

# Discrete Ziggurat: A Time-Memory Trade-off for Sampling from a Gaussian Distribution over the Integers

Johannes Buchmann, Daniel Cabarcas, Florian Göpfert,  
Andreas Hülsing, and Patrick Weiden

Technische Universität Darmstadt, Germany

**Abstract.** Several lattice-based cryptosystems require to sample from a discrete Gaussian distribution over the integers. Existing methods to sample from such a distribution either need large amounts of memory or they are very slow. In this paper we explore a different method that allows for a flexible time-memory trade-off, offering developers freedom in choosing how much space they can spare to store precomputed values. We prove that the generated distribution is close enough to a discrete Gaussian to be used in lattice-based cryptography. Moreover, we report on an implementation of the method and compare its performance to existing methods from the literature. We show that for large standard deviations, the Ziggurat algorithm outperforms all existing methods.

**Keywords** Lattice-Based Cryptography, Gaussian Sampling, Practicality, Implementation

## 1 Introduction

The object of study of this paper is the discrete Gaussian probability distribution over the integers. Sampling elements from such a distribution is widely used in lattice-based cryptography [GPV08,LP11,BGV12,GGH12]. It is a critical technical challenge to sample from a discrete Gaussian over the integers accurately and efficiently. Weiden et. al. [WHCB13] report that sampling from it takes more than 50% of the running time of the signing algorithm in their implementation of Lyubashevsky’s signature scheme [Lyu12].

All existing methods to sample from a Gaussian distribution over the integers either need large amounts of memory or they are very slow. For example, Galbraith and Dwarakanath estimate that Peikert’s sampler [Pei10] requires around 12MB of storage [GD12] for some parameters. Such a large memory requirement might be acceptable on a PC but not on the diversity of devices that demand cryptographic solutions today.

In this paper we explore a different alternative for sampling from a Gaussian distribution over the integers that offers a flexible trade-off between speed and memory. Moreover, for big standard deviations, this method beats commonly used methods. We call the method discrete Ziggurat because it adapts the Ziggurat algorithm [MT00] for the discrete case.

The discrete Ziggurat is specially appealing for cryptographic applications because of its flexibility. The method uses precomputed rectangles of equal ‘size’ to cover the area under the probability density function (PDF). Increasing the number of rectangles increases speed but also increases the memory used. Therefore, it offers an easy-to-tune trade-off between speed and memory. This is a desirable property because developers of cryptographic primitives can easily adjust it to fit the particular characteristics of different devices. On memory constraint devices like smartcards or microcontrollers they could use a low-memory low-speed setting, while on a high performing server they could use a high-memory high-speed configuration.

Originally, the Ziggurat sampler was developed for a continuous distribution. In order to adapt it to the discrete case some care must be taken. In particular the notion of ‘size’ of a rectangle must be redefined from the narrow concept of ‘area’ to the more general “probability to sample points inside the rectangle”. We discuss the implications of this generalization.

It is also challenging to analyze the quality of the discrete Ziggurat because of the subtleties of an actual implementation. In this paper we provide a careful analysis that takes into consideration the loss of precision due to the tailcut, the precision in sampling from the  $y$ -axis and the precision in calculating the PDF. The techniques used and the way they are combined in this analysis might show valuable for the analysis of other samplers. For developers we explain how to achieve a desired accuracy by setting the precision for representing numbers.

We implemented the discrete Ziggurat in C++ using the Number Theory Library (NTL) [Sho]. The implementation can be downloaded at the authors’ homepage<sup>1</sup>. We compare the efficiency of the discrete Ziggurat with existing methods and analyze the speed-memory trade-off. For example, we used the parameters proposed by Galbraith and Dwarakanath [GD12] for the normal distribution in Lyubashevsky’s signature scheme [Lyu12]. For this illustrative setting, the discrete Ziggurat

---

<sup>1</sup> In particular at <https://www.cdc.informatik.tu-darmstadt.de/~pschmidt/implementations/ziggurat/ziggurat-src.zip>.

produces about 1.13 million samples per second, using only 524 kilobytes of memory. In comparison, Peikert’s sampler outputs 281,000 samples per second for a memory usage of 33.55 megabytes. The Knuth-Yao algorithm is only slightly faster (it produces about 4% more samples), but increases the memory-consumption by a factor of more than 400.

