COSAC: COmpact and Scalable Arbitrary-Centered Discrete Gaussian Sampling over Integers

Raymond K. Zhao, Ron Steinfeld, and Amin Sakzad

Faculty of Information Technology, Monash University,
{raymond.zhao,ron.steinfeld,amin.sakzad}@monash.edu

Abstract. The arbitrary-centered discrete Gaussian sampler is a fundamental subroutine in implementing lattice trapdoor sampling algorithms. However, existing approaches typically rely on either a fast implementation of another discrete Gaussian sampler or pre-computations with regards to some specific discrete Gaussian distributions with fixed centers and standard deviations. These approaches may only support sampling from standard deviations within a limited range, or cannot efficiently sample from arbitrary standard deviations determined on-the-fly at run-time.

In this paper, we propose a compact and scalable rejection sampling algorithm by sampling from a continuous normal distribution and performing rejection sampling on rounded samples. Our scheme does not require pre-computations related to any specific discrete Gaussian distributions. Our scheme can sample from both arbitrary centers and arbitrary standard deviations determined on-the-fly at run-time. In addition, we show that our scheme only requires a low number of trials close to 2 per sample on average, and our scheme maintains good performance when scaling up the standard deviation. We also provide a concrete error analysis of our scheme based on the Rényi divergence. We implement our sampler and analyse its performance in terms of storage and speed compared to previous results. Our sampler's running time is center-independent and is therefore applicable to implementation of convolution-style lattice trapdoor sampling and identity-based encryption resistant against timing side-channel attacks.

Keywords: Lattice-based crypto · Discrete Gaussian sampling · Implementation · Efficiency

1 Introduction

The arbitrary-centered discrete Gaussian sampling algorithm is an important subroutine in implementing lattice trapdoor samplers, which is a fundamental tool employed by lattice-based cryptography applications such as digital signature [20] and identity-based encryption (IBE) [3, 7]. However, previous works focused more on optimising the lattice trapdoor sampling algorithms, but the

The final authenticated publication is available online at
https://doi.org/10.1007/978-3-030-44223-1_16.

implementation details of the arbitrary-centered discrete Gaussian sampling were not well addressed. Typically, arbitrary-centered discrete Gaussian sampling approaches need to perform either rejection sampling [5, 8, 12, 20, 21] or pre-computations related to some specific discrete Gaussian distributions [14, 15, 17]. However, both types of methods have issues in the implementation: rejection sampling based methods are either slow due to the large number of trials per sample on average (typically, about 8–10) [8], requiring high precision arithmetic for cryptography applications [12], or relying on a fast implementation of another discrete Gaussian sampler [5, 20, 21]. On the other hand, pre-computation based methods consume at least few kilobytes (KB) of memory to store the tables and have the following limitations: the pre-computation table size in [14, 15] grows significantly when scaling up the standard deviation and this approach cannot support arbitrary standard deviations determined on-the-fly at run-time, while it is unclear how to efficiently implement the offline phase in [17] if the full algorithm needs to be executed during the run-time.

Recently the rounded Gaussian sampling (i.e. sampling from a continuous normal distribution and rounding the samples) was adapted by lattice-based digital signatures [11, 25]. Compared with a previous discrete Gaussian sampling algorithm [6], the rounded Gaussian sampler in [11] showed impressive performance with regards to the running speed and can be implemented in constant-time. The implementation in [11] is also notably simple (within less than 40 lines of C++ source code). However, since it is unclear whether a rounded Gaussian distribution can be directly adapted to implement a lattice trapdoor, another interesting question is: can one employ the existing efficient (rounded) continuous Gaussian distribution sampling techniques to implement an arbitrary-centered discrete Gaussian sampler?

1.1 Contribution

In this paper, we introduce a novel arbitrary-centered discrete Gaussian sampling algorithm over integers by generalising ideas from [4]. Our scheme samples from a continuous normal distribution and performs rejection sampling on rounded samples by adapting techniques from [11, 25]. Compared to previous arbitrary-centered discrete Gaussian sampling techniques, our scheme has the following advantages:

- Our sampling algorithm does not require any pre-computations related to a specific discrete Gaussian distribution or a specific standard deviation, and both the center and the standard deviation can be arbitrary determined on-the-fly at run-time.
- In addition, we show in Section 4 that our sampling method only requires a low number of trials close to 2 per sample on average compared to about 8–10 on average in the rejection sampling with regards to a uniform distribution, and the rejection rate of our algorithm decreases when scaling up σ. Therefore, our sampling algorithm is not limited to small σ and can be adapted to sample from larger σ without affecting the efficiency.

– Since sampling from a continuous normal distribution is a well-studied topic [22] and the sampling algorithms are implemented in many existing software libraries (including the C++11 STL) and hardware devices, one can easily implement our scheme by employing existing tools.
– We provide a center-independent run-time implementation of our algorithm without timing leakage of the center and it can be adapted to achieve timing resistant implementation of convolution-style lattice trapdoor sampler [16, 18] and IBE [3].

2 Preliminaries

Let $\rho_{c,\sigma}(x) = \exp\left(-(x-c)^2/\left(2\sigma^2\right)\right)$ be the (continuous) Gaussian function with center c and standard deviation σ. We denote the continuous Gaussian (normal) distribution with center c and standard deviation σ by $\mathcal{N}\left(c,\sigma^2\right)$, which has the probability density function $\rho_{c,\sigma}(x)/\left(\sigma\sqrt{2\pi}\right)$. We denote the discrete Gaussian distribution on integer lattices with center c and standard deviation σ by: $\mathcal{D}_{c,\sigma}(x) = \rho_{c,\sigma}(x)/S$, where $S = \rho_{c,\sigma}(\mathbb{Z}) = \sum_{k\in\mathbb{Z}} \rho_{c,\sigma}(k)$ is the normalisation factor. We omit the center in notations (i.e. $\rho_\sigma(x)$ and $\mathcal{D}_\sigma(x)$) if the center is zero. In addition, we denote the uniform distribution on set S by $\mathcal{U}(S)$. Sampling from a distribution $\mathcal{P}$ is denoted by $x \hookleftarrow \mathcal{P}$. We define $\lfloor x \rceil$ as the nearest integer to $x \in \mathbb{R}$. We denote $\mathbb{Z}^+$ as the integer set $\{1,\ldots,\infty\}$ and $\mathbb{Z}^-$ as the integer set $\{-\infty,\ldots,-1\}$, respectively. Also, for a lattice Λ and any $\epsilon \in \mathbb{R}^+$, we denote the smoothing parameter $\eta_\epsilon(\Lambda)$ as the smallest $s \in \mathbb{R}^+$ such that $\rho_{1/(s\sqrt{2\pi})}(\Lambda^* \setminus \{\mathbf{0}\}) \leq \epsilon$, where Λ^* is the dual lattice of Λ: $\Lambda^* = \{\boldsymbol{w} \in \mathbb{R}^n : \forall \boldsymbol{x} \in \Lambda, \boldsymbol{x}\cdot\boldsymbol{w} \in \mathbb{Z}\}$ [18]. An upper bound on $\eta_\epsilon(\mathbb{Z})$ is given by [18]: $\eta_\epsilon(\mathbb{Z}) \leq \sqrt{\ln(2+2/\epsilon)/\pi}$.

