes.

One second contribution stems from a simple observation: implementations of otherwise perfectly secure schemes have failed in spectacular ways by introducing weaknesses, a common one being randomness failure: this is epitomized by nonce reuses in ECDSA, leading to jailbreaking Sony PS3 consoles¹ and exposing Bitcoin wallets [6]. The post-quantum community is aware of this point of failure but does not seem to have converged on a systematic way to mitigate it [46]. Randomness failures have been manually discovered and fixed in implementations of Dilithium [45], Falcon [51,47] and other schemes; the case of Falcon is particularly relevant to us because the sampler implemented was the one described in this document!

Our second contribution is a first step at systematically detecting such failures: we propose a statistical test suite called SAGA for validating discrete Gaussians. This test suite can check univariate samples; we therefore use it to validate our own implementation of our sampler. In addition, our test suite can check multivariate Gaussians as well; this enables validation at a higher level: if the base sampler over the integers is validated, but the output of the high-level scheme does not behave like a multivariate Gaussian even though the theory predicts it should, then this is indicative of an implementation mistake somewhere else in the implementation (or, at the worst case, that the theory is deficient). We illustrate that with a simple example of a (purportedly) deficient implementation of Falcon [53], however it can be used for any other

¹ https://media.ccc.de/v/27c3-4087-en-console_hacking_2010.

scheme using sampling multivariate discrete Gaussians, including but not limited to [40,17,25,10]. The test suite is publicly available at: <https://github.com/PQShield/SAGA>.

2 Related Works

In the recent years, there has been a surge of works related to Gaussian sampling over the integers. Building on convolution techniques from [50], [42] proposed an arbitrary-center Gaussian sampler base, as well as a statistical tool (the max-log distance) to analyse it. [3,52,39] revisited classical techniques with the Rényi divergence. Polynomial-based methods were further studied by [52,59,4]. The use of rounded Gaussians was proposed in [30]. Knuth-Yao's DDG trees have been considered in [19,31]. Lazy floating-point precision was studied in [18,15]. We note that techniques dating back to von Neumann [57] allow to generate (continuous) Gaussians elegantly using finite automata [23,2,32]. While these have been considered in the context of lattice-based cryptography [16,14] they are also notoriously hard to make isochronous. Finally, [58] studied previously cited techniques with the goal of minimizing their relative error.

3 Preliminaries

3.1 Gaussians

For $\sigma, \mu \in \mathbb{R}$ with $\sigma > 0$, we call Gaussian function of parameters σ, μ and denote by $\rho_{\sigma, \mu}$ the function defined over \mathbb{R} as $\rho_{\sigma, \mu}(x) = \exp\left(-\frac{(x-\mu)^2}{2\sigma^2}\right)$. Note that when $\mu = 0$ we omit it in index notation, e.g. $\rho_{\sigma}(x) = \rho_{\sigma, 0}(x)$. The parameter σ (resp. μ) is often called the standard deviation (resp. center) of the Gaussian. In addition, for any countable set $S \subseteq \mathbb{R}$ we abusively denote by $\rho_{\sigma, \mu}(S)$ the sum $\sum_{z \in S} \rho_{\sigma, \mu}(z)$. When $\sum_{z \in S} \rho_{\sigma, \mu}(z)$ is finite, we denote by $D_{S, \sigma, \mu}$ and call Gaussian distribution of parameters σ, μ the distribution over S defined by $D_{S, \sigma, \mu}(z) = \rho_{\sigma, \mu}(z) / \rho_{\sigma, \mu}(S)$. Here too, when $\mu = 0$ we omit it in index notation, e.g. $D_{S, \sigma, \mu}(z) = D_{S, \sigma}(z)$. We use the notation \mathcal{B}_p to denote the Bernoulli distribution of parameter p .

3.2 Rényi Divergence

We recall the definition of the Rényi divergence, which we will use massively in our security proofs.

Definition 1 (Rényi Divergence). *Let \mathcal{P}, \mathcal{Q} be two distributions such that $\text{Supp}(\mathcal{P}) \subseteq \text{Supp}(\mathcal{Q})$. For $a \in (1, +\infty)$, we define the Rényi divergence of order a by*

$$R_a(\mathcal{P}, \mathcal{Q}) = \left(\sum_{x \in \text{Supp}(\mathcal{P})} \frac{\mathcal{P}(x)^a}{\mathcal{Q}(x)^{a-1}} \right)^{\frac{1}{a-1}}.$$

In addition, we define the Rényi divergence of order $+\infty$ by

$$R_{\infty}(\mathcal{P}, \mathcal{Q}) = \max_{x \in \text{Supp}(\mathcal{P})} \frac{\mathcal{P}(x)}{\mathcal{Q}(x)}.$$

The Rényi divergence is not a distance; for example, it is neither symmetric nor does it verify the triangle inequality, which makes it less convenient than the statistical distance. On the other hand, it does verify cryptographically useful properties, including a few listed below.

Lemma 2 ([3]). *For two distributions \mathcal{P}, \mathcal{Q} and two families of distributions $(\mathcal{P}_i)_i, (\mathcal{Q}_i)_i$, the Rényi divergence verifies these properties:*

- **Data processing inequality.** For any function f , $R_a(f(\mathcal{P}), f(\mathcal{Q})) \leq R_a(\mathcal{P}, \mathcal{Q})$.
- **Multiplicativity.** $R_a(\prod_i \mathcal{P}_i, \prod_i \mathcal{Q}_i) = \prod_i R_a(\mathcal{P}_i, \mathcal{Q}_i)$.
- **Probability preservation.** For any event $E \subseteq \text{Supp}(\mathcal{Q})$ and $a \in (1, +\infty)$,

$$\mathcal{Q}(E) \geq \mathcal{P}(E)^{\frac{a}{a-1}} / R_a(\mathcal{P}, \mathcal{Q}), \quad (1)$$

$$\mathcal{Q}(E) \geq \mathcal{P}(E) / R_\infty(\mathcal{P}, \mathcal{Q}). \quad (2)$$

The following lemma shows that a bound of δ on the relative error between two distributions implies a bound $O(\delta^2)$ on the log of the Rényi divergence (as opposed to a bound $O(\delta)$ on the statistical distance).

Lemma 3 (Lemma 3 of [52]). *Let \mathcal{P}, \mathcal{Q} be two distributions of same support Ω . Suppose that the relative error between \mathcal{P} and \mathcal{Q} is bounded: $\exists \delta > 0$ such that $|\frac{\mathcal{P}}{\mathcal{Q}} - 1| \leq \delta$ over Ω . Then, for $a \in (1, +\infty)$:*

$$R_a(\mathcal{P}, \mathcal{Q}) \leq \left(1 + \frac{a(a-1)\delta^2}{2(1-\delta)^{a+1}}\right)^{\frac{1}{a-1}} \underset{\delta \rightarrow 0}{\sim} 1 + \frac{a\delta^2}{2}$$

3.3 Smoothing Parameter

For $\epsilon > 0$, the smoothing parameter $\eta_\epsilon(\Lambda)$ of a lattice Λ is the smallest value $\sigma > 0$ such that $\rho_{\frac{1}{\sigma\sqrt{2\pi}}}(\Lambda^* \setminus \{\mathbf{0}\}) \leq \epsilon$, where Λ^* denotes the dual of Λ . In the literature, some definitions of the smoothing parameter scale our definition by a factor $\sqrt{2\pi}$. It is shown in [41] that $\eta_\epsilon(\mathbb{Z}^n) \leq \eta_\epsilon^+(\mathbb{Z}^n)$, where:

$$\eta_\epsilon^+(\mathbb{Z}^n) = \frac{1}{\pi} \sqrt{\frac{1}{2} \log \left(2n \left(1 + \frac{1}{\epsilon} \right) \right)}. \quad (3)$$

3.4 Isochronous algorithms

We now give a semi-formal definition of isochronous algorithms.

Definition 4. *Let \mathcal{A} be a (probabilistic or deterministic) algorithm with set of input variables \mathcal{I} , set of output variables \mathcal{O} , and let $\mathcal{S} \subseteq \mathcal{I} \cup \mathcal{O}$ be the set of sensitive variables. We say that \mathcal{A} is perfectly isochronous with respect to \mathcal{S} if its running time is independent of any variable in \mathcal{S} .*

In addition, we say that \mathcal{A} statistically isochronous with respect to \mathcal{S} if there exists a distribution \mathcal{D} independent of all the variables in \mathcal{S} , such that the running time of \mathcal{A} is statistically close (for a clearly identified divergence) to \mathcal{D} .

We note that we can define a notion of computationally isochronous algorithm. For such an algorithm, it is computationally hard to recover the sensitive variables even given the distribution of the running time of the algorithm. We can even come up with a contrived example of such an algorithm: let $\mathcal{A}()$ select in isochronous an x uniformly in a space of min-entropy $\geq \lambda$, compute $y = H(x)$ and wait a time y before outputting x . One can show that recovering x given the running time of \mathcal{A} is hard if H is a one-way function.

4 The sampler

In this section, we describe our new sampler with arbitrary standard deviation and center. The main assumption of our setting is to consider that all the standard deviations are bounded and that the center is in $[0, 1]$. In other words, denoting the upper bound and lower bound on the

Algorithm 1 SamplerZ(σ, μ)

Require: $\mu \in [0, 1], \sigma \leq \sigma_{\max}$
Ensure: $z \sim D_{\mathbb{Z}, \sigma, \mu}$
 1: **while** True **do**
 2: $z_0 \leftarrow \text{BaseSampler}()$
 3: $b \leftarrow \{0, 1\}$ uniformly
 4: $z := (2b - 1) \cdot z_0 + b$
 5: $x := \frac{z_0^2}{2\sigma_{\max}^2} - \frac{(z-\mu)^2}{2\sigma^2}$
 6: **if** AcceptSample(σ, x) **then**
 7: **return** z

Algorithm 2 AcceptSample(σ, x)

Require: $\sigma_{\min} \leq \sigma \leq \sigma_{\max}, x < 0$
Ensure: $b \sim \mathcal{B}_{\frac{\sigma_{\min}}{\sigma} \cdot \text{ApproxExp}(x)}$
 1: $p := \frac{\sigma_{\min}}{\sigma} \cdot \text{ApproxExp}(x)$
Lazy Bernoulli sampling
 2: $i := 1$
 3: **do**
 4: $i := i \cdot 2^8$
 5: $u \leftarrow \llbracket 0, 2^8 - 1 \rrbracket$ uniformly
 6: $v := \lfloor p \cdot i \rfloor$ & 0xff
 7: **while** $u = v$
 8: **return** ($u < v$)

standard deviation as $\sigma_{\max} > \sigma_{\min} > 0$, we present an algorithm that samples the distribution $D_{\mathbb{Z}, \sigma, \mu}$ for any $\mu \in [0, 1]$ and $\sigma_{\min} \leq \sigma \leq \sigma_{\max}$.