*Related Work.* We briefly survey existing alternatives to sample from a discrete Gaussian probability distribution over the integers, denoted  $D_\sigma$ . For parameter  $\sigma > 0$ ,  $D_\sigma$  assigns  $x \in \mathbb{Z}$  a probability proportional to  $\rho_\sigma(x) = \exp(-\frac{1}{2}x^2/\sigma^2)$ . It is important to note that sampling from  $D_\sigma$  is different to sampling from a (continuous) normal distribution [TLLV07]. Another related problem is that of sampling from a Gaussian distribution over a generic lattice, a more complex problem, whose solutions often require sampling from  $D_\sigma$  as a subroutine [GPV08,Pei10,DN12,AGHS12].

For cryptographic applications it is sufficient to sample from the bounded subset  $B := \mathbb{Z} \cap [-t\sigma, t\sigma]$ , where the tailcut  $t > 0$  is chosen large enough to guarantee a desired precision [GPV08]. One alternative to sample from  $D_\sigma$  is to do rejection sampling on  $B$ . Another alternative is to precompute the cumulative distribution function (CDF) for  $x \in B$ , sample a uniform element  $y \in [0, 1)$  and perform a binary search on the CDF table to output the “inverse CDF” of  $y$  [Pei10]. To the best of our knowledge, no work analyzes the accuracy or efficiency of any of these methods in detail.

Yet another alternative, explored by Galbraith and Dwarakanath [GD12], is the Knuth-Yao algorithm. The algorithm first precomputes a binary tree with leaves labeled by the elements of  $B$ . For  $x \in B$ , if the probability of sampling  $x$  has a one in the  $i$ -th place of its binary representation, there is a leaf labeled  $x$  at height  $i$  of the tree. Then it samples by walking down the tree using one uniform bit at each step to decide which of the two children to move to. Galbraith and Dwarakanath present a very detailed analysis of the accuracy of the sampler and of the number of random bits it uses. They also propose ways to reduce the memory requirements. However, they do not assess the speed of the sampler.

Ducas and Nguyen propose an enhancement for rejection sampling. They observe that the sampler can compute at a low precision by default and only use high precision computation when a certain threshold is reached [DN12]. To the best of our knowledge, no work evaluates the effect of this enhancement in detail.

*Organization.* In Section 2 we explain the Ziggurat algorithm and we describe in detail its discrete variant. In Section 3 we analyze the quality of the distribution. Finally, in Section 4 we report on experimental results.

## 2 The Discrete Ziggurat Algorithm

The Ziggurat algorithm belongs to the class of rejection sampling algorithms and was introduced by Marsaglia and Tsang for sampling from a continuous Gaussian distribution [MT00]. Here we adapt it for the discrete case. After explaining the setting, we give a short overview over Ziggurat in the continuous case and shortly explain how to control the trade-off. Afterwards, we show how we adapted it to the discrete case and explain how to perform the necessary precomputing. Subsequently, we discuss the implementation-details and finish the section with further improvements.

### 2.1 Setting

We are concerned with sampling from a discrete Gaussian distribution centered at zero with bounded support  $B := [-t\sigma, t\sigma] \cap \mathbb{Z}$  for some parameter  $t > 0$ . This bounded support is sufficient for the application in lattice-based cryptography as long as  $t$  is chosen large enough. Moreover, we show in Section 3.2 how to select parameters such that the sampled distribution is within a certain statistical distance to a (truly) discrete Gaussian distribution. The assumption that the distribution is centered at zero is also fine, as we can add a constant offset to transform samples into a distribution centered around any other integer.