Theorem 1 (Adapted from [18], Lemma 2.4). *For any $\epsilon \in (0,1)$ and $c \in \mathbb{R}$, if $\sigma \geq \eta_\epsilon(\mathbb{Z})$, then $\rho_{c,\sigma}(\mathbb{Z}) = \left[\frac{1-\epsilon}{1+\epsilon}, 1\right] \cdot \rho_\sigma(\mathbb{Z})$, and $\rho_\sigma(\mathbb{Z})$ is approximately $\int_{-\infty}^{\infty} \rho_\sigma(x)\,\mathrm{d}x = \sigma\sqrt{2\pi}$.*

Definition 1 (Relative Error). *For two distributions $\mathcal{P}$ and $\mathcal{Q}$ such that $\mathrm{Supp}(\mathcal{P}) = \mathrm{Supp}(\mathcal{Q})$, the relative error between $\mathcal{P}$ and $\mathcal{Q}$ is defined as: $\Delta(\mathcal{P}||\mathcal{Q}) = \max_{x\in\mathrm{Supp}(\mathcal{P})} \frac{|\mathcal{P}(x)-\mathcal{Q}(x)|}{\mathcal{Q}(x)}$.*

Definition 2 (Rényi Divergence [2, 19]). *For two discrete distributions $\mathcal{P}$ and $\mathcal{Q}$ such that $\mathrm{Supp}(\mathcal{P}) \subseteq \mathrm{Supp}(\mathcal{Q})$, the Rényi divergence (RD) of order $\alpha \in (1,+\infty)$ is defined as: $R_\alpha(\mathcal{P}||\mathcal{Q}) = \left(\sum_{x\in\mathrm{Supp}(\mathcal{P})} \frac{\mathcal{P}(x)^\alpha}{\mathcal{Q}(x)^{\alpha-1}}\right)^{\frac{1}{\alpha-1}}$.*

Theorem 2 (Relative Error Bound, Adapted from [19], Lemma 3 and Eq. 4). *For two distributions $\mathcal{P}$ and $\mathcal{Q}$ such that $\mathrm{Supp}(\mathcal{P}) = \mathrm{Supp}(\mathcal{Q})$, we have: $R_\alpha(\mathcal{P}||\mathcal{Q}) \leq \left(1 + \frac{\alpha(\alpha-1)\cdot(\Delta(\mathcal{P}||\mathcal{Q}))^2}{2(1-\Delta(\mathcal{P}||\mathcal{Q}))^{\alpha+1}}\right)^{\frac{1}{\alpha-1}}$. The right-hand side is asymptotically*

equivalent to $1 + \alpha \cdot (\Delta(\mathcal{P}||\mathcal{Q}))^2/2$ *as* $\Delta(\mathcal{P}||\mathcal{Q}) \to 0$. *In addition, if a cryptographic search problem using* M *independent samples from* $\mathcal{Q}$ *is* $(\lambda+1)$*-bit secure, then the same problem sampling from* $\mathcal{P}$ *will be* λ*-bit secure if* $R_{2\lambda}(\mathcal{P}||\mathcal{Q}) \leq 1 + 1/(4M)$.

3 Previous Work

3.1 Rejection Sampling

The classic rejection sampling algorithm [8, 23] can sample from an arbitrary-centered discrete Gaussian distribution. To sample from $\mathcal{D}_{c,\sigma}$, one can sample $x \hookleftarrow \mathcal{U}([c-\tau\sigma, c+\tau\sigma] \cap \mathbb{Z})$ and accept x with probability $\rho_{c,\sigma}(x)$ as the output, where τ is the tail-cut factor (typically, about 10–12). However, this method is slow as the number of trials is $2\tau/\sqrt{2\pi}$ on average (about 8–10 for typical τ). Recently an algorithm sampling exactly from $\mathcal{D}_{c,\sigma}$ without floating-point arithmetic was presented by [12], which also has a lower rejection rate compared to the classic rejection sampling algorithm. However, this algorithm relies on high precision integer arithmetic to satisfy the precision requirements in cryptography applications.

To reduce the rejection rate, recent works performed rejection sampling with regards to some distributions much closer to $\mathcal{D}_{c,\sigma}$ compared to a uniform distribution: The Falcon signature [20] and its constant-time variant [21] adapted a rejection sampling method with regards to bimodal Gaussians: to sample from $\mathcal{D}_{c,\sigma}$ where $c \in [0,1]$, one can choose some $\sigma' \geq \sigma$ and sample $x \hookleftarrow \mathcal{D}^+_{\sigma'}$ (i.e. the discrete Gaussian distribution $\mathcal{D}_{\sigma'}$ restricted to the domain $\mathbb{Z}^+ \cup \{0\}$). The algorithm computes $x' = b + (2b-1) \cdot x$ where $b \hookleftarrow \mathcal{U}(\{0,1\})$. The authors of [20, 21] showed that x' has a bimodal Gaussian distribution close to the target distribution. The algorithm then accepts x' with probability $C(\sigma) \cdot \exp\left(\frac{x^2}{2\sigma'^2} - \frac{(x'-c)^2}{2\sigma^2}\right)$ as the output, where the scaling factor $C(\sigma) = \min(\sigma)/\sigma$ when sampling from multiple σ. This scheme has the average acceptance rate $C(\sigma) \cdot \rho_{c,\sigma}(\mathbb{Z}) / (2\rho_{\sigma'}(\mathbb{Z}^+))$, which is proportional to $\min(\sigma)/\sigma'$ [20, 21]. However, if the application needs to sample from different σ, the acceptance probability is high only when $\min(\sigma)$ and $\max(\sigma)$ are sufficiently close. This is not an issue in the Falcon signature, since the parameters in Falcon implies σ' is very close to $\max(\sigma)$ and $\min(\sigma)/\max(\sigma) \approx 0.73$ [21]. However, if the gap between $\min(\sigma)$ and $\max(\sigma)$ is large, since $\sigma' \geq \max(\sigma)$, this algorithm might have a low acceptance rate.[1]

A recent work [5] extended the binary sampling algorithm from the BLISS signature [6] to support non-zero arbitrary centers. For any center $c \in \mathbb{R}$, sampling from $\mathcal{D}_{c,\sigma}$ is equivalent to sampling from $\mathcal{D}_{c_F,\sigma} + \lfloor c \rfloor$, where $c_F = c - \lfloor c \rfloor \in [0,1)$ is the fractional part of c. In addition, for $c_F \in [1/2, 1)$, sampling from $\mathcal{D}_{c_F,\sigma}$ is

[1] One may employ different implementations for different σ, similar to the implementation of Falcon.

equivalent to sampling from $1 - \mathcal{D}_{c'_F,\sigma}$ where $c'_F = 1 - c_F \in (0, 1/2]$. A modified binary sampling scheme [5] can then be adapted to sample from $\mathcal{D}_{c'_F,\sigma}$ with any $c'_F \in (0, 1/2]$, in which the average number of trials is upper-bounded by: $\frac{\sigma^2}{\sigma_0\sigma - \sigma_0^2} \cdot \frac{\rho_{\sigma_0}(\mathbb{Z}^+)}{\sigma\sqrt{\pi/2} - 1}$, where $\sigma_0 = \sqrt{1/(2\ln 2)}$ is a fixed parameter used by the binary sampling algorithm [5, 6] and $\sigma = k\sigma_0$ for some $k \in \mathbb{Z}^+$. This upper-bound is about 1.47 for large σ [5].

3.2 TwinCDT

The authors of [14, 15] suggested a variant of the Cumulative Distribution Table (CDT) method [4] with multiple pre-computed tables. These algorithms will have two phases: online and offline. To be more specific, for $c \in [0, 1)$, during the offline phase, the algorithm pre-computes multiple CDT of $\mathcal{D}_{i/n,\sigma}$, where $i \in \{0, \ldots, n-1\}$ and $n \in \mathbb{Z}^+$ is sufficiently large. During the online phase, the algorithm picks a sample generated from either $\mathcal{D}_{\lfloor n(c-\lfloor c \rfloor)\rfloor/n,\sigma}$ or $\mathcal{D}_{\lceil n(c-\lfloor c \rfloor)\rceil/n,\sigma}$ as the output. Although the algorithm is very fast compared to other approaches, however, σ is fixed during the offline computation and thus this algorithm cannot support sampling from $\mathcal{D}_{c,\sigma}$ with both arbitrary c and σ determined on-the-fly at run-time. Another issue is that the pre-computation table size grows significantly when scaling up σ (see Table 2 in Section 5) and therefore the algorithm is not scalable.

3.3 Convolution

A recursive convolution sampling scheme for $\mathcal{D}_{c,\sigma}$ was presented in [17] as follows: suppose the center c has k fractional bits. Let $\sigma_0 = \sigma/\sqrt{\sum_{i=0}^{k-1} 2^{-2i}}$. One can sample $x_k \hookleftarrow \mathcal{D}_{c_k,\sigma_0}$ where $c_k = 2^{k-1} \cdot c$, then use $y_k = 2^{-k+1} \cdot x_k$ to round c to a new center $c' = c - y_k$ with $k' = k - 1$ fractional bits. Set $c = c'$ and $k = k'$ in the next iteration until $k = 0$, and $\sum_{i=1}^{k} y_i$ will be a sample distributed as $\mathcal{D}_{c,\sigma}$. The authors of [17] separated this algorithm into an online phase and an offline phase, where the offline phase will generate samples x_i in batch and the online phase will compute the linear combinations of x_i for $i \in \{1, \ldots, k\}$. The online phase is very fast and can be implemented in constant-time. However, for implementations where both sampling from $\mathcal{D}_{c_i,\sigma_0}$ and computing the linear combinations need to be carried during the run-time, it is unclear how to efficiently implement the $\mathcal{D}_{c_i,\sigma_0}$ sampling algorithm in constant-time (which is another discrete Gaussian sampler supporting a small amount of centers c_i). The offline batch sampler also consumes significant amount of memory (see Table 2 in Section 5).