Our sampling algorithm is called **SamplerZ** and is described in Algorithm 1. We denote by **BaseSampler** an algorithm that samples an element with the fixed half Gaussian distribution $D_{\mathbb{Z}^+, \sigma_{\max}}$. The first step consists in using **BaseSampler**. The obtained z_0 sample is then transformed into $z := (2b - 1) \cdot z_0 + b$ where b is a bit drawn uniformly in $\{0, 1\}$. Let us denote by $BG_{\sigma_{\max}}$ the distribution of z . The distribution of z is a discrete bimodal half-Gaussian of centers 0 and 1. More formally,

$$BG_{\sigma_{\max}}(z) = \frac{1}{2} \begin{cases} D_{\mathbb{Z}^+, \sigma_{\max}}(-z) & \text{if } z \leq 0 \\ D_{\mathbb{Z}^+, \sigma_{\max}}(z - 1) & \text{if } z \geq 1. \end{cases} \quad (4)$$

The main interest of getting the latter distribution is that $BG_{\sigma_{\max}}(z) \geq D_{\mathbb{Z}, \sigma, \mu}(z)$ for any $z \in \mathbb{Z}$, $\mu \in [0, 1]$ and $\sigma \leq \sigma_{\max}$. Then, to recover the desired distribution $D_{\mathbb{Z}, \sigma, \mu}$ for the inputs (σ, μ) , one might want to apply the classical rejection sampling technique applied to lattice based schemes [34] and accept z with probability

$$\begin{aligned} \frac{D_{\mathbb{Z}, \sigma, \mu}(z)}{BG_{\sigma_{\max}}(z)} &= \begin{cases} \exp\left(\frac{z^2}{2\sigma_{\max}^2} - \frac{(z-\mu)^2}{2\sigma^2}\right) & \text{if } z \leq 0 \\ \exp\left(\frac{(z-1)^2}{2\sigma_{\max}^2} - \frac{(z-\mu)^2}{2\sigma^2}\right) & \text{if } z \geq 1 \end{cases} \\ &= \exp\left(\frac{z_0^2}{2\sigma_{\max}^2} - \frac{(z-\mu)^2}{2\sigma^2}\right). \end{aligned}$$

The element inside the exp is computed in step 5. Next, we also introduce an algorithm denoted **AcceptSample**. The latter performs the rejection sampling (Algorithm 2): using **ApproxExp** an algorithm that returns $\exp(\cdot)$, it returns a Bernoulli sample with the according probability. Actually, for isochrony matters, detailed in Section 6, the latter acceptance probability is rescaled by a factor $\frac{\sigma_{\min}}{\sigma}$. As z follows the $BG_{\sigma_{\max}}$ distribution, after the rejection sampling, the final distribution of **SamplerZ**(σ, μ) is then proportional to $\frac{\sigma_{\min}}{\sigma} \cdot D_{\mathbb{Z}, \sigma, \mu}$ which is, after normalization exactly equal to $D_{\mathbb{Z}, \sigma, \mu}$. Thus, with this construction, one can derive the following proposition.

Proposition 5 (Correctness). *Assume that all the uniform distributions are perfect and that $\text{BaseSampler} = D_{\mathbb{Z}^+, \sigma_{\max}}$ and $\text{ApproxExp} = \exp$, then the construction of **SamplerZ** (in Algorithms 1 and 2) is such that $\text{SamplerZ}(\sigma, \mu) = D_{\mathbb{Z}, \sigma, \mu}$.*

In practical implementations, one cannot achieve perfect distributions. Only achieving $\text{BaseSampler} \approx D_{\mathbb{Z}^+, \sigma_{\max}}$ and $\text{ApproxExp} \approx \exp$ is possible. Section 6 proves that, under certain conditions on **BaseSampler** and **ApproxExp** and on the number of sampling queries, the final distribution remains indistinguishable from $D_{\mathbb{Z}, \sigma, \mu}$.

Table 1: Number of calls to `SamplerZ`, `BaseSampler` and `ApproxExp`

	Notation	Value for Falcon
Calls to sign (as per NIST)	Q_s	$\leq 2^{64}$
Calls to <code>SamplerZ</code>	Q_{samplerZ}	$Q_s \cdot 2 \cdot n \leq 2^{75}$
Calls to <code>BaseSampler</code>	Q_{bs}	$\mathcal{N}_{\text{iter}} \cdot Q_{\text{samplerZ}} \leq 2^{76}$
Calls to <code>ApproxExp</code>	Q_{exp}	$Q_{\text{bs}} \leq 2^{76}$

5 Proof of Security

Table 1 gives the notations for the number of calls to `SamplerZ`, `BaseSampler` and `ApproxExp` and the considered values when the sampler is instantiated for Falcon. Due to the rejection sampling in step 6, there will be a (potentially infinite) number of iterations of the **while** loop. We will show later in Lemma 7, that the number of iterations follows a geometric law of parameter $\approx \frac{\sigma_{\min} \cdot \sqrt{2\pi}}{2 \cdot \rho_{\sigma_{\max}}(\mathbb{Z}^+)}$. We note $\mathcal{N}_{\text{iter}}$ a heuristic considered maximum number of iterations. By a central limit argument, $\mathcal{N}_{\text{iter}}$ will only be marginally higher than the expected number of iterations. To instantiate the values $Q_{\text{exp}} = Q_{\text{bs}} = \mathcal{N}_{\text{iter}} \cdot Q_{\text{samplerZ}}$ for the example of Falcon, we take $\mathcal{N}_{\text{iter}} = 2$. In fact, $\frac{\sigma_{\min} \cdot \sqrt{2\pi}}{2 \cdot \rho_{\sigma_{\max}}(\mathbb{Z}^+)} \leq 2$ for Falcon's parameters.

The following Theorem estimates the security of `SamplerZ`, it is independant of the chosen values for the number of calls.

Theorem 6 (Security of `SamplerZ`). *Let λ_{IDEAL} (resp. λ_{REAL}) be the security parameter of an implementation using the perfect distribution $D_{\mathbb{Z}, \sigma, \mu}$ (resp. the real distribution `SamplerZ`). If both following conditions are respected, at most two bits of security are lost. In other words, $\Delta\lambda := \lambda_{\text{IDEAL}} - \lambda_{\text{REAL}} \leq 2$.*

$$\forall x < 0, \left| \frac{\text{ApproxExp}(x) - \exp(x)}{\exp(x)} \right| \leq \sqrt{\frac{2 \cdot \lambda_{\text{REAL}}}{2 \cdot (2 \cdot \lambda_{\text{REAL}} + 1)^2 \cdot Q_{\text{exp}}}} \quad (\text{Cond. (1)})$$

$$R_{2 \cdot \lambda_{\text{REAL}} + 1}(\text{BaseSampler}, D_{\mathbb{Z}^+, \sigma_{\max}}) \leq 1 + \frac{1}{4 \cdot Q_{\text{bs}}} \quad (\text{Cond. (2)})$$

The proof of this Theorem is given in Appendix A.

To get concrete numerical values, we assume that 256 bits are claimed on the original scheme, thus 254 bits of security are claimed for the real implementation. Then for an implementation of Falcon, the numerical values are

$$\sqrt{\frac{2 \cdot \lambda_{\text{REAL}}}{2 \cdot (2 \cdot \lambda_{\text{REAL}} + 1)^2 \cdot Q_{\text{exp}}}} \approx 2^{-43} \quad \text{and} \quad \frac{1}{4 \cdot Q_{\text{bs}}} \approx 2^{-78}.$$

5.1 Instantiating the `ApproxExp`

To achieve condition (1) with `ApproxExp`, we use a polynomial approximation of the exponential on $[-\ln(2), 0]$. In fact, one can reduce the parameter x modulo $\ln(2)$ such that $x = -r - s \ln(2)$. Compute the exponential remains to compute $\exp(x) = 2^{-s} \exp(-r)$. Noting that $s \geq 64$ happen very rarely, thus s can be saturated at 63 to avoid overflow without loss in precision.

We use the polynomial approximation tool provided in GALACTICS [4]. This tool generates polynomial approximations that allow a computation in fixed precision with chosen size of coefficients and degree. As an example, for 32-bit coefficients and a degree 10, we obtain a polynomial $P_{\text{gal}}(x) := \sum_{i=0}^{10} a_i \cdot x^i$, with:

- $a_0 = 1$;
- $a_1 = 1$;
- $a_2 = 2^{-1}$;
- $a_3 = 2863311530 \cdot 2^{-34}$;
- $a_4 = 2863311481 \cdot 2^{-36}$;
- $a_5 = 2290647631 \cdot 2^{-38}$;
- $a_6 = 3054141714 \cdot 2^{-41}$;
- $a_7 = 3489252544 \cdot 2^{-44}$;
- $a_8 = 3473028713 \cdot 2^{-47}$;
- $a_9 = 2952269371 \cdot 2^{-50}$;
- $a_{10} = 3466184740 \cdot 2^{-54}$.

For any $x \in [-\ln(2), 0]$, P_{gal} verifies $\left| \frac{P_{\text{gal}}(x) - \exp(x)}{\exp(x)} \right| \leq 2^{-47}$, which is sufficient to verify condition (1) for Falcon implementation.

Flexibility on the implementation of the polynomial Depending on the platform and the requirement for the signature, one can adapt the polynomial to fit their constraints. For example, if one wants to minimize the number of multiplications, implementing the polynomial with Horner’s form is the best option. The polynomial is written in the following form :

$$P_{\text{gal}}(x) = a_0 + x \cdot (a_1 + x \cdot (a_2 + x \cdot (a_3 + x \cdot (a_4 + x \cdot (a_5 + x \cdot (a_6 + x \cdot (a_7 + x \cdot (a_8 + x \cdot (a_9 + x \cdot a_{10}))))))))).$$

Evaluating P_{gal} is then done serially as follows:

$$\begin{aligned} y &\leftarrow a_{10} \\ y &\leftarrow a_9 + y \times x \\ &\vdots \\ y &\leftarrow a_1 + y \times x \\ y &\leftarrow a_0 + y \times x \end{aligned}$$

Some architectures with small register sizes may be faster if the size of the coefficients of the polynomial is minimized, thus GALACTICS tool can be used to generate a polynomial with smaller coefficients. For example, we propose an alternative polynomial approximation on $[0, \frac{\ln(2)}{64}]$ with 25 bits coefficients.

$$P = 1 + x + 2^{-1}x^2 + 699051 \cdot 2^{-22} \cdot x^3 + 699299 \cdot 2^{-24} \cdot x^4 + 605552 \cdot 2^{-26} \cdot x^5$$

To recover the polynomial approximation on $[0, \ln(2)]$, we compute $P(\frac{x}{64})^{64}$.