### 2.2 Intuition

We briefly review the continuous Ziggurat for the above setting to give some intuition. As the target distribution is symmetric, we can proceed as follows. We use the method to sample a value  $x \leq t\sigma$  within  $\mathbb{R}_0^+$ . Afterwards, if  $x = 0$  we accept with probability  $1/2$ . Otherwise, we sample a sign  $s \in \{-1, 1\}$  and return the signed value  $sx$ .

Now, how do we sample  $x$  within  $\mathbb{R}_0^+$ ? During set-up, we enclose the area of the probability density function (PDF) in an area  $A$  consisting of  $m$  horizontal rectangles with equal area as shown in Figure 1. How the rectangles are computed is described below. Next, we store the coordinates  $(x_i, y_i)$  of the lower right corner of each rectangle  $R_i$ ,  $1 < i < m - 1$ .

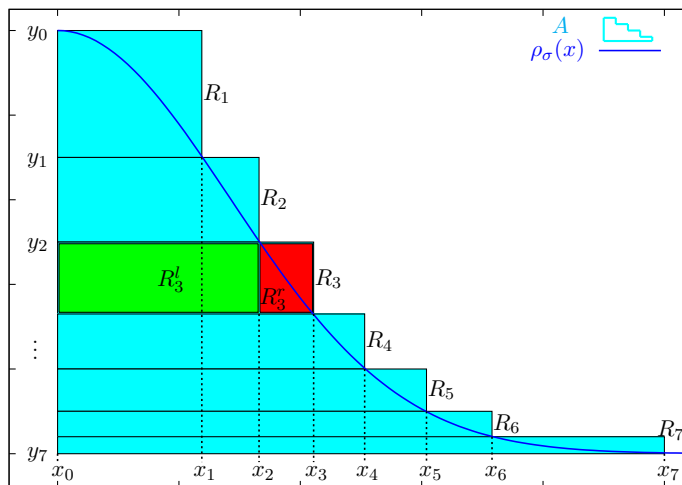


Fig. 1. Ziggurat for  $m = 7$  with covering area  $A$  and the partition into rectangles.

Please note that each rectangle  $R_i$  can be split into a left rectangle  $R_i^l$  that lies completely within the area of the PDF and a right rectangle  $R_i^r$  that is only partially covered by the PDF. For an example, see  $R_3$  in Figure 1.

Now, to actually sample a value  $x \leq t\sigma$  within  $\mathbb{R}_0^+$  we first sample an integer  $1 \leq i \leq m$  uniformly at random, to select a random rectangle. Next, we sample an  $x$ -coordinate inside  $R_i$ , by sampling a uniformly random  $x'$  within  $[0, x_i]$ . If  $x' \leq x_{i-1}$ , i.e. if  $x'$  is inside  $R_i^l$ , we directly accept and return  $x'$ . Otherwise,  $x'$  lies within  $R_i^r$ . In this case, we do rejection sampling. Namely, we sample a value  $\gamma$  within  $[y_{i+1}, y_i]$  uniformly at random. Then, if  $\gamma + y_{i+1} \leq \rho_\sigma(x')$ , i.e. we hit a point in the area below the PDF, we accept and return  $x'$ . Otherwise, we reject and restart the whole process by sampling a new  $i$ .

In order to understand why this sampling-algorithm works, think of it as an efficient implementation of rejection-sampling in the area  $A$ . More precisely, the implementation of the first step (sampling a point in the enclosing area) is improved. Since all the rectangles have equal size, the probabilities of sampling a point in a given rectangle are equal. Therefore, one can sample the rectangle first and a point in the rectangle afterwards.

The expensive part of the algorithm is computing  $\rho_\sigma(x')$  if  $x'$  does not lie within  $R_i^l$ . It becomes even more expensive whenever a value is rejected. For this reason Ziggurat provides a time-memory trade-off, which is controlled by the number of rectangles used, as follows. If we use more rectangles, the ratio between the left and the right rectangle within one

rectangle is changed in such a way that the left rectangle becomes comparatively bigger. Hence, we accept an  $x'$  without computing  $\rho_\sigma(x')$  with higher probability. Moreover, using more rectangles, the area  $A$  tighter encloses the area  $C$  below the PDF. Thereby, the area  $A \setminus C$  that leads to a rejection shrinks and with it the overall probability of a rejection. However, for each additional rectangle the coordinates of one additional point have to be stored, increasing the memory requirements.