4 Proposed Algorithm

In the textbook [4], the author defined a variant of the discrete Gaussian distribution as $\Pr[X = z] = c \cdot \exp\left(-(|z| + 1/2)^2 / (2\sigma^2)\right)$, where $z \in \mathbb{Z}$ and

Algorithm 1 Rejection sampler adapted from [4], pg. 117, ch. 3

Input: Standard deviation $\sigma \in \mathbb{R}^+$.
Output: A sample z distributed as $\Pr[X = z] = c \cdot \exp\left(-\left(|z| + 1/2\right)^2 / \left(2\sigma^2\right)\right)$.

1: **function** SAMPLER(σ)
2: Sample $x \hookleftarrow \mathcal{N}\left(0, \sigma^2\right)$.
3: Sample $r \hookleftarrow \mathcal{U}\left([0, 1)\right)$.
4: Let $Y = \left(|\lfloor x \rceil| + 1/2\right)^2 - x^2$.
5: **if** $r < \exp\left(-Y / \left(2\sigma^2\right)\right)$ **then**
6: Let $z = \lfloor x \rceil$.
7: **else**
8: **goto** 2.
9: **end if**
10: **return** z.
11: **end function**

c is the normalisation constant, i.e. $\Pr[X = z] \propto \rho_{-1/2,\sigma}(z)$ for $z \geq 0$ and $\Pr[X = z] \propto \rho_{1/2,\sigma}(z)$ for $z < 0$. A rejection sampling algorithm (see Algorithm 1) was provided by [4] with rejection probability less than $(2/\sigma) \cdot \sqrt{2/\pi}$ for such a distribution, which is fast for large σ (see Appendix B for the proof).

Here we generalise Algorithm 1 to sample from $\mathcal{D}_{c,\sigma}(z)$. By removing the absolute value and replacing the fixed center $-1/2$ with a generic center c in Algorithm 1, we observe that $Y' = \left(\lfloor x \rceil + c\right)^2 - x^2 \geq 0$ when $(c \geq 1/2, x \geq 0)$ or $(c \leq -1/2, x < 0)$. Therefore, we can replace Y with Y' and perform a similar rejection sampling to Algorithm 1 when sampling from $\mathcal{D}_{c,\sigma}(z)$ for some c and $z = \lfloor x \rceil$. To extend Algorithm 1 to support all $c \in \mathbb{R}$ and $z \in \mathbb{Z}$, we first compute $c_I = \lfloor c \rceil$ and $c_F = c_I - c$, where $c_F \in [-1/2, 1/2]$. Then we can sample from $\mathcal{D}_{-c_F,\sigma}$ instead, since $\mathcal{D}_{c,\sigma} = \mathcal{D}_{-c_F,\sigma} + c_I$. To sample from $\mathcal{D}_{-c_F,\sigma}$ for all $c_F \in [-1/2, 1/2]$, we shift the center of the underlying continuous normal distribution, i.e. sampling $y \hookleftarrow \mathcal{N}\left(\pm 1, \sigma^2\right)$, and perform a rejection sampling over $z = \lfloor y \rceil$ with acceptance rate $\exp\left(-Y'' / \left(2\sigma^2\right)\right)$ where $Y'' = \left(\lfloor y \rceil + c_F\right)^2 - \left(y \mp 1\right)^2$ (we also need to ensure $Y'' \geq 0$ before performing this rejection sampling). The sampling algorithm for $\mathcal{D}_{-c_F,\sigma}$ is presented in Algorithm 2. Note that the output of Algorithm 2 is restricted to the domain $\mathbb{Z} \setminus \{0\}$. Therefore, the algorithm needs to output 0 with probability $\mathcal{D}_{-c_F,\sigma}(0)$. We present the full algorithm in Algorithm 3. Since both Algorithm 2 and Algorithm 3 do not require pre-computations related to σ, our scheme can support arbitrary standard deviations determined on-the-fly at run-time in addition to arbitrary centers.

Theorem 3. *The output z sampled by Algorithm 2 is distributed as $\mathcal{D}_{-c_F,\sigma}\left(\mathbb{Z} \setminus \{0\}\right)$. The output of Algorithm 3 is distributed as $\mathcal{D}_{c,\sigma}(\mathbb{Z})$.*

Proof. When $b = 0$, y is distributed as $\mathcal{N}\left(-1, \sigma^2\right)$. For step 11 in Algorithm 2, we have $Y_1 = \left(\lfloor y \rceil + c_F\right)^2 - \left(y + 1\right)^2 \geq 0$ for any $c_F \in [-1/2, 1/2]$ when $y \leq -1/2$. Therefore, the rejection condition $\exp\left(-Y_1 / \left(2\sigma^2\right)\right) \in (0, 1]$. Let $z_0 = \lfloor y \rceil$. We

Algorithm 2 $\mathcal{D}_{-c_F,\sigma}\left(\mathbb{Z}\setminus\{0\}\right)$ sampler

Input: Center $c_F \in [-1/2, 1/2]$. Standard deviation $\sigma \in \mathbb{R}^+$.
Output: A sample z distributed as $\mathcal{D}_{-c_F,\sigma}$ restricted to the domain $\mathbb{Z}\setminus\{0\}$.

1: **function** RoundingSampler(c_F, σ)
2: Sample $x \hookleftarrow \mathcal{N}(0,1)$.
3: Sample $b \hookleftarrow \mathcal{U}(\{0,1\})$.
4: **if** $b = 0$ **then**
5: Let $y = \sigma \cdot x - 1$.
6: **if** $y > -1/2$ **then**
7: **goto** 2.
8: **end if**
9: Sample $r \hookleftarrow \mathcal{U}([0,1))$.
10: Let $Y_1 = (\lfloor y \rceil + c_F)^2 - (y+1)^2$.
11: **if** $r < \exp\left(-Y_1/\left(2\sigma^2\right)\right)$ **then**
12: Let $z = \lfloor y \rceil$.
13: **else**
14: **goto** 2.
15: **end if**
16: **else**
17: Let $y = \sigma \cdot x + 1$.
18: **if** $y < 1/2$ **then**
19: **goto** 2.
20: **end if**
21: Sample $r \hookleftarrow \mathcal{U}([0,1))$.
22: Let $Y_2 = (\lfloor y \rceil + c_F)^2 - (y-1)^2$.
23: **if** $r < \exp\left(-Y_2/\left(2\sigma^2\right)\right)$ **then**
24: Let $z = \lfloor y \rceil$.
25: **else**
26: **goto** 2.
27: **end if**
28: **end if**
29: **return** z.
30: **end function**

have the output distribution:

$$\begin{aligned}\Pr\left[z = z_0\right] &\propto \int_{z_0-1/2}^{z_0+1/2} \exp\left(-\frac{(y+1)^2}{2\sigma^2}\right) \cdot \exp\left(-\frac{(z_0+c_F)^2 - (y+1)^2}{2\sigma^2}\right) \mathrm{d}y \\ &= \int_{z_0-1/2}^{z_0+1/2} \exp\left(-\frac{(z_0+c_F)^2}{2\sigma^2}\right) \mathrm{d}y = \rho_{-c_F,\sigma}\left(z_0\right). \end{aligned} \tag{1}$$

In this case, the distribution of $z = z_0$ is $\mathcal{D}_{-c_F,\sigma}$ restricted to the domain $\mathbb{Z}^-$ (due to the rejection of y to $(-\infty, -1/2]$).

Similarly, when $b = 1$, y is distributed as $\mathcal{N}\left(1, \sigma^2\right)$. For step 23 in Algorithm 2, we have $Y_2 = (\lfloor y \rceil + c_F)^2 - (y-1)^2 \geq 0$ for any $c_F \in [-1/2, 1/2]$ when $y \geq 1/2$. Therefore, the rejection condition $\exp\left(-Y_2/\left(2\sigma^2\right)\right) \in (0,1]$. Let $z_0 = \lfloor y \rceil$.