Some architectures enjoy some level of parallelism, in which case it is desirable to minimise the depth of the circuit computing the polynomial². Writing P_{gal} in Estrin’s form [21] is helpful in this regard.

$$\begin{aligned} \mathbf{x}_2 &\leftarrow \mathbf{x} \times \mathbf{x} \\ \mathbf{x}_4 &\leftarrow \mathbf{x}_2 \times \mathbf{x}_2 \\ P_{\text{gal}}(\mathbf{x}) &\leftarrow (\mathbf{x}_4 \times \mathbf{x}_4) \times ((\mathbf{a}_8 + \mathbf{a}_9 \times \mathbf{x}) + \mathbf{x}_2 \times \mathbf{a}_{10}) \\ &\quad + (((\mathbf{a}_0 + \mathbf{a}_1 \times \mathbf{x}) + \mathbf{x}_2 \times (\mathbf{a}_2 + \mathbf{a}_3 \times \mathbf{x})) + \mathbf{x}_4 \times ((\mathbf{a}_4 + \mathbf{a}_5 \times \mathbf{x}) + \mathbf{x}_2 \times (\mathbf{a}_6 + \mathbf{a}_7 \times \mathbf{x}))) \end{aligned}$$

5.2 Instantiating the BaseSampler

To achieve condition (2) with BaseSampler, we rely on a cumulative distribution table (CDT). We precompute a table of the cumulative distribution function of $D_{\mathbb{Z}^+, \sigma_{\max}}$ with a certain precision; then, to produce a sample, we generate a random value in $[0, 1]$ with the same precision, and return the index of the last entry in the table that is greater than that value. In variable time, this can be done relatively efficiently with a binary search, but a constant-time implementation has essentially no choice but to read the entire table each time and carry out each comparison. This process is summed up in Algorithm 3. The parameters w and θ are respectively the number

Algorithm 3 SampleCDT: full-table access CDT

```

 $z \leftarrow 0$ 
 $u \leftarrow [0, 1)$  uniformly with  $\theta$  bits of absolute precision
for  $0 \leq i \leq w$  do
     $b \leftarrow (\text{CDT}[w] \geq u)$   $\triangleright b = 1$  if it is true and 0 otherwise
     $z \leftarrow z + b$ 
return  $z$ 

```

of elements of the CDT and the precision of its coefficients. Let $a = 2 \cdot \lambda_{\text{REAL}} + 1$. To derive the parameters w and θ we use a simple script that, given σ_{\max} and θ as inputs:

1. Computes the smallest tailcut w such that the Renyi divergence R_a between the ideal distribution $D_{\mathbb{Z}^+, \sigma_{\max}}$ and its restriction to $\{0, \dots, w\}$ (noted $D_{[w], \sigma_{\max}}$) verifies $R_a(D_{[w], \sigma_{\max}}, D_{\mathbb{Z}^+, \sigma_{\max}}) \leq 1 + 1/(4Q_{\text{bs}})$;
2. Rounds the probability density table (PDT) of $D_{[w], \sigma_{\max}}$ with θ bits of absolute precision. This rounding is done “cleverly” by truncating all the PDT values except the largest:
 - for $z \geq 1$, the value $D_{[w], \sigma_{\max}}(z)$ is truncated: $PDT(z) = 2^{-\theta} \lfloor 2^\theta D_{[w], \sigma_{\max}}(z) \rfloor$.
 - in order to have a probability distribution, $PDT(0) = 1 - \sum_{z \geq 1} PDT(z)$.
 Since $R_a(PDT, D_{[w], \sigma_{\max}}) \leq R_\infty(PDT, D_{[w], \sigma_{\max}})$, it is then easy to show that

$$R_a(PDT, D_{[w], \sigma_{\max}}) \leq \frac{PDT(0)}{D_{[w], \sigma_{\max}}(0)} \leq \frac{2^{-\theta w}}{D_{[w], \sigma_{\max}}(0)} \approx \frac{2^{-\theta w}}{\sigma_{\max} \sqrt{\pi/2} + 1}.$$

3. Derives the CDT from the PDT.

Taking $\sigma_{\max} = 1.8205$ and $\theta = 72$ as inputs, we found $w = 19$.

◦ $PDT(0) = 2^{-72} \times 1697680241746640300030$	◦ $PDT(10) = 2^{-72} \times 476288472308334$
◦ $PDT(1) = 2^{-72} \times 1459943456642912959616$	◦ $PDT(11) = 2^{-72} \times 20042553305308$
◦ $PDT(2) = 2^{-72} \times 928488355018011056515$	◦ $PDT(12) = 2^{-72} \times 623729532807$
◦ $PDT(3) = 2^{-72} \times 436693944817054414619$	◦ $PDT(13) = 2^{-72} \times 14354889437$
◦ $PDT(4) = 2^{-72} \times 151893140790369201013$	◦ $PDT(14) = 2^{-72} \times 244322621$
◦ $PDT(5) = 2^{-72} \times 39071441848292237840$	◦ $PDT(15) = 2^{-72} \times 3075302$
◦ $PDT(6) = 2^{-72} \times 7432604049020375675$	◦ $PDT(16) = 2^{-72} \times 28626$
◦ $PDT(7) = 2^{-72} \times 1045641569992574730$	◦ $PDT(17) = 2^{-72} \times 197$
◦ $PDT(8) = 2^{-72} \times 108788995549429682$	◦ $PDT(18) = 2^{-72} \times 1$
◦ $PDT(9) = 2^{-72} \times 8370422445201343$	

One can check that for any $a \geq 509$, $R_a(\text{SampleCDT}_{w=19, \theta=72}, D_{\mathbb{Z}^+, \sigma_{\max}}) \leq 1 + 2^{-80} \leq 1 + \frac{1}{4Q_{\text{bs}}}$, which validates condition (2) for Falcon implementation.

6 Analysis of resistance against timing attacks

In this section, we show that Algorithm 1 is impervious against timing attacks. We formally prove that it is isochronous with respect to σ, μ and the output z (in the sense of Definition 4). We first prove a technical lemma which shows that the number of iterations in the **while** loop of Algorithm 1 is (almost) independent of σ, μ, z .

² We are thankful to Thomas Pornin for bringing up this fact.

Lemma 7. Let $\epsilon \in (0, 1)$, $\mu \in [0, 1]$ and let $\sigma_{\min}, \sigma, \sigma_0$ be standard deviations such that $\eta_\epsilon^+(\mathbb{Z}^n) = \sigma_{\min} \leq \sigma \leq \sigma_0$. Let $p = \frac{\sigma_{\min} \cdot \sqrt{2\pi}}{2 \cdot \rho_{\sigma_{\max}}(\mathbb{Z}^+)}$. The number of iterations of the **while** loop in `SamplerZ`(σ, μ) follows a geometric law of parameter

$$P_{\text{true}}(\sigma, \mu) \in p \cdot \left[1, 1 + \frac{(1 + 2^{-80})\epsilon}{n} \right].$$

The proof of Lemma 7 can be found in Appendix B. Next, we show that Algorithm 1 is perfectly isochronous with respect to z and statistically isochronous (for the Rényi divergence) with respect to σ, μ .

Theorem 8. Let $\epsilon \in (0, 1)$, $\mu \in \mathbb{R}$ and let $\sigma_{\min}, \sigma, \sigma_0$ be standard deviations such that $\eta_\epsilon^+(\mathbb{Z}^n) = \sigma_{\min} \leq \sigma \leq \sigma_0$ and $p = \frac{\sigma_{\min} \cdot \sqrt{2\pi}}{2 \cdot \rho_{\sigma_{\max}}(\mathbb{Z}^+)}$ be a constant in $(0, 1)$. Suppose that the elementary operations $\{+, -, \times, /\}$ over integer and floating-point numbers are isochronous. The running time of Algorithm 1 follows a distribution $T_{\sigma, \mu}$ such that:

$$R_a(T_{\sigma, \mu} \| T) \lesssim 1 + \frac{a\epsilon^2 \max(1, \frac{1-p}{p})^2}{n^2(1-p)} = 1 + O\left(\frac{a\epsilon^2}{n^2}\right)$$

for some distribution T independent of σ, μ and the output z .

Finally, we leverage Theorem 8 to prove that the running time of `SamplerZ`(σ, μ) does not help an adversary to break a cryptographic scheme. We consider that the adversary has access to some function $g(\text{SamplerZ}(\sigma, \mu))$ as well as the running time of `SamplerZ`(σ, μ): this is intended to capture the fact that in practice the output of `SamplerZ`(σ, μ) is not given directly to the adversary, but processed by some function before. For example, in the signature scheme Falcon, samples are processed by algorithms depending on the signer's private key. On the other hand, we consider that the adversary has powerful timing attack capabilities by allowing him to learn the exact runtime of each call to `SamplerZ`(σ, μ).

Corollary 9. Consider an adversary \mathcal{A} making Q_s queries to $g(\text{SamplerZ}(\sigma, \mu))$ for some randomized function g , and solving a search problem with success probability $2^{-\lambda}$ for some $\lambda \geq 1$. With the notations of Theorem 8, suppose that $\max(1, \frac{1-p}{p})^2 \leq n(1-p)$ and $\epsilon \leq \frac{1}{\sqrt{\lambda Q_s}}$. Learning the running time of each call to `SamplerZ`(σ, μ) does not increase the success probability of \mathcal{A} by more than a constant factor.

The proof of Corollary 9 can be found in Appendix D. A nice thing about Corollary 9 is that the conditions required to make it effective are *already met in practice* since they are also required for black-box security of cryptographic schemes. For example, it is systematic to set $\sigma \geq \eta_\epsilon^+(\mathbb{Z}^n)$.

Impact of the scaling factor. The scaling factor $\frac{\sigma_{\min}}{\sigma} \leq \frac{\sigma_{\min}}{\sigma_{\max}}$ is crucial in making our sampler isochronous, as it decorrelates the running time from σ . However, it also impacts the running time T , as one can easily show that T is proportional to the scaling factor. It is therefore desirable to make it as small as possible. The maximal value of the scaling factor is actually dependent on the cryptographic scheme in which our sampler is used. In Appendix E, we show that for the case of the signature scheme Falcon, $\frac{\sigma_{\min}}{\sigma_{\max}} \leq 1.17^{-2} \approx 0.73$ and the impact of the scaling factor is limited. Moreover, one can easily show that for Peikert's sampler [48], the scaling factor is equal to 1 and has no impact.

7 “Err on the side of Gaussian”

This Section focuses on ensuring correct and verified implementations of our proposed isochronous Gaussian sampler. The motivation for this section is based minimising implementation bugs, such as implementation issues with Falcon [51,47] or the famous Heartbleed (CVE-2014-0160) or ROCA vulnerabilities [44] (CVE-2017-15361). We propose a test suite named SAGA (Statistically Acceptable GAussians) in order to verify correct univariate or multivariate Gaussian variables. At the very least, SAGA can act as a ‘sanity check’ for implementors and practitioners; as bugs or errors in implementations of Gaussian samplers can potentially be very costly. Furthermore, SAGA is designed to run in a generic fashion, agnostic to the technique used, by only requiring a list of univariate (i.e., `BaseSampler`) or multivariate (i.e., `SamplerZ`) Gaussian samples as input. Although we evaluate SAGA by applying it to Falcon, SAGA is equally as applicable to any lattice-based cryptographic schemes which require Gaussian sampling, such as Frodo-KEM [43], identity-based encryption [17,8], and in fully homomorphic encryption [54].