### 2.3 Adaption to the Discrete Case

In the discrete case, the algorithm works quite similar. The whole pseudocode can be found in Appendix A.1. As before, a sign  $s$ , a rectangle with index  $i$  and a potential sample  $x'$  are sampled. If  $x'$  lies in a left rectangle and is non-zero,  $sx'$  is returned immediately. If  $x'$  equals zero, it is returned with probability  $1/2$ , like in the continuous case. If not, exactly the same rejection sampling procedure as in the continuous case is used to decide whether  $sx'$  is returned or the whole process is restarted.

In contrast to the continuous case, the notion of ‘size’ defined using the area of a rectangle can not be used in the discrete case. We have seen in the last section that the size of a rectangle has to be proportional to the probability to sample a point in it. In the discrete case, we therefore define the size of a rectangle as the number of integer  $x$ -coordinates in the rectangle times its height. For instance, the rectangle  $R_3$  has size  $(1 + \lfloor x_3 \rfloor) \cdot (y_2 - y_3)$ .

The second difference between the continuous and the discrete case is the way the rectangles are computed. While we did not explain how this is done in the continuous case, as it would go beyond the scope of this work, we give a description for the discrete case. We explain how to obtain a partition for the Ziggurat algorithm for a given number of  $m$  rectangles where each rectangle has exactly the same ‘size’  $S$ . Therefore, we set

$$y_m := 0, \quad x_0 := 0 \quad \text{and} \quad x_m := t\sigma,$$

and we iteratively compute a possible partition “from right to left” via

$$\begin{aligned} y_{m-1} &= \frac{S}{1 + \lfloor x_m \rfloor}, & x_{m-1} &= \rho_\sigma^{-1}(y_{m-1}), \\ \text{for } i = m-2, \dots, 1 : & & y_i &= \frac{S}{1 + \lfloor x_{i+1} \rfloor} + y_{i+1}, & x_i &= \rho_\sigma^{-1}(y_i), \\ y_0 &= \frac{S}{1 + \lfloor x_1 \rfloor} + y_1. \end{aligned}$$

Recall that  $\rho_\sigma$  is a scaled density function with  $\rho_\sigma(0) = 1$ . Therefore, a valid partition for Ziggurat requires  $y_0 \geq 1$ , since only then the partition completely covers the area under the curve  $\rho_\sigma$  on the support  $B_0^+ := [0, t\sigma] \cap \mathbb{Z}_0^+$ . Since the value  $y_0$  depends on the ‘size’  $S$  of the rectangles, any value of  $S$  for which  $y_0 \geq 1$  leads to a valid partition. We heuristically determine  $S$  as follows. We set  $S = \sigma / (m \cdot \sqrt{\pi/2}) \cdot c$  with initial value  $c = 1$ , compute the corresponding partition, and increase  $c$  stepwise as long as  $y_0 < 1$ . (To improve the quality of the input partition, i.e. minimizing  $y_0 - 1$ , one can perform a binary search for  $S$  in  $[\sigma / (m \cdot \sqrt{\pi/2}), t\sigma + 1]$ .) In the case that no valid partition is found, we increase  $x_m$  by one and restart the whole process. Reaching  $x_m = (t + 1)\sigma$ , we abort. We note that this method ended with no partition being output in only about 1.3% of our computations. In these cases, i.e. when no valid partition is found, one can re-run the procedure with one or more of the following changed: number of rectangles  $m$ , Gaussian parameter  $\sigma$  (if possible), or upper bound on  $x_m$ .