Algorithm 3 $\mathcal{D}_{c,\sigma}(\mathbb{Z})$ sampler

Input: Center $c \in \mathbb{R}$. Standard deviation $\sigma \in \mathbb{R}^+$. Normalisation factor $S = \rho_{c,\sigma}(\mathbb{Z}) \approx \sigma\sqrt{2\pi}$.
Output: A sample distributed as $\mathcal{D}_{c,\sigma}(\mathbb{Z})$.

function RoundingSamplerFull(c, σ)
  Let $c_I = \lfloor c \rceil$ and $c_F = c_I - c$.
  Sample $r \hookleftarrow \mathcal{U}([0,1))$.
  if $r < \exp\left(-c_F^2/\left(2\sigma^2\right)\right)/S$ **then**
    Let $z' = 0$.
  else
    Let $z' =$ RoundingSampler(c_F, σ).
  end if
  return $z' + c_I$.
end function

We have the output distribution:

$$\begin{aligned}\Pr[z = z_0] &\propto \int_{z_0-1/2}^{z_0+1/2} \exp\left(-\frac{(y-1)^2}{2\sigma^2}\right) \cdot \exp\left(-\frac{(z_0+c_F)^2 - (y-1)^2}{2\sigma^2}\right) \mathrm{d}y \\ &= \int_{z_0-1/2}^{z_0+1/2} \exp\left(-\frac{(z_0+c_F)^2}{2\sigma^2}\right) \mathrm{d}y = \rho_{-c_F,\sigma}(z_0). \end{aligned} \tag{2}$$

In this case, the distribution of $z = z_0$ is $\mathcal{D}_{-c_F,\sigma}$ restricted to the domain $\mathbb{Z}^+$ (due to the rejection of y to $[1/2, \infty)$). Therefore, the output z in Algorithm 2 is distributed as $\mathcal{D}_{-c_F,\sigma}$ restricted to the domain $\mathbb{Z} \setminus \{0\}$.

In Algorithm 3, the probability $\Pr[z' = 0] = \exp\left(-c_F^2/\left(2\sigma^2\right)\right)/S = \mathcal{D}_{-c_F,\sigma}(0)$. Therefore, variable z' is distributed as $\mathcal{D}_{-c_F,\sigma}(\mathbb{Z})$. Since $c = c_I - c_F$, we have the output $z' + c_I$ distributed as $\mathcal{D}_{c,\sigma}(\mathbb{Z})$.

□

To prove the rejection rate of Algorithm 2, we need the following lemma:

Lemma 1. *For any $\epsilon \in (0,1)$ and $c \in [-1/2, 1/2]$, if $\sigma \geq \eta_\epsilon(\mathbb{Z})$, then both $\rho_{c,\sigma}(\mathbb{Z}^-)$ and $\rho_{c,\sigma}(\mathbb{Z}^+)$ have the lower bound: $\frac{1}{2} \cdot \frac{1-\epsilon}{1+\epsilon} \cdot \rho_\sigma(\mathbb{Z}) - 1$.*

Proof. When $c \in [-1/2, 1/2]$, for $\rho_{c,\sigma}(\mathbb{Z}^-)$, we have:

$$\rho_{c,\sigma}(\mathbb{Z}) = \rho_{c,\sigma}\left(\mathbb{Z}^+\right) + \rho_{c,\sigma}\left(\mathbb{Z}^- \cup \{0\}\right) \leq 2\rho_{c,\sigma}\left(\mathbb{Z}^- \cup \{0\}\right) = 2\rho_{c,\sigma}\left(\mathbb{Z}^-\right) + 2\rho_{c,\sigma}(0).$$

Therefore,

$$\begin{aligned}\rho_{c,\sigma}\left(\mathbb{Z}^-\right) &\geq \frac{1}{2} \cdot \rho_{c,\sigma}(\mathbb{Z}) - \rho_{c,\sigma}(0) \\ &\geq \frac{1}{2} \cdot \frac{1-\epsilon}{1+\epsilon} \cdot \rho_\sigma(\mathbb{Z}) - \rho_{c,\sigma}(0) \quad \text{(By Theorem 1).}\end{aligned}$$

We have $\rho_\sigma(0) \geq \rho_{c,\sigma}(0)$ for $c \in [-1/2, 1/2]$. Therefore,

$$\rho_{c,\sigma}\left(\mathbb{Z}^-\right) \geq \frac{1}{2} \cdot \frac{1-\epsilon}{1+\epsilon} \cdot \rho_\sigma(\mathbb{Z}) - 1.$$

Similarly, when $c \in [-1/2, 1/2]$, for $\rho_{c,\sigma}(\mathbb{Z}^+)$, we have:

$$\rho_{c,\sigma}(\mathbb{Z}) = \rho_{c,\sigma}\left(\mathbb{Z}^-\right) + \rho_{c,\sigma}\left(\mathbb{Z}^+ \cup \{0\}\right) \leq 2\rho_{c,\sigma}\left(\mathbb{Z}^+ \cup \{0\}\right) = 2\rho_{c,\sigma}\left(\mathbb{Z}^+\right) + 2\rho_{c,\sigma}(0).$$

Therefore, since $c \in [-1/2, 1/2]$, we have:

$$\begin{aligned}
\rho_{c,\sigma}\left(\mathbb{Z}^+\right) &\geq \frac{1}{2} \cdot \rho_{c,\sigma}(\mathbb{Z}) - \rho_{c,\sigma}(0) \\
&\geq \frac{1}{2} \cdot \frac{1-\epsilon}{1+\epsilon} \cdot \rho_\sigma(\mathbb{Z}) - \rho_{c,\sigma}(0) \quad \text{(By Theorem 1)} \\
&\geq \frac{1}{2} \cdot \frac{1-\epsilon}{1+\epsilon} \cdot \rho_\sigma(\mathbb{Z}) - 1 \quad (\rho_\sigma(0) \geq \rho_{c,\sigma}(0) \text{ when } c \in [-1/2, 1/2]).
\end{aligned}$$

□

Theorem 4. *For $\sigma \geq \eta_\epsilon(\mathbb{Z})$, the expected number of trials M in Algorithm 2 has the upper bound: $M \leq 2 \cdot \frac{1+\epsilon}{1-\epsilon} \cdot \frac{\sigma\sqrt{2\pi}}{\sigma\sqrt{2\pi} - 1 - 2 \cdot \frac{1+\epsilon}{1-\epsilon}}$. If σ is much greater than $\left(1 + 2 \cdot \frac{1+\epsilon}{1-\epsilon}\right)/\sqrt{2\pi}$, then $M \leq 2 \cdot (1 + O(\epsilon) + O(1/\sigma))$.*

Proof. By Theorem 3, when $b = 0$, we have the output probability density function $f(y) = \rho_{-c_F,\sigma}(\lfloor y \rceil)/\rho_{-c_F,\sigma}(\mathbb{Z}^-)$ and the input probability density function $g(y) = \rho_{-1,\sigma}(y)/\left(\sigma\sqrt{2\pi}\right)$. The expected number of trials can be written as:

$$M = \max \frac{f(y)}{g(y)} = \max \left(\frac{\rho_{-c_F,\sigma}(\lfloor y \rceil)}{\rho_{-1,\sigma}(y)} \cdot \frac{\sigma\sqrt{2\pi}}{\rho_{-c_F,\sigma}(\mathbb{Z}^-)} \right).$$

We have:

$$\frac{\rho_{-c_F,\sigma}(\lfloor y \rceil)}{\rho_{-1,\sigma}(y)} = \frac{\exp\left(-\frac{(\lfloor y \rceil + c_F)^2}{2\sigma^2}\right)}{\exp\left(-\frac{(y+1)^2}{2\sigma^2}\right)} = \exp\left(-\frac{(\lfloor y \rceil + c_F)^2 - (y+1)^2}{2\sigma^2}\right) \leq 1.$$

Therefore,

$$M \leq \frac{\sigma\sqrt{2\pi}}{\rho_{-c_F,\sigma}(\mathbb{Z}^-)} \leq 2 \cdot \frac{1+\epsilon}{1-\epsilon} \cdot \frac{\sigma\sqrt{2\pi}}{\rho_\sigma(\mathbb{Z}) - 2 \cdot \frac{1+\epsilon}{1-\epsilon}} \leq 2 \cdot \frac{1+\epsilon}{1-\epsilon} \cdot \frac{\sigma\sqrt{2\pi}}{\sigma\sqrt{2\pi} - 1 - 2 \cdot \frac{1+\epsilon}{1-\epsilon}},$$

where the second inequality follows from Lemma 1, and the third inequality follows from $\rho_\sigma(\mathbb{Z}) = \rho_\sigma(\mathbb{Z}^- \cup \{0\}) + \rho_\sigma(\mathbb{Z}^+ \cup \{0\}) - 1$ and the sum-integral comparison: $\rho_\sigma(\mathbb{Z}^- \cup \{0\}) \geq \int_{-\infty}^{0} \rho_\sigma(x)\,\mathrm{d}x = \sigma\sqrt{\pi/2}$ and $\rho_\sigma(\mathbb{Z}^+ \cup \{0\}) \geq \int_0^\infty \rho_\sigma(x)\,\mathrm{d}x = \sigma\sqrt{\pi/2}$.