7.1 Univariate tests

The statistical tests we implement here are inspired by a previous test suite proposal called GLITCH [29]. We use standard statistical tools to validate a Gaussian sampler is operating with the correct mean, standard deviation, skewness, and kurtosis, and finally we check whether it passes a chi-square normality test. Skewness and kurtosis are descriptors of a normal distribution that respectively measure the symmetry and peakedness of a distribution. To view the full statistical analysis of these tests we created a python class, `UnivariateSamples(mu, sigma, data)`, which take as input the expected mean (`mu`), expected standard deviation (`sigma`), and the list of observed univariate Gaussian samples (`data`). An example of how this works, as well as its output, is shown in Appendix F.1.

7.2 Multivariate tests

This section details multivariate normality tests. The motivation for these tests is to detect situations where the base Gaussian sampler over the integers is correctly implemented, yet the high-level scheme (e.g. a signature scheme) uses it incorrectly way and ends up with a defective multivariate Gaussian.

Multivariate normality: There are a number of statistical tests which evaluate the normality of multivariate distributions. We found that multivariate normality tests predominantly used in other fields [38,28,11] suffer with size and scaling issues. That is, the large sample sizes we expect to use and the poor power properties of these tests will make a type II error highly likely³. In fact, we implemented the Mardia [38] and Henze-Zirkler [28] tests and found, although they worked for small sample sizes, they diverged to produce false negatives for sample sizes ≥ 50 in small dimensions $n = 64$.

However, the Doornik-Hansen test [13] minimises these issues by using transformed versions of the skewness and kurtosis of the multivariate data, increasing the test’s power. We also note that it has a much faster computation time (essentially linear in the sample size) compared to [38,28] (essentially quadratic in the sample size). As with the univariate tests, we created a python class, denoted `MultivariateSamples(sigma,data)`, which can output four results; two based on the covariance matrix, and two based on the data’s normality. An example of how this works, as well as its output, is shown in Appendix F.2.

³ Type I and type II errors are, respectively, rejection of a true null hypothesis and the non-rejection of a false null hypothesis.

A glitch in the (covariance) matrix: Our second test asks the following question: how would someone implement correctly the base sampler, yet subsequently fail to use it properly? There is no universal answer to that, and one usually has to rely on context, experience and common sense to establish the most likely way this could happen. For example, in the signature scheme Falcon, univariate Gaussian samples are combined according to values contained in a balanced binary tree computed at key generation; there could be an implementation mistake in the procedure computing the tree (during key generation), or in the procedure combining the samples (during signing). Such mistakes have a very recognizable effect on the empiric covariance matrix of Falcon signatures: they make them look like block-Toeplitz matrices, as illustrated by Figure 1a. We therefore devised a test which differentiate block Toeplitz covariance matrices from the ones we would expect from spherical Gaussians.

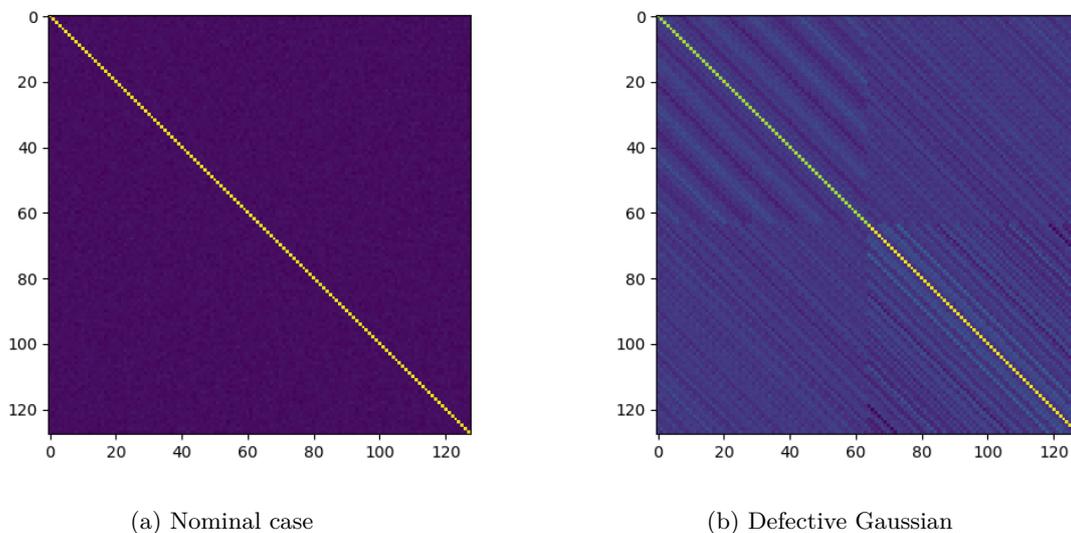


Figure 1: Empiric covariance matrices of Falcon signatures. Figure 1a corresponds to a correct implementation of Falcon. Figure 1b corresponds to an implementation where there is a mistake when constructing the Falcon tree.

Supplementary tests: In the case where normality has been rejected, SAGA also provides a number of extra tests to aid in finding the issues. More details for this can be found in Appendix F.3.

8 Application and Limitations

Our sampler has been implemented by Pornin as part of the new isochronous implementation of Falcon [51]. This implementation can use floating-point hardware or AVX2 instructions when available, but also includes floating-point emulation code that uses only usual integer operations. We perform benchmarks of this sampler implementation on a single Intel Core i7-6500U CPU core clocked at 2.5 GHz. In Table 2 we present the running times of our isochronous sampler. To compare with [60], we scale the numbers to be based on 2.5GHz. Note that for our sampler the number of samples per second is in average for $1.2915 < \sigma \leq 1.8502$ while for [60] $\sigma = 2$ is fixed.

Table 2: Number of samples per second at 2.5 GHz for our sampler and [60].

Algorithm	Number of samples
This work ⁴	$1.84 \times 10^6/\text{sec}$
This work (AVX2) ⁵	$7.74 \times 10^6/\text{sec}$
[60] (AVX2) ⁶	$5.43 \times 10^6/\text{sec}$

In Table 3 we present the running times of the Falcon isochronous implementation [51] that contains our sampler and compare it with a second non-isochronous implementation nearly identical excepting the base sampler which is a faster lazy CDT sampler, and the rejection sampling which is not scaled by a constant. Compared to the non-isochronous implementation, the isochronous one is about 22% slower, but remains very competitive speed-wise.

Table 3: Falcon signature generation time at 2.5 GHz.

Degree	Non-isochronous (using AVX2)	isochronous (using AVX2)
512	210.88 μs (153.64 μs)	257.33 μs (180.04 μs)
1024	418.76 μs (311.33 μs)	515.28 μs (361.39 μs)

Advantages and limitations. Our sampler has an acceptance rate $\approx \frac{\sigma_{\min}}{\sigma_{\max}+0.4}$ making it especially suitable when σ_{\min} and σ_{\max} are close. In particular, our sampler is, so far, the fastest isochronous sampler for the parameters in Falcon. However, the larger the gap between σ_{\min} and σ_{\max} , the lower the acceptance rate. In addition, our sampler uses a cumulative distribution table (CDT) which is accessed in an isochronous way. This table grows when σ_{\max} grows, while making both running time and memory usage larger. When σ_{\max} is large or far from σ_{\min} , there exist faster isochronous samplers based on convolution [42] and rejection sampling [60]⁷ techniques.

Acknowledgements

We thank Léo Ducas for helpful suggestions. We also thank Thomas Pornin and Mehdi Tibouchi for useful discussions. The first and second authors were supported by the project PQ Cybersecurity (Innovate UK research grant 104423). The third and fourth authors were supported by BPI-France in the context of the national project RISQ (P141580), and by the European Union PROMETHEUS project (Horizon 2020 Research and Innovation Program, grant 780701). The fourth author was also supported by ANRT under the program CIFRE N2016/1583.

⁴ [51] standard double-precision floating-point (IEEE 754) implementation that runs SHAKE256 to generate the randomness.

⁵ [51] AVX2 implementation that runs eight ChaCha20 instances in parallel (using AVX2 instructions) to generate the randomness.

⁶ [60] constant-time implementation that runs hardware AES256 (AES-NI) instructions to generate the randomness.

⁷ The constant-time sampler in [60] may still reveal σ .