## 2.4 Implementation

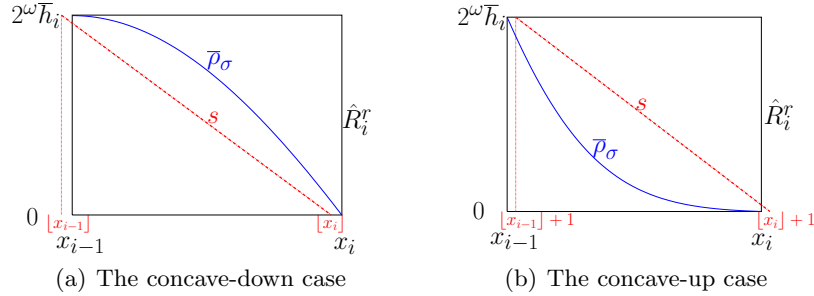
For an implementation, we have to analyze the effect of computing with limited precision. We use a dash over numbers or functions to indicate the use of their corresponding  $n$ -bit fixed-point approximation, e.g.  $\bar{y}$  and  $\bar{\rho}_\sigma$  denote the  $n$ -bit approximation of  $y \in \mathbb{R}$  and the function  $\rho_\sigma$ , respectively. Since we can exactly calculate  $\bar{\rho}_\sigma$ , we can find a partition such that the rectangles have exactly the same ‘size’ and represent it with the vertical bounds  $\bar{y}_i$  (which we store with  $n$  bits fixed point precision) and the rounded horizontal borders  $\lfloor x_i \rfloor$ . The last problem is to sample uniformly at random in the infinite sets  $[\bar{y}_i, \bar{y}_{i-1}]$ . Our solution is to discretize the set: We define  $\bar{h}_i := \bar{y}_{i-1} - \bar{y}_i$  to be the height of the  $i$ -th rectangle, sample  $y' \stackrel{\$}{\leftarrow} \{0, 1, \dots, 2^\omega - 1\}$  for a parameter  $\omega \in \mathbb{Z}_0^+$  and transform the samples to  $\bar{y} = \bar{h}_i y' \in [0, 2^\omega \bar{h}_i]$ . Instead of transforming  $\bar{y}$  into the interval  $[\bar{y}_i, \bar{y}_{i-1}]$  we replace the condition  $\bar{y} \leq \bar{\rho}_\sigma(x)$  for  $\bar{y} \in [\bar{y}_i, \bar{y}_{i-1}]$  with  $\bar{y} \leq 2^\omega (\bar{\rho}_\sigma(x) - \bar{y}_i)$  for  $\bar{y} \in [0, 2^\omega \bar{h}_i]$ . We show in Section 3 how to choose the parameters  $t$ ,  $\omega$  and  $n$  in order to bound the statistical distance between the distribution defined by our algorithm and  $D_\sigma$  by a given value.

## 2.5 Further Improvement

Since the most time-consuming part of the discrete Ziggurat is the computation of  $\bar{\rho}_\sigma$ , we want to avoid it as often as possible. As mentioned

above, it is only necessary if  $(x, y)$  is contained in a right rectangle  $R_i^r$ . But even in this case, depending on the shape of  $\bar{\rho}_\sigma$  inside of  $R_i^r$ , we can avoid the evaluation of  $\bar{\rho}_\sigma$  in nearly half of the cases and more easily reject or accept  $x$  as follows.

We divide  $R_i^r$  by connecting its upper left and lower right corner by a straight line  $s$ . Since  $\bar{\rho}_\sigma$  has inflection point  $\sigma$ , it is concave-down for  $x \leq \sigma$ , and concave-up otherwise. In the concave-down case ( $x_i \leq \sigma$ ) all points  $(x, y)$  in  $R_i^r$  below  $s$  implicitly fulfill the acceptance condition, thus  $x$  is instantly output. In the concave-up case ( $\sigma \leq x_{i-1}$ ) all points above  $s$  lead to immediate rejection. In all other cases we have to evaluate  $\bar{\rho}_\sigma(x)$  and check the acceptance condition. For the discrete Ziggurat we have to adjust this approach to our way of sampling  $\bar{y}_i$  and our use of the values  $\lfloor x_i \rfloor$  instead of  $x_i$  (for an idea how to accomplish this see Figure 2).