Similarly, when $b = 1$, we have the output probability density function $f(y) = \rho_{-c_F,\sigma}(\lfloor y \rceil)/\rho_{-c_F,\sigma}(\mathbb{Z}^+)$ and the input probability density function $g(y) = \rho_{1,\sigma}(y)/\left(\sigma\sqrt{2\pi}\right)$. The expected number of trials can be written as:

$$M = \max \frac{f(y)}{g(y)} = \max \left(\frac{\rho_{-c_F,\sigma}(\lfloor y \rceil)}{\rho_{1,\sigma}(y)} \cdot \frac{\sigma\sqrt{2\pi}}{\rho_{-c_F,\sigma}(\mathbb{Z}^+)} \right).$$

We have:

$$\frac{\rho_{-c_F,\sigma}(\lfloor y \rceil)}{\rho_{1,\sigma}(y)} = \frac{\exp\left(-\frac{(\lfloor y \rceil + c_F)^2}{2\sigma^2}\right)}{\exp\left(-\frac{(y-1)^2}{2\sigma^2}\right)} = \exp\left(-\frac{(\lfloor y \rceil + c_F)^2 - (y-1)^2}{2\sigma^2}\right) \leq 1.$$

Therefore,

$$M \leq \frac{\sigma\sqrt{2\pi}}{\rho_{-c_F,\sigma}(\mathbb{Z}^+)} \leq 2 \cdot \frac{1+\epsilon}{1-\epsilon} \cdot \frac{\sigma\sqrt{2\pi}}{\rho_\sigma(\mathbb{Z}) - 2 \cdot \frac{1+\epsilon}{1-\epsilon}} \leq 2 \cdot \frac{1+\epsilon}{1-\epsilon} \cdot \frac{\sigma\sqrt{2\pi}}{\sigma\sqrt{2\pi} - 1 - 2 \cdot \frac{1+\epsilon}{1-\epsilon}},$$

where the second inequality follows from Lemma 1, and the third inequality follows from $\rho_\sigma(\mathbb{Z}) \geq \sigma\sqrt{2\pi} - 1$.

When σ is much greater than $\left(1 + 2 \cdot \frac{1+\epsilon}{1-\epsilon}\right)/\sqrt{2\pi}$, $\sigma\sqrt{2\pi}$ is much greater than $1 + 2 \cdot \frac{1+\epsilon}{1-\epsilon}$. Thus,

$$M \leq 2 \cdot \frac{1+\epsilon}{1-\epsilon} \cdot \frac{\sigma\sqrt{2\pi}}{\sigma\sqrt{2\pi} - 1 - 2 \cdot \frac{1+\epsilon}{1-\epsilon}} \leq 2 \cdot (1 + O(\epsilon) + O(1/\sigma)).$$

□

4.1 Accuracy Analysis

We now analyse the relative error of Algorithm 2 here. Let the absolute error of the continuous Gaussian sample x be e_x: $x' = x + e$, where x' is the actual sample, x is the ideal sample, and the error $|e| \leq e_x$. We denote the actual distribution by $\mathcal{P}_{\text{actual}}$ and the ideal distribution by $\mathcal{P}_{\text{ideal}}$. Since the variable y might be rounded to an incorrect integer due to the error from x when y is close to the boundaries $z_0 \pm 1/2$ [11], we have:

$$\begin{aligned}
&\Delta(\mathcal{P}_{\text{actual}} || \mathcal{P}_{\text{ideal}}) = \max \left| \frac{\mathcal{P}_{\text{actual}}}{\mathcal{P}_{\text{ideal}}} - 1 \right| \\
&= \max_{z_0} \left| \frac{\int_{z_0-1/2-\sigma e_x}^{z_0+1/2+\sigma e_x} \exp\left(-\frac{(z_0+c_F)^2}{2\sigma^2}\right) \mathrm{d}y}{\rho_{-c_F,\sigma}(z_0)} - 1 \right| \quad \text{(by (1), (2), and } y = \sigma x \pm 1) \\
&= \max_{z_0} \left| \frac{(1 + 2\sigma e_x) \cdot \rho_{-c_F,\sigma}(z_0)}{\rho_{-c_F,\sigma}(z_0)} - 1 \right| = 2\sigma e_x.
\end{aligned}$$

By Theorem 2, for λ-bit security, we need:

$$R_{2\lambda}\left(\mathcal{P}_{\text{actual}}||\mathcal{P}_{\text{ideal}}\right) \leq 1+\frac{1}{4M} \implies 1+2\lambda\cdot\frac{\left(\Delta\left(\mathcal{P}_{\text{actual}}||\mathcal{P}_{\text{ideal}}\right)\right)^2}{2} \leq 1+\frac{1}{4M}$$
$$\implies e_x \leq \frac{1}{4\sigma\sqrt{\lambda M}}.$$

Note that both $\mathcal{P}_{\text{actual}}$ and $\mathcal{P}_{\text{ideal}}$ have the same normalisation factor, since $\mathcal{P}_{\text{actual}}$ is obtained by the imperfect continuous Gaussian distribution with the rounding error contributed to the interval of the integral [11].

5 Evaluation

Side-channel Resistance Our implementation is not fully constant-time because the rejection rate may still reveal σ due to Theorem 4. However, since the rejection rate is independent of the center, our implementation can achieve fully constant-time with respect to the secret if σ is public. The σ in convolution-style lattice trapdoor samplers [16, 18] is typically a public constant, but σ in GPV-style sampler [10] depends on the secret. Note that the IBE implementation in [3] adapted a variant of [16], but it appears that the implementation source code[2] of [3] used a different distribution and the side-channel resistance perspective is unclear. Our sampling algorithm can be applied in the IBE implementation of [3] to give a fully constant-time IBE implementation.

We perform benchmarks of Algorithm 3 with fixed σ and random arbitrary centers. We employ the Box-Muller continuous Gaussian sampler [11, 25] implemented by using the VCL library [9], which provides $e_x \leq 2^{-48}$ [11]. To compare with [15], we select $\sigma = \{2, 4, 8, 16, 32\}$, and to compare with [17], we choose $\sigma = 2^{15}$. In addition, we also compare with variants [5, 26] of the binary sampling algorithm [6] for additional $\sigma = \left\{2^{17}, 2^{20}\right\}$. From the error analysis in Section 4.1, for given e_x and λ, $M \leq \frac{1}{16\lambda e_x^2\sigma^2}$. For $\sigma \in \left[2, 2^{20}\right]$ and $\lambda = 128$, we have $M \leq 2^{45}$. We adapt techniques similar to [26] to avoid high precision arithmetic (see Appendix A for details) and the scheme[3] is implemented by using the double precision i.e. $\delta_f = 52$. We also compute the normalisation factor S in double precision. We use the AES256 counter mode with hardware AES instructions (AES-NI) to generate the randomness in our implementations. We provide both the non-constant time reference implementation and the center-independent run-time implementation. We take care of all the branches for the center-independent run-time implementation by adapting constant-time selection techniques [1]. For the non-constant time reference implementation (the "Ref." column in Table 1), we use the $\exp(x)$ from the C library, which provides about 50-bit precision [20], while for the center-independent run-time implementation (the "Center-independent" column in Table 1), we adapt the techniques from [26] with about 45-bit precision. From the precision analysis in [19, 26], the

[2] https://github.com/lbibe/code

[3] Our implementation is available at https://github.com/raykzhao/gaussian_ac

Table 1. Number of Samples per Second for Our Scheme with Fixed σ at 4.2GHz (with $\lambda = 128$).

σ	Ref. ($\times 10^6$)	Center-independent ($\times 10^6$)
2	10.33 ± 0.18	8.96 ± 0.16
4	11.57 ± 0.18	10.87 ± 0.15
8	11.95 ± 0.17	11.61 ± 0.13
16	12.14 ± 0.16	12.00 ± 0.12
32	12.19 ± 0.15	12.21 ± 0.11
2^{15}	11.70 ± 0.13	11.57 ± 0.09
2^{17}	11.20 ± 0.14	11.63 ± 0.10
2^{20}	11.17 ± 0.13	11.28 ± 0.09

above precisions (including the precision of S) are sufficient for $\lambda = 128$ and $M \leq 2^{45}$.