References

1. Shweta Agrawal, Dan Boneh, and Xavier Boyen. Efficient lattice (H)IBE in the standard model. In Gilbert [27], pages 553–572.
2. Joachim Ahrens and Ulrich Dieter. Extension of forsythe’s method for random sampling from the normal distribution. *Mathematics of computation*, 27:927–937, 1973.
3. Shi Bai, Adeline Langlois, Tancrede Lepoint, Damien Stehlé, and Ron Steinfeld. Improved security proofs in lattice-based cryptography: Using the Rényi divergence rather than the statistical distance. In Tetsu Iwata and Jung Hee Cheon, editors, *ASIACRYPT 2015, Part I*, volume 9452 of *LNCS*, pages 3–24. Springer, Heidelberg, November / December 2015.
4. Gilles Barthe, Sonia Belad, Thomas Espitau, Pierre-Alain Fouque, Mélissa Rossi, and Mehdi Tibouchi. GALACTICS: Gaussian Sampling for Lattice-Based Constant-Time Implementation of Cryptographic Signatures, Revisited. Cryptology ePrint Archive, Report 2019/511, 2019. <https://eprint.iacr.org/2019/511>.
5. Nina Bindel, Sedat Akleylek, Erdem Alkim, Paulo S. L. M. Barreto, Johannes Buchmann, Edward Eaton, Gus Gutoski, Juliane Kramer, Patrick Longa, Harun Polat, Jefferson E. Ricardini, and Gustavo Zanon. qTESLA. Technical report, National Institute of Standards and Technology, 2019. available at <https://csrc.nist.gov/projects/post-quantum-cryptography/round-2-submissions>.
6. Joachim Breitner and Nadia Heninger. Biased nonce sense: Lattice attacks against weak ecdsa signatures in cryptocurrencies. In Ian Goldberg and Tyler Moore, editors, *Financial Cryptography and Data Security*, pages 3–20, Cham, 2019. Springer International Publishing.
7. Leon Groot Bruinderink, Andreas Hülsing, Tanja Lange, and Yuval Yarom. Flush, gauss, and reload - A cache attack on the BLISS lattice-based signature scheme. In Benedikt Gierlich and Axel Y. Poschmann, editors, *CHES 2016*, volume 9813 of *LNCS*, pages 323–345. Springer, Heidelberg, August 2016.
8. Peter Campbell and Michael Groves. Practical post-quantum hierarchical identity-based encryption. 16th IMA International Conference on Cryptography and Coding, 2017. <http://www.qub.ac.uk/sites/CSIT/FileStore/Fileupload,785752,en.pdf>.
9. David Cash, Dennis Hofheinz, Eike Kiltz, and Chris Peikert. Bonsai trees, or how to delegate a lattice basis. In Gilbert [27], pages 523–552.
10. Yilei Chen, Nicholas Genise, and Pratyay Mukherjee. Approximate trapdoors for lattices and smaller hash-and-sign signatures. *IACR Cryptology ePrint Archive*, 2019:1029, 2019.
11. David Roxbee Cox and NJH Small. Testing multivariate normality. *Biometrika*, 65(2):263–272, 1978.
12. Ralph B D’Agostino, Albert Belanger, and Ralph B D’Agostino Jr. A suggestion for using powerful and informative tests of normality. *The American Statistician*, 44(4):316–321, 1990.
13. Jurgen A Doornik and Henrik Hansen. An omnibus test for univariate and multivariate normality. *Oxford Bulletin of Economics and Statistics*, 70:927–939, 2008.
14. Yusong Du, Baodian Wei, and Huang Zhang. A rejection sampling algorithm for off-centered discrete gaussian distributions over the integers. *Science China Information Sciences*, 62(3):39103, Sep 2018.
15. Léo Ducas. *Signatures fondées sur les réseaux euclidiens : attaques, analyses et optimisations*. Theses, École Normale Supérieure, 2013.
16. Léo Ducas, Alain Durmus, Tancrede Lepoint, and Vadim Lyubashevsky. Lattice signatures and bimodal Gaussians. In Ran Canetti and Juan A. Garay, editors, *CRYPTO 2013, Part I*, volume 8042 of *LNCS*, pages 40–56. Springer, Heidelberg, August 2013.
17. Léo Ducas, Vadim Lyubashevsky, and Thomas Prest. Efficient identity-based encryption over NTRU lattices. In Palash Sarkar and Tetsu Iwata, editors, *ASIACRYPT 2014, Part II*, volume 8874 of *LNCS*, pages 22–41. Springer, Heidelberg, December 2014.
18. Léo Ducas and Phong Q. Nguyen. Faster Gaussian lattice sampling using lazy floating-point arithmetic. In Xiaoyun Wang and Kazue Sako, editors, *ASIACRYPT 2012*, volume 7658 of *LNCS*, pages 415–432. Springer, Heidelberg, December 2012.
19. Nagarjun C Dwarakanath and Steven D Galbraith. Sampling from discrete gaussians for lattice-based cryptography on a constrained device. *Applicable Algebra in Engineering, Communication and Computing*, 25(3):159–180, 2014.
20. Thomas Espitau, Pierre-Alain Fouque, Benoît Gérard, and Mehdi Tibouchi. Side-channel attacks on BLISS lattice-based signatures: Exploiting branch tracing against strongSwan and electromagnetic emanations in microcontrollers. In Thuraisingham et al. [55], pages 1857–1874.
21. Gerald Estrin. Organization of computer systems: The fixed plus variable structure computer. In *Papers Presented at the May 3-5, 1960, Western Joint IRE-AIEE-ACM Computer Conference*, IRE-AIEE-ACM ’60 (Western), pages 33–40, New York, NY, USA, 1960. ACM.
22. James J. Filliben and Alan Heckert. 1.3.3.24. *Quantile-Quantile Plot*. 2013.
23. George E. Forsythe. Von neumann’s comparison method for random sampling from the normal and other distributions. *Mathematics of Computation*, 26(120):817–826, 1972.

24. Nicolas Gama, Nick Howgrave-Graham, and Phong Q. Nguyen. Symplectic lattice reduction and NTRU. In Serge Vaudenay, editor, *EUROCRYPT 2006*, volume 4004 of *LNCS*, pages 233–253. Springer, Heidelberg, May / June 2006.
25. Nicholas Genise and Daniele Micciancio. Faster gaussian sampling for trapdoor lattices with arbitrary modulus. In Jesper Buus Nielsen and Vincent Rijmen, editors, *EUROCRYPT 2018, Part I*, volume 10820 of *LNCS*, pages 174–203. Springer, Heidelberg, April / May 2018.
26. Craig Gentry, Chris Peikert, and Vinod Vaikuntanathan. Trapdoors for hard lattices and new cryptographic constructions. In Richard E. Ladner and Cynthia Dwork, editors, *40th ACM STOC*, pages 197–206. ACM Press, May 2008.
27. Henri Gilbert, editor. *EUROCRYPT 2010*, volume 6110 of *LNCS*. Springer, Heidelberg, May / June 2010.
28. N Henze and B Zirkler. A class of invariant consistent tests for multivariate normality. *Communications in Statistics-Theory and Methods*, 19(10):3595–3617, 1990.
29. James Howe and Máire O’Neill. GLITCH: A discrete gaussian testing suite for lattice-based cryptography. In *Proceedings of the 14th International Joint Conference on e-Business and Telecommunications (ICETE 2017) - Volume 4: SECRYPT, Madrid, Spain, July 24-26, 2017.*, pages 413–419, 2017.
30. Andreas Hülsing, Tanja Lange, and Kit Smeets. Rounded gaussians - fast and secure constant-time sampling for lattice-based crypto. In Michel Abdalla and Ricardo Dahab, editors, *PKC 2018, Part II*, volume 10770 of *LNCS*, pages 728–757. Springer, Heidelberg, March 2018.
31. Angshuman Karmakar, Sujoy Sinha Roy, Frederik Vercauteren, and Ingrid Verbauwhede. Pushing the speed limit of constant-time discrete gaussian sampling. a case study on falcon. Cryptology ePrint Archive, Report 2019/267, 2019. <https://eprint.iacr.org/2019/267>.
32. Charles F. F. Karney. Sampling exactly from the normal distribution. *ACM Trans. Math. Softw.*, 42(1):3:1–3:14, January 2016.
33. Selcuk Korkmaz, Dincer Goksuluk, and Gokmen Zararsiz. Mvn: An r package for assessing multivariate normality. *The R Journal*, 6(2):151–162, 2014.
34. Vadim Lyubashevsky. Fiat-Shamir with aborts: Applications to lattice and factoring-based signatures. In Mitsuru Matsui, editor, *ASIACRYPT 2009*, volume 5912 of *LNCS*, pages 598–616. Springer, Heidelberg, December 2009.
35. Vadim Lyubashevsky, Léo Ducas, Eike Kiltz, Tancrede Lepoint, Peter Schwabe, Gregor Seiler, and Damien Stehlé. CRYSTALS-DILITHIUM. Technical report, National Institute of Standards and Technology, 2019. available at <https://csrc.nist.gov/projects/post-quantum-cryptography/round-2-submissions>.
36. Vadim Lyubashevsky and Daniel Wichs. Simple lattice trapdoor sampling from a broad class of distributions. In Jonathan Katz, editor, *PKC 2015*, volume 9020 of *LNCS*, pages 716–730. Springer, Heidelberg, March / April 2015.
37. Prasanta Chandra Mahalanobis. On the generalized distance in statistics. National Institute of Science of India, 1936.
38. Kanti V Mardia. Measures of multivariate skewness and kurtosis with applications. *Biometrika*, 57(3):519–530, 1970.
39. Carlos Aguilar Melchor and Thomas Ricosset. Cdt-based gaussian sampling: From multi to double precision. *IEEE Trans. Computers*, 67(11):1610–1621, 2018.
40. Daniele Micciancio and Chris Peikert. Trapdoors for lattices: Simpler, tighter, faster, smaller. In David Pointcheval and Thomas Johansson, editors, *EUROCRYPT 2012*, volume 7237 of *LNCS*, pages 700–718. Springer, Heidelberg, April 2012.
41. Daniele Micciancio and Oded Regev. Worst-case to average-case reductions based on gaussian measures. *SIAM J. Comput.*, 37(1):267–302, 2007.
42. Daniele Micciancio and Michael Walter. Gaussian sampling over the integers: Efficient, generic, constant-time. In Jonathan Katz and Hovav Shacham, editors, *CRYPTO 2017, Part II*, volume 10402 of *LNCS*, pages 455–485. Springer, Heidelberg, August 2017.
43. Michael Naehrig, Erdem Alkim, Joppe Bos, Léo Ducas, Karen Easterbrook, Brian LaMacchia, Patrick Longa, Ilya Mironov, Valeria Nikolaenko, Christopher Peikert, Ananth Raghunathan, and Douglas Stebila. FrodoKEM. Technical report, National Institute of Standards and Technology, 2019. available at <https://csrc.nist.gov/projects/post-quantum-cryptography/round-2-submissions>.
44. Matúš Nemeč, Marek Šýs, Petr Svenda, Dusan Klinec, and Vashek Matyas. The return of coppersmith’s attack: Practical factorization of widely used RSA moduli. In Thuraisingham et al. [55], pages 1631–1648.
45. NIST et al. OFFICIAL COMMENT: CRYSTALS-DILITHIUM, 2018. <https://groups.google.com/a/list.nist.gov/d/msg/pqc-forum/aWxC2ynJDLE/Y0sMJ2ewAAAJ>.
46. NIST et al. Footguns as an axis for security analysis, 2019. <https://groups.google.com/a/list.nist.gov/forum/#!topic/pqc-forum/l2iYk-8sGnI> last accessed 23-09-2019.
47. NIST et al. OFFICIAL COMMENT: Falcon (bug & fixes), 2019. <https://groups.google.com/a/list.nist.gov/forum/#!topic/pqc-forum/7Z8x5AMXy8s> last accessed on 23-09-2019.
48. Chris Peikert. An efficient and parallel Gaussian sampler for lattices. In Tal Rabin, editor, *CRYPTO 2010*, volume 6223 of *LNCS*, pages 80–97. Springer, Heidelberg, August 2010.

49. Peter Pessl, Leon Groot Bruinderink, and Yuval Yarom. To BLISS-B or not to be: Attacking strongSwan’s implementation of post-quantum signatures. In Thuraisingham et al. [55], pages 1843–1855.
50. Thomas Pöppelmann, Léo Ducas, and Tim Güneysu. Enhanced lattice-based signatures on reconfigurable hardware. In Lejla Batina and Matthew Robshaw, editors, *CHES 2014*, volume 8731 of *LNCS*, pages 353–370. Springer, Heidelberg, September 2014.
51. Thomas Pornin. New Efficient, Constant-Time Implementations of Falcon. Cryptology ePrint Archive, Report 2019/893, 2019. <https://eprint.iacr.org/2019/893>.
52. Thomas Prest. Sharper bounds in lattice-based cryptography using the Rényi divergence. In Tsuyoshi Takagi and Thomas Peyrin, editors, *ASIACRYPT 2017, Part I*, volume 10624 of *LNCS*, pages 347–374. Springer, Heidelberg, December 2017.
53. Thomas Prest, Pierre-Alain Fouque, Jeffrey Hoffstein, Paul Kirchner, Vadim Lyubashevsky, Thomas Pornin, Thomas Ricosset, Gregor Seiler, William Whyte, and Zhenfei Zhang. FALCON. Technical report, National Institute of Standards and Technology, 2019. available at <https://csrc.nist.gov/projects/post-quantum-cryptography/round-2-submissions>.
54. Microsoft SEAL (release 3.4). <https://github.com/Microsoft/SEAL>, October 2019. Microsoft Research, Redmond, WA.
55. Bhavani M. Thuraisingham, David Evans, Tal Malkin, and Dongyan Xu, editors. *ACM CCS 2017*. ACM Press, October / November 2017.
56. Mehdi Tibouchi and Alexandre Wallet. One bit is all it takes: A devastating timing attack on bliss’s non-constant time sign flips. Cryptology ePrint Archive, Report 2019/898, 2019. <https://eprint.iacr.org/2019/898>.
57. John von Neumann. Various techniques used in connection with random digits. *National Bureau of Standards, Applied Math Series*, 12:36–38, 11 1950.
58. Michael Walter. Sampling the integers with low relative error. Cryptology ePrint Archive, Report 2019/068, 2019. <https://eprint.iacr.org/2019/068>.
59. Raymond K. Zhao, Ron Steinfeld, and Amin Sakzad. Facct: Fast, compact, and constant-time discrete gaussian sampler over integers. Cryptology ePrint Archive, Report 2018/1234, 2018. <https://eprint.iacr.org/2018/1234>.
60. Raymond K. Zhao, Ron Steinfeld, and Amin Sakzad. Compact and scalable arbitrary-centered discrete gaussian sampling over integers. Cryptology ePrint Archive, Report 2019/1011, 2019. <https://eprint.iacr.org/2019/1011>.

A Proof of Theorem 6

To estimate the security loss with the replacement of $D_{\mathbb{Z},\sigma,\mu}$ by `SamplerZ`, we introduce an intermediate case where `ApproxExp` outputs a perfect value. The 3 cases are defined as follows.

1. (IDEAL) The ideal $D_{\mathbb{Z},\sigma,\mu}$ is called. By Proposition 5, it is the same as considering `ApproxExp` = `exp` and `BaseSampler` = $D_{\mathbb{Z}^+, \sigma_{\max}}$.
2. (INTER) Only `ApproxExp` = `exp` is assumed.
3. (REAL) The imperfect `SamplerZ` is called.

We recall that λ_{REAL} (resp. λ_{IDEAL}) is the security parameter of the REAL (resp. IDEAL) case. We aim at computing $\Delta\lambda = \lambda_{\text{IDEAL}} - \lambda_{\text{REAL}}$.

We denote by $a := 2 \cdot \lambda_{\text{REAL}} + 1$. The values $R_a(\text{REAL}, \text{INTER})$ and $R_a(\text{INTER}, \text{IDEAL})$ will be used to quantify the distance between each case. Let E be an event breaking the scheme. Let δ_{IDEAL} (resp. δ_{INTER} , δ_{REAL}) be the probability that this event occurs in the use of the IDEAL (resp. INTER, REAL) case. By probability preservation of the Rényi divergence:

$$\delta_{\text{IDEAL}} \geq \delta_{\text{INTER}}^{\frac{a}{a-1}} / R_a(\text{INTER}, \text{IDEAL})^{Q_{\text{bs}}}$$

$$\delta_{\text{INTER}} \geq \delta_{\text{REAL}}^{\frac{a}{a-1}} / R_a(\text{REAL}, \text{INTER})^{Q_{\text{exp}}}.$$

By definition, $\delta_{\text{REAL}} \geq 2^{-\lambda_{\text{REAL}}}$, thus, the second equation can be upper bounded using $\delta_{\text{REAL}}^{\frac{a}{a-1}} \geq \delta_{\text{REAL}} / \sqrt{2}$. By combining it,

$$\delta_{\text{INTER}} \geq \delta_{\text{REAL}} / \left(\sqrt{2} \cdot R_a(\text{REAL}, \text{INTER})^{Q_{\text{exp}}} \right).$$

And thus,

$$\begin{aligned}\delta_{\text{IDEAL}} &\geq \delta_{\text{REAL}}^{\frac{a}{a-1}} \cdot \left(\sqrt{2^{\frac{a}{a-1}}} \cdot R_a(\text{REAL}, \text{INTER})^{\frac{aQ_{\text{exp}}}{a-1}} R_a(\text{INTER}, \text{IDEAL})^{Q_{\text{bs}}} \right)^{-1} \\ &\geq \delta_{\text{REAL}} \cdot \left(\sqrt{2^{\frac{a}{a-1}+1}} \cdot R_a(\text{REAL}, \text{INTER})^{\frac{aQ_{\text{exp}}}{a-1}} R_a(\text{INTER}, \text{IDEAL})^{Q_{\text{bs}}} \right)^{-1}\end{aligned}$$

So,

$$\Delta\lambda = \log_2 \left(\sqrt{2^{\frac{a}{a-1}+1}} \cdot R_a(\text{REAL}, \text{INTER})^{\frac{aQ_{\text{exp}}}{a-1}} \cdot R_a(\text{INTER}, \text{IDEAL})^{Q_{\text{bs}}} \right)$$

Let us now use the conditions to get a concrete upper bound on $\Delta\lambda$. First, suppose that condition (1) is verified. We use $a := 2 \cdot \lambda_{\text{REAL}} + 1$, then, for all $x < 0$:

$$1 - \sqrt{\frac{a-1}{2 \cdot a^2 \cdot Q_{\text{exp}}}} \leq \frac{P(x)}{\text{ApproxExp}(x)} \leq 1 + \sqrt{\frac{a-1}{2 \cdot a^2 \cdot Q_{\text{exp}}}}.$$

An application of Lemma 3 yields to $R_a(\text{REAL}, \text{INTER}) \leq 1 + \frac{a-1}{4aQ_{\text{exp}}}$.

Secondly, suppose that condition (2) is verified. Recall that $BG_{\sigma_{\text{max}}}$ denotes the ideal distribution of z before rejection sampling (Step 4). Let us denote by $\overline{BG}_{\sigma_{\text{max}}}$ the distribution of z before the rejection sampling when `BaseSampler` is not perfect. The next steps in `SamplerZ` algorithm consist in multiplying the output distribution by $z \mapsto \frac{\sigma_{\text{min}}}{\sigma} \exp\left(\frac{(z-\mu)^2}{2\sigma^2} - \frac{z_0^2}{2\sigma_{\text{max}}^2}\right)$. By data processing, we get

$$R_a(\text{INTER}, \text{IDEAL}) \leq R_a(\overline{BG}_{\sigma_{\text{max}}}, BG_{\sigma_{\text{max}}})$$

Then, since (considering the distribution of b as perfectly uniform)

$$R_a(\text{INTER}, \text{IDEAL}) \leq R_a((2b-1)\text{BaseSampler} + b, (2b-1)D_{\mathbb{Z}^+, \sigma_{\text{max}}} + b),$$

we re-apply data processing and obtain

$$\begin{aligned}R_a(\text{INTER}, \text{IDEAL}) &\leq R_a(\text{BaseSampler}, D_{\mathbb{Z}^+, \sigma_{\text{max}}}) \\ &\leq 1 + \frac{1}{4Q_{\text{bs}}}\end{aligned}$$

Wrapping up,

$$\begin{aligned}\Delta\lambda &= \log_2 \left(\sqrt{2^{\frac{a}{a-1}+1}} \cdot R_a(\text{REAL}, \text{INTER})^{\frac{aQ_{\text{exp}}}{a-1}} \cdot R_a(\text{INTER}, \text{IDEAL})^{Q_{\text{bs}}} \right) \\ &\leq \log_2 \left(\sqrt{2^{\frac{a}{a-1}+1}} \cdot \left(1 + \frac{a-1}{4aQ_{\text{exp}}}\right)^{\frac{aQ_{\text{exp}}}{a-1}} \cdot \left(1 + \frac{1}{4Q_{\text{bs}}}\right)^{Q_{\text{bs}}} \right).\end{aligned}$$

Using the inequality $(1 + \frac{x}{n})^n \leq \exp(x)$ for $x, n > 0$,

$$\begin{aligned}\Delta\lambda &\leq \log_2 \left(\sqrt{2^{\frac{a}{a-1}+1}} \cdot \exp(1/4)^2 \right) \\ &= \log_2 \left(\sqrt{2^{\frac{a}{a-1}+1}} \cdot 2 \right) \\ &\leq 2.\end{aligned}$$

B Proof of Lemma 7

Proof. We note $p_1(z)$ the probability that $z \in \mathbb{Z}$ (with a uniquely associated z_0) is output by `SamplerZ` in a given iteration. It holds that:

$$p_1(z) = \underbrace{\frac{\rho_{\sigma_{\max}}(z_0)}{\rho_{\sigma_{\max}}(\mathbb{Z}^+)}}_{\mathbb{P}[\text{BaseSampler} \rightarrow z_0]} \cdot \underbrace{\frac{1}{2}}_{\mathbb{P}[b]} \cdot \underbrace{\frac{\sigma_{\min}}{\sigma} \cdot \frac{\rho_{\sigma, \mu}(z)}{\rho_{\sigma_{\max}}(z_0)}}_{\mathbb{P}[\text{AcceptSample} \rightarrow \text{true} | z]} = \frac{\sigma_{\min} \cdot \rho_{\sigma, \mu}(z)}{2 \cdot \sigma \cdot \rho_{\sigma_{\max}}(\mathbb{Z}^+)}.$$

The probability P_{true} that (`AcceptSample` \rightarrow `true`) in a given iteration is:

$$P_{\text{true}} = \mathbb{P}[\text{AcceptSample} \rightarrow \text{true}] = \sum_z p_1(z) = \frac{\sigma_{\min} \cdot \rho_{\sigma, \mu}(\mathbb{Z})}{2 \cdot \sigma \cdot \rho_{\sigma_{\max}}(\mathbb{Z}^+)}. \quad (5)$$

One can see in (5) that P_{true} is independent of the output z . This is unsurprising since a different z is picked at each new iteration of the `while` loop, and each iteration's running time is constant.