The benchmark is carried on as follows: we use `g++` 9.1.1 to compile our implementations with the compiling options `-O3 -march=native` enabled. The benchmark is running on an Intel i7-7700K CPU at 4.2GHz, with the Hyper-threading and the Turbo Boost disabled. We generate 1024 samples (with a random arbitrary center per sample) for 1000 times and measure the consumed CPU cycles, with the exception that we fix $c = 0$ and compare our center-independent run-time implementation with [26], since the scheme in [26] is essentially a constant-time zero-centered discrete Gaussian sampler. Then we convert the CPU cycles to the average number of samples per second for the comparison purpose with previous works.

The benchmark results of our scheme are shown in Table 1 (in the format of mean $\pm$ standard deviation). We also summarise the performance of previous works in Table 2, and show the comparison with [26] in Table 3 when $c = 0$. Since previous works [5, 15, 17] measured the number of generated samples per second running on CPUs with different frequencies, we scale all the numbers to be based on 4.2GHz.[4] In addition, since some previous works [15, 17] require pre-computations to implement the sampling schemes, we summarise the pre-computation memory storage consumptions in Table 2. Because the TwinCDT method [15] provided different tradeoffs between the running speed and the pre-computation storage consumption, we show all 3 different sampling speeds and the corresponding pre-computation storage consumptions for each σ from [15]. Note that although our sampling scheme does not require pre-computations, however, the $\exp(x)$ implementation typically consumes a small amount of memory to store the coefficients of the polynomial approximation. For example, the polynomial approximation of the $\exp(x)$ in our center-independent run-time implementation (adapted from [26]) has degree 10 with double precision coefficients, and therefore it consumes $(10 + 1) \cdot 8 = 88$ bytes.

[4] The online+offline benchmark result is obtained and scaled from the variant implemented by [5].

Table 2. Summary of Previous Works for Fixed σ at 4.2GHz (with $\lambda = 128$).

σ	Num. of samples ($\times 10^6$/sec)	Pre-computation storage (KB)
2 [15]	51.01/62.45/76.43	1.4/4.6/46
4 [15]	45.50/56.44/69.09	1.9/6.3/63
8 [15]	37.70/53.31/63.51	3/10/100
16 [15]	31.29/37.63/52.29	5.2/17/172
32 [15]	34.38/39.76/42.60	9.5/32/318
2^{15} [17]	$\approx$ 12.35 (online)[5], 1.78 (online+offline)	$2^{5.4}$
4–2^{20} [5]	$\approx$ 16.3	–[6]

Table 3. Number of Samples per Second Compared with [26] for Fixed σ and $c = 0$ at 4.2GHz (with $\lambda = 128$).

σ	Our Scheme ($\times 10^6$/sec)	[26] ($\times 10^6$/sec)
2	9.44	19.87
4	11.10	19.04
8	12.08	19.04
16	12.63	18.62
32	12.93	18.80
2^{15}	12.67	18.36
2^{17}	12.67	18.90
2^{20}	13.04	18.70

From Table 1, our scheme has good performance for both small and large σ (11.53×10^6 samples per second for the non-constant time reference implementation and 11.27×10^6 samples per second for the center-independent run-time implementation on average). In particular, our scheme has better performance for large σ since the number of trials becomes lower by Theorem 4. Note that the amount of randomness required by the comparison steps in Appendix A will significantly increase for very small or very large σ. Therefore, our implementation consumes different amount of randomness in comparison steps for each σ based on Appendix A, and the performance for some larger σ is slightly slower than smaller σ in Table 1 due to the increased amount of randomness required. The overhead introduced by the center-independent run-time implementation is at most 13.33% in our benchmarks. Note that the overhead of the center-independent run-time implementation is smaller for large σ due to the lower probability of outputting $z' = 0$ in Algorithm 3.

For $\sigma \in [2, 32]$, although the TwinCDT method [15] is 2.5x–7.3x faster than our non-constant time reference implementation, however, this method requires a pre-computation with at least 1.4 KB memory consumption to store the CDT, while our scheme only requires at most several hundred bytes if considering all

[5] The result in [17] is based on the authors' reference implementation, which is not claimed to be optimal [24].

[6] The base sampler and the Bernoulli sampler may require pre-computations depending on the implementation techniques.

the polynomial approximation coefficients (including those functions used by the Box-Muller continuous Gaussian sampler). When scaling up σ, the TwinCDT method [15] also costs much larger amount of memory (the pre-computation storage size increases by a factor of 6.7–6.9 when σ changes from 2 to 32), and the performance becomes significantly worse (the number of samples per second decreases by 32.6–44.3% when σ changes from 2 to 32). In contrary, the pre-computation storage of our scheme is independent of σ and only relies on the precision requirements. Our scheme is also scalable and maintains good performance even for large $\sigma = 2^{15}$. In addition, for applications sampling from various σ such as [7], one sampler subroutine implemented by using our scheme is able to serve all σ since the implementation does not require any pre-computations depending on σ, while the TwinCDT method [15] needs to pre-compute a different CDT for each σ.

Compared with [17] for $\sigma = 2^{15}$, if we measure both the online and offline phase running speed in total, our center-independent run-time implementation achieves better performance in terms of both timing (6.5x faster) and pre-computation storage (the implementation in [17] requires about 42 KB to implement the Knuth-Yao [13] offline batch sampler).[7] The online-phase only running speed in [17] is slightly (1.07x) faster than our scheme. On the other hand, our scheme requires no offline pre-computations related to a specific discrete Gaussian distribution. In addition, our scheme can also be accelerated if we generate all the continuous Gaussian samples during the offline phase and only perform the rejection during the online phase. In this case, our center-independent run-time implementation generates 13.73×10^6 samples per second during the online phase, which is 1.11x faster than [17].

For the comparison with variants of the binary sampling algorithm, in Table 2, our non-constant time reference implementation is about 28.2% slower than [5] for $\sigma \in \left[4, 2^{20}\right]$ with arbitrary centers, and from Table 3, our center-independent run-time implementation is 30.3%–52.5% slower than [26] when $c = 0$ and $\sigma \in \left[2, 2^{20}\right]$. However, the scheme in [26] does not support an arbitrary center, while the side-channel resistance perspective of [5] is unclear. We expect that our implementation can achieve at most about 73.5% of the running speed of [5, 26] on average for large σ, since both binary sampling variants [5, 26] require less than 1.47 trials per sample on average, while the average number of trials per sample is close to 2 in our scheme for large σ.

6 Conclusion

In conclusion, we generalise the idea from [4] and present a compact and scalable arbitrary-centered discrete Gaussian sampling scheme over integers. Our scheme performs rejection sampling on rounded samples from a continuous normal distribution, which does not rely on any discrete Gaussian sampling implementations. We show that our scheme maintains good performance for $\sigma \in \left[2, 2^{20}\right]$ and needs

[7] Here we compare the performance with our center-independent run-time implementation because the implementation in [17] is constant-time.

no pre-computations related to any specific σ, which is suitable to implement applications that requires sampling from multiple different σ. In addition, we provide concrete rejection rate and error analysis of our scheme.

The performance of our scheme heavily relies on the underlying continuous Gaussian sampling algorithm. However, the Box-Muller sampler [11, 25] employed in our implementation does not have the fastest sampling speed compared to other algorithms according to a survey [22]. The main reason behind the choice of the continuous Gaussian sampler in our implementation is because the Box-Muller sampler is very simple to implement in constant-time [11]. If the side-channel perspective is not a concern, one may employ other more efficient non-constant time algorithms from the survey [22] to achieve a faster implementation of our scheme.

Acknowledgments. Ron Steinfeld was supported in part by ARC Discovery Project grant DP180102199.