However, it is not obvious from (5) that P_{true} is independent of σ and μ ; we now show that it is essentially the case. Since $\sigma \geq \eta_{\epsilon}^+(\mathbb{Z}^n) \geq \eta_{\epsilon/n}^+(\mathbb{Z}) \geq \eta_{\epsilon/n}(\mathbb{Z})$, it holds from [26, Lemma 2.7] that:

$$\rho_{\sigma, \mu}(\mathbb{Z}) \in \left[\frac{1 - \epsilon/n}{1 + \epsilon/n}, 1 \right] \cdot \rho_{\sigma}(\mathbb{Z}). \quad (6)$$

It is now helpful to bound $\rho_{\sigma}(\mathbb{Z})$. By the Poisson summation formula:

$$\rho_{\sigma}(\mathbb{Z}) = \sigma \sqrt{2\pi} \cdot \left(1 + 2 \sum_{i \geq 1} \exp(-2i^2 \pi^2 \sigma^2) \right). \quad (7)$$

For any $\sigma > 1$, it holds that $\sum_{i \geq 1} \exp(-2i^2 \pi^2 \sigma^2) \in \exp(-2\pi^2 \sigma^2) \cdot [1, 1 + 2^{-80}]$. Moreover, it follows from (3) that $\exp(-2\pi^2 \sigma^2) \leq \frac{\epsilon}{2n}$. Combined this fact with (7) yields:

$$\sigma \sqrt{2\pi} \leq \rho_{\sigma}(\mathbb{Z}) \leq \sigma \sqrt{2\pi} \cdot \left(1 + (1 + 2^{-80}) \cdot \frac{\epsilon}{n} \right) \quad (8)$$

Finally, combining (5), (6) and (8) yields:

$$P_{\text{true}}(\sigma, \mu) \in \frac{\sigma_{\min} \cdot \sqrt{2\pi}}{2 \cdot \rho_{\sigma_{\max}}(\mathbb{Z}^+)} \cdot \left(1, 1 + \frac{(1 + 2^{-80})\epsilon}{n} \right). \quad (9)$$

This concludes the proof. \square

C Proof of Theorem 8

Before proving Theorem 8, we will need a preliminary lemma. Note that this lemma cannot be proven in a black-box way using Lemma 3 since the relative error between any two distinct geometric distributions is infinite.

Lemma 10. *Let \mathcal{P} and \mathcal{Q} be geometric distributions of parameters $p, q \geq C$ for a constant $C \geq 0$. Suppose there exists $\delta = o(1/(a+1))$ such that:*

$$\begin{aligned} e^{-\delta} &\leq \frac{p}{q} \leq e^{\delta}, \\ e^{-\delta} &\leq (1-p)/(1-q) \leq e^{\delta}. \end{aligned}$$

Then the Rényi divergence between \mathcal{P} and \mathcal{Q} is bounded as follows:

$$R_a(\mathcal{P} \parallel \mathcal{Q}) \lesssim 1 + \frac{a(1-p)\delta^2}{p^2} \left(\sim 1 + \frac{a(1-q)\delta^2}{q^2} \right).$$

Proof. The beginning of the proof follows the one of [52, Lemma 3], but makes a more precise estimation. Let $f_a : (x, y) \mapsto \frac{y^a}{(x+y)^{a-1}}$. We compute values of f_a and its derivatives around $(0, y)$:

$$\begin{aligned} f_a(x, y) &= y && \text{for } x = 0 \\ \frac{\partial f_a}{\partial x}(x, y) &= 1 - a && \text{for } x = 0 \\ \frac{\partial^2 f_a}{\partial x^2}(x, y) &= a(a-1)y^a(x+y)^{-a-1} \\ &\leq \frac{a(a-1)}{e^{-(a+1)\delta}y} && \text{for } |x+y| \leq e^\delta \cdot y \end{aligned}$$

We now use partial Taylor bounds. If $|x_k| \leq (e^{\delta k} - 1) \cdot y_k$, then:

$$f_a(x_k, y_k) \leq f_a(0, y) + \frac{\partial f_a}{\partial x}(0, y_k) \cdot x_k + \frac{a(a-1)(e^{\delta k} - 1)^2}{2e^{-(a+1)\delta k}} \cdot y_k$$

We take $y_k = \mathcal{P}(k)$ and $x_k = \mathcal{Q}(k) - \mathcal{P}(k)$, and note that $e^{-k\delta} \leq \mathcal{P}(k)/\mathcal{Q}(k) \leq e^{k\delta}$, hence $\delta_k \leq k\delta$. Summing all over the support of \mathcal{P} gives:

$$R_a^a(\mathcal{P} \parallel \mathcal{Q}) \leq 1 + \frac{a(a-1)}{2} \sum_{k \in \mathbb{Z}^+} \frac{(e^{k\delta} - 1)^2}{e^{-(a+1)k\delta}} \cdot (1-p)^{k-1} p \quad (10)$$

$$\begin{aligned} &\leq 1 + \frac{pa(a-1)}{2(1-p)} \cdot \left(\frac{2(1-p)^2\delta^2}{p^3} + O\left(\frac{a\delta^3(p-1)^3}{p^4}\right) \right) \quad (11) \\ &\lesssim 1 + \frac{a(a-1)(1-p)\delta^2}{p^2} \end{aligned}$$

To compute the sum in (10), we expanded $(e^{k\delta} - 1)^2$, which then gives us three distinct geometric sums which all converge since $\delta = o(1/(a+1))$, and for which closed formulae are known. A tedious but easy Taylor expansion then gives us 11, at which point we can conclude. \square

We now prove Theorem 8.

Proof. Let T_0 denote the running time of one iteration of the **while** loop in Algorithm 1. It is clear that steps 2 to 5 are isochronous. On the other hand, it is also clear that **AcceptSample** (Algorithm 2) is isochronous; indeed, all its atomic operations are isochronous and each iteration of its **while** loop has a constant probability $1 - 2^{-8}$ of being the last one. We can conclude that T_0 follows a distribution which is independent of σ, μ, z .

Let us denote by $I_{\sigma, \mu}$ (resp. I) the number of iterations of the **while** loop when each iteration accepts with probability P_{true} (resp. p). By Lemma 7, $I_{\sigma, \mu}$ (resp. I) follows a geometric law of parameter P_{true} (resp. p). The relative error between P_{true} and p is upper bounded by $\delta = \frac{(1+2^{-80})^\epsilon}{n} \cdot \max(1, \frac{p}{1-p})$. It then follows from Lemma 10 that:

$$\begin{aligned} R_a(I_{\sigma, \mu} \parallel I) &\lesssim 1 + \frac{a(1-p)\delta^2}{p^2} \\ &\lesssim 1 + \frac{a\epsilon^2 \max(1, \frac{1-p}{p})^2}{n^2(1-p)} \end{aligned}$$

Finally, the total running time T of Algorithm 1 is a function of the running time of each iteration and the number of iterations: $T = f(T_0, I)$ for some function f . This allows to apply once again the data-processing inequality:

$$R_a(T_{\sigma, \mu} \parallel T) = R_a(f(T_0, I_{\sigma, \mu}) \parallel f(T_0, I)) \leq R_a(I_{\sigma, \mu} \parallel I),$$

which concludes the proof. \square

D Proof of Corollary 9

Proof. Let D denote the output distribution of $g(\text{SamplerZ}(\sigma, \mu))$. In the ideal case, we can consider without loss of generality that the adversary can query the joint distribution (D, T) , where T is as in the proof of Theorem 8 and is thus independent from σ, μ . In the real case, the adversary learns the runtime of each call to $\text{SamplerZ}(\sigma, \mu)$. Since we showed in the proof of Theorem 8 that the runtime is independent of the output z , we can model it as the distribution $T_{\sigma, \mu}$ described in the proof of Theorem 8, and the adversary has access to $(D, T_{\sigma, \mu})$.

Let P_0, P_1 denote the success probability of \mathcal{A} in the ideal and real cases, respectively. Taking $a = \lambda$ gives:

$$P_1 \leq [P_0 \cdot R_a((D, T_{\sigma, \mu})^{Q_s} \parallel (D, T)^{Q_s})]^{(a-1)/a} \quad (12)$$

$$\leq [P_0 \cdot R_a((D, T_{\sigma, \mu}) \parallel (D, T))^{Q_s}]^{(a-1)/a} \quad (13)$$

$$\leq [P_0 \cdot R_a(T_{\sigma, \mu} \parallel T)^{Q_s}]^{(a-1)/a} \quad (14)$$

$$\lesssim \left[P_0 \cdot \left(1 + \frac{a\epsilon^2 \max(1, \frac{1-p}{p})^2}{n^2(1-p)} \right)^{Q_s} \right]^{(a-1)/a} \quad (15)$$

$$\lesssim \left[P_0 \cdot \left(1 + \frac{1}{n \cdot Q_s} \right)^{Q_s} \right]^{(\lambda-1)/\lambda} \quad (16)$$

$$\lesssim 2^{-(\lambda-1)} \cdot e^{1/n} \quad (17)$$

Hereabove, (12) uses the probability preservation property of the Renyi divergence, and (13) uses its multiplicativity. We previously showed that D and $T_{\sigma, \mu}$ are independent, thus we can discard D in (14). We then apply Theorem 8 to get (15). We replace ϵ by $\frac{1}{\sqrt{\lambda Q_s}}$, take $a = \lambda$ and $\max(1, \frac{1-p}{p})^2 \leq n(1-p)$ to obtain (16). Finally, we use the identity $(1 + x/k)^k \leq e^x$ and replace P_0 by $2^{-\lambda}$ to get (17). \square

E Impact of the scaling factor in Falcon

We now study the impact of the scaling factor $\frac{\sigma_{\min}}{\sigma} \leq \frac{\sigma_{\min}}{\sigma_{\max}}$ on the running time for the particular case of Falcon. There, each σ verifies $\sigma_{\min} \leq \sigma \leq \sigma_{\max}$, where $\sigma_{\min} = \eta_e^+(\mathbb{Z}^n)$ and $\sigma_{\max} = \sigma_{\min} \cdot \frac{\max_i \|\tilde{\mathbf{b}}_i\|}{\min_i \|\tilde{\mathbf{b}}_i\|}$. The $\tilde{\mathbf{b}}_i$ are the Gram-Schmidt vectors of the secret, short basis \mathbf{B} . In Falcon, it holds that:

$$\max_i \|\tilde{\mathbf{b}}_i\| \leq 1.17\sqrt{q} \quad (18)$$

$$\min_i \|\tilde{\mathbf{b}}_i\| \geq \sqrt{q}/1.17 \quad (19)$$

By construction, (18) is true (Falcon enforces it at key generation). To prove (19), we rely on a peculiar property of Falcon's private bases: *symplecticity*. Let f, g, F, G be such that $fG - gF = q$. Let:

$$\mathbf{J} = \begin{bmatrix} 0 & 1 \\ -1 & 0 \end{bmatrix} \quad \text{and} \quad \mathbf{B} = \begin{bmatrix} g & -f \\ F & -G \end{bmatrix}.$$

This form of \mathbf{B} is indeed the one used in Falcon. It has been observed in [24] that \mathbf{B} is q -symplectic, that is, it verifies:

$$\mathbf{B}^t \times \mathbf{J} \times \mathbf{B} = q \cdot \mathbf{J}. \quad (20)$$

As per [24, Corollary 1], this implies that for any i , $\|\tilde{\mathbf{b}}_{2n+1-i}\| = q/\|\tilde{\mathbf{b}}_i\|$. Combining this with (18) yields (19). Thus $\frac{\sigma_{\min}}{\sigma_{\max}} \leq (1.17)^{-2} \approx 0.73$, which means a non-negligible but reasonable impact on the running time of the sampler.

F Additional information on SAGA

F.1 Univariate tests

An example of the standard outputs of the univariate tests in SAGA is shown in Listing 1.1.

```
>> python3                                     # initialize python
>> import saga as test                         # import the test suite
>> a = test.UnivariateSamples(0, 1.5, data)    # run tests on data
>> a                                           # print the results

Testing a Gaussian sampler with center = 0 and sigma = 1.5
Number of samples: 10000

Moments | Expected      Empiric
-----+-----
Mean:   | 0.00000      -0.01260
St. dev. | 1.50000      1.49514
Skewness | 0.00000      -0.03414
Kurtosis | 0.00000      -0.00494

Chi-2 statistic: 5.98872465536634
Chi-2 p-value: 0.8162096584746129 (should be > 0.001)

How many outliers? 0

Is the sample valid? True
```

Listing 1.1: Output statistics on univariate Gaussian samples.

F.2 Multivariate tests

An example of the standard outputs of the multivariate tests in SAGA is shown in Listing 1.2.

```

>> python3                                     # initialize python environment
>> import saga as test                         # import the test suite
>> sk,u = test.test_sig(256, 1000)           # generate your multivariate data
>> u                                           # print the results

Testing a centered multivariate Gaussian of dimension = 2048 and sigma = 163.072
Number of samples: 10000

The test checks that the data corresponds to a multivariate Gaussian, by doing the
following:
1 - Print the covariance matrix (visual check). One can also plot
the covariance matrix by using self.show_covariance().
2 - Perform the Doornik-Hansen test of multivariate normality.
The p-value obtained should be > 0.001
3 - Perform a custom test called covariance diagonals test.
4 - Run a test of univariate normality on each coordinate

1 - Covariance matrix (2048 x 2048):
[[ 1.0115 -0.0268 -0.0112 ... 0.0102 -0.0043 0.004 ]
 [-0.0268 1.0067 -0.0057 ... 0.004 0.0063 0.0286]
 [-0.0112 -0.0057 1.0059 ... 0.0008 -0.0023 0.0052]
 ...
 [ 0.0102 0.004 0.0008 ... 0.986 0.0192 -0.0062]
 [-0.0043 0.0063 -0.0023 ... 0.0192 0.9905 0.0012]
 [ 0.004 0.0286 0.0052 ... -0.0062 0.0012 1.0129]]

2 - P-value of Doornik-Hansen test:           0.7064
3 - P-value of covariance diagonals test:     0.0497
4 - Gaussian coordinates (w/ st. dev. = sigma)? 2042 out of 2048

>> u.univariates                               # returns univariate tests on each 'row'

Testing a Gaussian sampler with center = 0 and sigma = 164.46976732471182
Number of samples: 1000

Moments | Expected Empiric
-----+-----
Mean:   | 0.00000 -2.01000
St. dev. | 164.46977 160.92979
Skewness | 0.00000 0.01555
Kurtosis | 0.00000 -0.00831

Chi-2 statistic: 97.98787878787878
Chi-2 p-value: 0.2650065102842649 (should be > 0.001)

How many outliers? 0

Is the sample valid? True

```

Listing 1.2: Output statistics on multivariate Gaussian samples.

F.3 Supplementary, visual, ‘sanity check’ tests

We also provide further sanity check functionality in the code, which will be particularly useful if any of the statistical tests above fail to conform to the expected values. D’Agostino et al. [12] suggested the best way to do this is via graphical methods, thus we provide visuals for both univariate and multivariate Gaussians.

For univariate Gaussians, we provide plots of the observed and expected probability density functions (Figure 2a) and quantile-quantile (QQ) plots (Figure 2b), shown in Figure 2. The QQ plot also provides the coefficient of determination, $R^2 \in [0, 1)$, which measures how well the observed data follows the distribution we expect. There are many different errors which can be

observed in QQ plots, from sampling biases to data skewness, but there are some useful guides [22], helpful websites^{8,9}, and online tools¹⁰ to help analysing these plots.

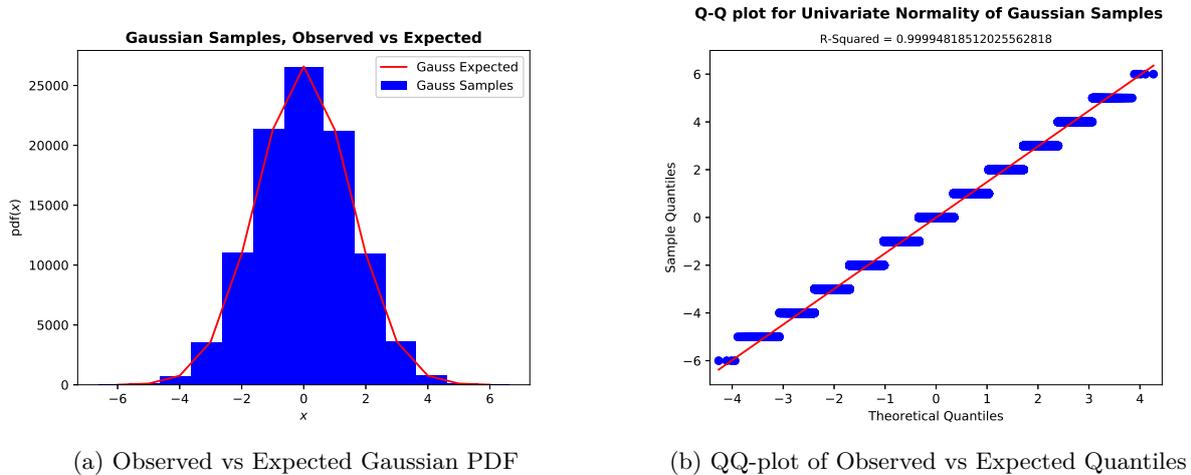


Figure 2: Visual representation of the expected form of univariate Gaussian samples.

For multivariate distributions, it is typically difficult to visualise any statistical properties of a distribution beyond three dimensions. However, there is a method for checking multivariate normality using a distribution’s Mahalanobis distance [37], which is a measure of the distance between certain points to a certain distribution. More specifically, it is a multi-dimensional generalisation of measuring how many standard deviations a point is away from a distribution. These sorted distances should follow a chi-square distribution [33], thus we can visualise this as a QQ-plot comparing the Mahalanobis distance versus the expected chi-square distribution. Figure 3 shows an example of our proposed multivariate normality graphical representation.

A final test we provide is for checking the rejection rate of the Gaussian samplers. Theoretically, the rejection rates of the Gaussian samples should decrease geometrically. An example of what the rejection rates should look like for a fixed σ are provided in Figure 4.

In order to check the rejection rates, the test suite requires the additional output of the rejection rate associated to each sample. An example of this is shown in Listings 1.3 and 1.4, which shows the typical Gaussian sampler (Listing 1.3) and the Gaussian sampler which also outputs the repetitions for each sample (Listing 1.4).

⁸ <https://stats.stackexchange.com/questions/101274/how-to-interpret-a-qq-plot>.

⁹ <https://data.library.virginia.edu/understanding-q-q-plots/>.

¹⁰ <https://xiongge.shinyapps.io/qqplots/>

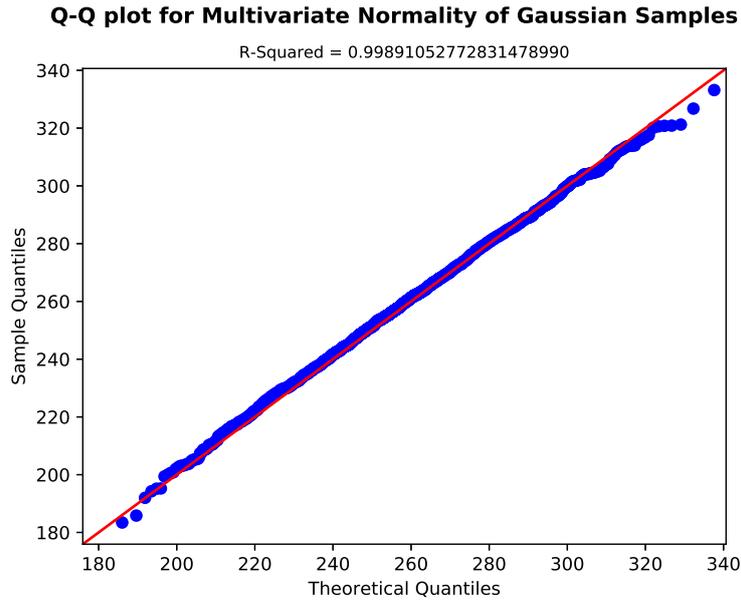


Figure 3: Visual representation of the expected form of univariate Gaussian samples.

```
def samplerz(center, sigma):
    assert(sigma < sigma0)
    assert(sigma >= sigmin)
    c0 = center - floor(center)
    sf = sigma / sigma0

    while(1):
        z0 = sampler0()
        b = randint(0, 1)
        z = ((b << 1) - 1) * z0 + b
        x = ((z - c0) ** 2) / (2 * (sigma ** 2)) -
            (z0 ** 2) / (2 * (sigma0 ** 2))

        if berexp(x, sf) is True:
            return floor(center) + z
```

Listing 1.3: Standard Gaussian sampler.

```
def samplerz_rep(center, sigma):
    assert(sigma < sigma0)
    assert(sigma >= sigmin)
    c0 = center - floor(center)
    sf = sigma / sigma0
    cnt = 0 # repetition counter
    while(1):
        z0 = sampler0()
        b = randint(0, 1)
        z = ((b << 1) - 1) * z0 + b
        x = ((z - c0) ** 2) / (2 * (sigma ** 2)) -
            (z0 ** 2) / (2 * (sigma0 ** 2))
        cnt += 1 # increment counter
        if berexp(x, sf) is True:
            # output count with value
            return floor(center) + z, cnt
```

Listing 1.4: Gaussian sampler with repetitions.

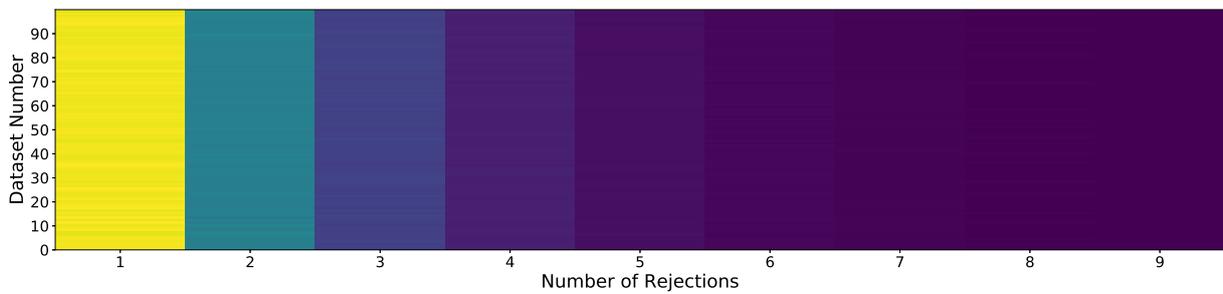


Figure 4: Expected geometric decrease in rejection numbers during Gaussian sampling.