References

1. Aumasson, J.P.: Guidelines for low-level cryptography software. https://github.com/veorq/cryptocoding (2019), accessed: 2020-01-28
2. Bai, S., Langlois, A., Lepoint, T., Stehlé, D., Steinfeld, R.: Improved security proofs in lattice-based cryptography: Using the Rényi divergence rather than the statistical distance. In: ASIACRYPT (1). Lecture Notes in Computer Science, vol. 9452, pp. 3–24. Springer (2015)
3. Bert, P., Fouque, P., Roux-Langlois, A., Sabt, M.: Practical implementation of ring-SIS/LWE based signature and IBE. In: PQCrypto. Lecture Notes in Computer Science, vol. 10786, pp. 271–291. Springer (2018)
4. Devroye, L.: Non-Uniform Random Variate Generation. Springer-Verlag, New York, NY, USA (1986)
5. Du, Y., Wei, B., Zhang, H.: A rejection sampling algorithm for off-centered discrete Gaussian distributions over the integers. SCIENCE CHINA Information Sciences **62**(3), 39103:1–39103:3 (2019)
6. Ducas, L., Durmus, A., Lepoint, T., Lyubashevsky, V.: Lattice signatures and bimodal Gaussians. In: CRYPTO (1). Lecture Notes in Computer Science, vol. 8042, pp. 40–56. Springer (2013)
7. Ducas, L., Lyubashevsky, V., Prest, T.: Efficient identity-based encryption over NTRU lattices. In: ASIACRYPT (2). Lecture Notes in Computer Science, vol. 8874, pp. 22–41. Springer (2014)
8. Ducas, L., Nguyen, P.Q.: Faster Gaussian lattice sampling using lazy floating-point arithmetic. In: ASIACRYPT. Lecture Notes in Computer Science, vol. 7658, pp. 415–432. Springer (2012)
9. Fog, A.: VCL C++ vector class library. www.agner.org/optimize/vectorclass.pdf, accessed: 2019-08-01
10. Gentry, C., Peikert, C., Vaikuntanathan, V.: Trapdoors for hard lattices and new cryptographic constructions. In: STOC. pp. 197–206. ACM (2008)
11. Hülsing, A., Lange, T., Smeets, K.: Rounded Gaussians - fast and secure constant-time sampling for lattice-based crypto. In: Public Key Cryptography (2). Lecture Notes in Computer Science, vol. 10770, pp. 728–757. Springer (2018)

12. Karney, C.F.F.: Sampling exactly from the normal distribution. ACM Trans. Math. Softw. **42**(1), 3:1–3:14 (2016)
13. Knuth, D., Yao, A.: Algorithms and Complexity: New Directions and Recent Results, chap. The complexity of nonuniform random number generation. Academic Press (1976)
14. Melchor, C.A., Albrecht, M.R., Ricosset, T.: Sampling from arbitrary centered discrete Gaussian for lattice-based cryptography. In: ACNS. Lecture Notes in Computer Science, vol. 10355, pp. 3–19. Springer (2017)
15. Melchor, C.A., Ricosset, T.: CDT-based Gaussian sampling: From multi to double precision. IEEE Trans. Computers **67**(11), 1610–1621 (2018)
16. Micciancio, D., Peikert, C.: Trapdoors for lattices: Simpler, tighter, faster, smaller. In: EUROCRYPT. Lecture Notes in Computer Science, vol. 7237, pp. 700–718. Springer (2012)
17. Micciancio, D., Walter, M.: Gaussian sampling over the integers: Efficient, generic, constant-time. In: CRYPTO (2). Lecture Notes in Computer Science, vol. 10402, pp. 455–485. Springer (2017)
18. Peikert, C.: An efficient and parallel Gaussian sampler for lattices. In: CRYPTO. Lecture Notes in Computer Science, vol. 6223, pp. 80–97. Springer (2010)
19. Prest, T.: Sharper bounds in lattice-based cryptography using the Rényi divergence. In: ASIACRYPT (1). Lecture Notes in Computer Science, vol. 10624, pp. 347–374. Springer (2017)
20. Prest, T., Fouque, P.A., Hoffstein, J., Kirchner, P., Lyubashevsky, V., Pornin, T., Ricosset, T., Seiler, G., Whyte, W., Zhang, Z.: Falcon: Fast-Fourier lattice-based compact signatures over NTRU. https://falcon-sign.info/ (2017), accessed: 2018-10-31
21. Prest, T., Ricosset, T., Rossi, M.: Simple, fast and constant-time Gaussian sampling over the integers for Falcon. Second PQC Standardization Conference, https://csrc.nist.gov/CSRC/media/Events/Second-PQC-Standardization-Conference/documents/accepted-papers/rossi-simple-fast-constant.pdf (2019), accessed: 2019-08-13
22. Thomas, D.B., Luk, W., Leong, P.H.W., Villasenor, J.D.: Gaussian random number generators. ACM Comput. Surv. **39**(4), 11 (2007)
23. von Neumann, J.: Various techniques used in connection with random digits. In: Householder, A., Forsythe, G., Germond, H. (eds.) Monte Carlo Method, pp. 36–38. National Bureau of Standards Applied Mathematics Series, 12, Washington, D.C.: U.S. Government Printing Office (1951)
24. Walter, M.: Private communication (2020), date: 2020-01-29
25. Zhang, Z., Chen, C., Hoffstein, J., Whyte, W.: NIST PQ submission: pqNTRUSign a modular lattice signature scheme. https://www.onboardsecurity.com/nist-post-quantum-crypto-submission (2017), accessed: 2019-08-01
26. Zhao, R.K., Steinfeld, R., Sakzad, A.: FACCT: fast, compact, and constant-time discrete Gaussian sampler over integers. IEEE Trans. Computers **69**(1), 126–137 (2020)

A Precision Analysis

To avoid sampling a uniformly random real r with high absolute precisions at rejection steps 11 and 23 in Algorithm 2, and step 4 in Algorithm 3, we adapt the comparison approach similar to [26]. Assume an IEEE-754 floating-point value

$f \in (0,1)$ with (δ_f+1)-bit precision is represented by $f = \left(1 + mantissa \cdot 2^{-\delta_f}\right) \cdot 2^{exponent}$, where integer $mantissa$ has δ_f bits and $exponent \in \mathbb{Z}^-$. To check $r < f$, one can sample $r_m \hookleftarrow \mathcal{U}\left(\{0,1\}^{\delta_f+1}\right)$, $r_e \hookleftarrow \mathcal{U}\left(\{0,1\}^{l}\right)$, and check $r_m < mantissa+2^{\delta_f}$ and $r_e < 2^{l+exponent+1}$ instead for some l such that $l+exponent+1 \geq 0$.

Here we analyse the precision requirement of r_e. We have the following theorem for the worst-case acceptance rate in Algorithm 2:

Theorem 5. *Assume $x \in [-\tau, \tau]$ and $y \in [-\tau\sigma - 1, \tau\sigma + 1]$. In worst case, step 11 in Algorithm 2 has the acceptance rate:*

$$p_1 \geq \exp\left(-\frac{(-2\tau\sigma + c_F - 3/2)(c_F - 3/2)}{2\sigma^2}\right),$$

and step 23 in Algorithm 2 has the acceptance rate:

$$p_2 \geq \exp\left(-\frac{(2\tau\sigma + c_F + 3/2)(c_F + 3/2)}{2\sigma^2}\right).$$

Proof. For $b = 0$ and $y \leq -1/2$, we have the acceptance rate $p_1 = \exp\left(-Y_1/\left(2\sigma^2\right)\right)$ at step 11 in Algorithm 2 where:

$$\begin{aligned}
Y_1 &= (\lfloor y \rceil + c_F)^2 - (y+1)^2 \\
&= (y + \delta + c_F)^2 - (y+1)^2 \quad (\lfloor y \rceil = y + \delta \text{ where } \delta \in [-1/2, 1/2]) \\
&= (2y + \delta + c_F + 1)(\delta + c_F - 1) \\
&\leq (-2\tau\sigma + c_F - 3/2)(c_F - 3/2). \quad (\text{when } \delta = -1/2 \text{ and } y = -\tau\sigma - 1)
\end{aligned}$$

Similarly, for $b = 1$ and $y \geq 1/2$, we have the acceptance rate $p_2 = \exp\left(-Y_2/\left(2\sigma^2\right)\right)$ at step 23 in Algorithm 2 where:

$$\begin{aligned}
Y_2 &= (\lfloor y \rceil + c_F)^2 - (y-1)^2 \\
&= (y + \delta + c_F)^2 - (y-1)^2 \quad (\lfloor y \rceil = y + \delta \text{ where } \delta \in [-1/2, 1/2]) \\
&= (2y + \delta + c_F - 1)(\delta + c_F + 1) \\
&\leq (2\tau\sigma + c_F + 3/2)(c_F + 3/2). \quad (\text{when } \delta = 1/2 \text{ and } y = \tau\sigma + 1)
\end{aligned}$$

□

Let $\Delta \leq 1/2$ be the maximum relative error of the right hand side computations at rejection steps 11 and 23 in Algorithm 2, and step 4 in Algorithm 3. For $\exp\left(-Y_1/\left(2\sigma^2\right)\right)$ at step 11 in Algorithm 2, we have:

$$\begin{aligned}
exponent_1 &\geq \left\lfloor \log_2\left((1-\Delta)\cdot\exp\left(-\frac{Y_1}{2\sigma^2}\right)\right)\right\rfloor \\
&\geq \left\lfloor -1 - \frac{(-2\tau\sigma + c_F - 3/2)(c_F - 3/2)}{2\sigma^2}\cdot\log_2 e\right\rfloor \quad (\text{by Thm. 5 and } \Delta \leq 1/2) \\
&\geq \left\lfloor -1 - \frac{2\tau\sigma + 2}{\sigma^2}\cdot\log_2 e\right\rfloor. \quad (\text{when } c_F = -1/2)
\end{aligned}$$

Similarly, for $\exp\left(-Y_2/\left(2\sigma^2\right)\right)$ at step 23 in Algorithm 2, we have:

$$\begin{aligned}
exponent_2 &\geq \left\lfloor \log_2\left((1-\Delta)\cdot\exp\left(-\frac{Y_2}{2\sigma^2}\right)\right)\right\rfloor \\
&\geq \left\lfloor -1-\frac{\left(2\tau\sigma+c_F+3/2\right)\left(c_F+3/2\right)}{2\sigma^2}\cdot\log_2 e\right\rfloor \quad \text{(by Thm. 5 and } \Delta\leq 1/2) \\
&\geq \left\lfloor -1-\frac{2\tau\sigma+2}{\sigma^2}\cdot\log_2 e\right\rfloor. \quad \text{(when } c_F=1/2)
\end{aligned}$$

For $\exp\left(-c_F^2/\left(2\sigma^2\right)\right)/S$ at step 4 in Algorithm 3, we have:

$$\begin{aligned}
exponent_3 &\geq \left\lfloor \log_2\left((1-\Delta)\cdot\exp\left(-\frac{c_F^2}{2\sigma^2}\right)/S\right)\right\rfloor \\
&\geq \left\lfloor -1-\frac{1}{8\sigma^2}\cdot\log_2 e-\log_2\left(\sigma\sqrt{2\pi}\right)\right\rfloor. \quad \text{(when } c_F=\pm 1/2 \text{ and } \Delta\leq 1/2)
\end{aligned}$$

Therefore, we have:

$$exponent \geq \min\left\{\left\lfloor -1-\frac{2\tau\sigma+2}{\sigma^2}\cdot\log_2 e\right\rfloor, \left\lfloor -1-\frac{1}{8\sigma^2}\cdot\log_2 e-\log_2\left(\sigma\sqrt{2\pi}\right)\right\rfloor\right\}.$$

Since the probability $\Pr\left[-\tau\leq x\leq\tau\right]=\operatorname{erf}\left(\tau/\sqrt{2}\right)$ for $x\hookleftarrow\mathcal{N}(0,1)$, to ensure $1-\Pr\left[-\tau\leq x\leq\tau\right]\leq 2^{-\lambda}$, we need $\tau\geq\sqrt{2}\cdot\operatorname{erf}^{-1}\left(1-2^{-\lambda}\right)$. Therefore, for $\lambda=128$ and $\sigma\in\left[2,2^{20}\right]$, we have $\tau\geq 13.11$, $exponent\geq -23$, and thus $l\geq 22$, i.e. r_e needs to have at least 22 bits.

B Proof of Algorithm 1

Since Algorithm 1 was an exercise in [4] without solutions, here we provide a brief proof of Algorithm 1.

Normalisation Factor By definition, we have the normalisation factor:

$$\begin{aligned}
\frac{1}{c} &= \sum_{k\in\mathbb{Z}}\exp\left(-\frac{\left(|k|+1/2\right)^2}{2\sigma^2}\right) \\
&= \sum_{k\in\mathbb{Z}^-}\exp\left(-\frac{\left(k-1/2\right)^2}{2\sigma^2}\right)+\exp\left(-\frac{1}{8\sigma^2}\right)+\sum_{k\in\mathbb{Z}^+}\exp\left(-\frac{\left(k+1/2\right)^2}{2\sigma^2}\right) \\
&= \rho_{1/2,\sigma}\left(\mathbb{Z}^-\right)+\rho_{-1/2,\sigma}\left(\mathbb{Z}^+\right)+\exp\left(-\frac{1}{8\sigma^2}\right) \\
&\geq \frac{1-\epsilon}{1+\epsilon}\cdot\rho_\sigma\left(\mathbb{Z}\right)+\exp\left(-\frac{1}{8\sigma^2}\right)-2. \quad \text{(By Lemma 1)}
\end{aligned}$$

Correctness Let $z_0 = \lfloor x \rceil$. We have $Y = (|z_0| + 1/2)^2 - x^2 \geq 0$ for any $x \in \mathbb{R}$. Therefore, the rejection condition $\exp\left(-Y/\left(2\sigma^2\right)\right) \in (0, 1]$. We have the output distribution:

$$\begin{aligned}\Pr\left[z = z_0\right] &\propto \int_{z_0-1/2}^{z_0+1/2} \exp\left(-\frac{x^2}{2\sigma^2}\right) \cdot \exp\left(-\frac{(|z_0| + 1/2)^2 - x^2}{2\sigma^2}\right) \mathrm{d}x \\ &= \int_{z_0-1/2}^{z_0+1/2} \exp\left(-\frac{(|z_0| + 1/2)^2}{2\sigma^2}\right) \mathrm{d}x = \exp\left(-\frac{(|z_0| + 1/2)^2}{2\sigma^2}\right).\end{aligned}$$

Rejection Rate By definition, we have the output probability density function $f(x) = c \cdot \exp\left(-\frac{(|\lfloor x \rceil| + 1/2)^2}{2\sigma^2}\right)$ and the input probability density function $g(x) = \rho_\sigma(x) / \left(\sigma\sqrt{2\pi}\right)$. The expected number of trials can be written as:

$$M = \max \frac{f(x)}{g(x)} = \max\left(\frac{\exp\left(-\frac{(|\lfloor x \rceil| + 1/2)^2}{2\sigma^2}\right)}{\rho_\sigma(x)} \cdot \frac{\sigma\sqrt{2\pi}}{1/c}\right).$$

We have:

$$\frac{\exp\left(-\frac{(|\lfloor x \rceil| + 1/2)^2}{2\sigma^2}\right)}{\rho_\sigma(x)} = \frac{\exp\left(-\frac{(|\lfloor x \rceil| + 1/2)^2}{2\sigma^2}\right)}{\exp\left(-\frac{x^2}{2\sigma^2}\right)} = \exp\left(-\frac{(|\lfloor x \rceil| + 1/2)^2 - x^2}{2\sigma^2}\right) \leq 1.$$

Therefore,

$$\begin{aligned}M \leq \frac{\sigma\sqrt{2\pi}}{1/c} &\leq \frac{\sigma\sqrt{2\pi}}{\frac{1-\epsilon}{1+\epsilon} \cdot \rho_\sigma(\mathbb{Z}) + \exp\left(-\frac{1}{8\sigma^2}\right) - 2} \\ &\leq \frac{\sigma\sqrt{2\pi}}{\frac{1-\epsilon}{1+\epsilon} \cdot \left(\sigma\sqrt{2\pi} - 1\right) + \exp\left(-\frac{1}{8\sigma^2}\right) - 2},\end{aligned}$$

where the second inequality follows from the inequality of $1/c$ and the third inequality follows from the fact that $\rho_\sigma(\mathbb{Z}) \geq \sigma\sqrt{2\pi} - 1$. Thus, we have the rejection probability:

$$1 - \frac{1}{M} \leq \frac{\left(1 - \frac{1-\epsilon}{1+\epsilon}\right) \cdot \sigma\sqrt{2\pi} + \frac{1-\epsilon}{1+\epsilon} - \exp\left(-\frac{1}{8\sigma^2}\right) + 2}{\sigma\sqrt{2\pi}} \approx \frac{3 - \exp\left(-\frac{1}{8\sigma^2}\right)}{\sigma\sqrt{2\pi}} \leq \frac{2}{\sigma}\sqrt{\frac{2}{\pi}},$$

when ϵ is small.