**Fig. 2.** Optimization to discrete Ziggurat ( $\hat{R}_i^r$  is  $R_i^r$  vertically shifted and stretched)

### 3 Quality of our Sampler

In this section, we show how to choose parameters for the algorithm such that it achieves a given quality in the sense of statistical distance to a discrete normal distribution. We begin with a theorem that bounds the statistical distance between the distribution produced by the sampler and a discrete normal distribution. Afterwards, we show as an example how to select parameters such that the statistical distance is smaller than  $2^{-100}$ . The methodology can be used to select parameters for any given statistical distance.



### 3.1 Statistical Distance Between Sampled and Gaussian Distribution

No practical algorithm outputs samples exactly distributed according to a Gaussian distribution. Therefore, it is important to understand how much the produced output distribution differs from the normal distribution. This difference is measured by the statistical distance. Recall that  $t$  determines the tailcut and  $\omega$  the precision of the sampled  $y$ -values. As explained before, we use numbers with  $n$ -bit fixed-point precision. Similar to the definition of the support  $B_0^+ = [0, t\sigma] \cap \mathbb{Z}_0^+$ , we define  $B^+ := [0, t\sigma] \cap \mathbb{Z}^+$ . The next theorem gives a lower bound on the quality of our sampler depending on the used parameters.

**Theorem 1.** *The statistical distance between the discrete Gaussian distribution  $D_\sigma$  and the distribution  $\bar{D}_\sigma$  output by our algorithm is bounded by*

$$\Delta(D_\sigma, \bar{D}_\sigma) < te^{(1-t^2)/2} + \frac{|B_0^+|}{\bar{\rho}_\sigma(B^+) + \frac{1}{2}}(2^{-\omega+1} + 2^{-n}). \quad (1)$$

For readability we solely give an intuition of the proof here, while the complete proof is given Appendix A.3. The proof follows a hybrid argument. We introduce intermediary distributions. The first intermediary distribution differs from a Gaussian distribution by the tailcut. The second intermediary distribution takes the limited precision of the stored numbers and the sampled  $y$ -values into consideration. After bounding the statistical distances between the consecutive distributions, we apply the triangle inequality to show the main result.

### 3.2 Parameter Selection

We now show how to choose  $t$ ,  $n$  and  $\omega$  such that the statistical distance of our distribution and the discrete Gaussian distribution is below  $2^{-100}$  for  $\sigma = 10$ . We choose  $t$  to be the smallest positive integer such that  $t \exp((1-t^2)/2) < 2^{-101}$ , which is  $t = 13$ . Furthermore, we choose  $\omega = n + 1$  and obtain  $2^{-\omega+1} + 2^{-n} = 2^{-n+1}$ . We can now find an  $n$  such that the second addend of inequality (1) is bounded by  $2^{-101}$ . Since this calculation is a little bit complex, we try to get a feeling for the expected result first. Since  $t$  was chosen such that the probability of sampling an element in the tail is extremely small, we obtain

$$\bar{\rho}_\sigma(B^+) + \frac{1}{2} \approx \rho_\sigma(B^+) + \frac{1}{2} \approx \rho_\sigma(\mathbb{Z}_0^+) \approx \int_0^\infty \rho_\sigma(x) dx = \sigma \sqrt{\frac{\pi}{2}}$$

and expect

$$2^{-n+1} \frac{|B_0^+|}{\bar{\rho}_\sigma(B^+)} \approx 2^{-n+1} \frac{t\sigma}{\sigma\sqrt{\pi/2}} \approx 2^{-n+1}t \approx 2^{-n+5}.$$

The smallest  $n$  satisfying  $5 - n \leq -101$  is  $n = 106$ . An exact calculation shows indeed that  $n = 106$  suffices.

## 4 Experiments and Results

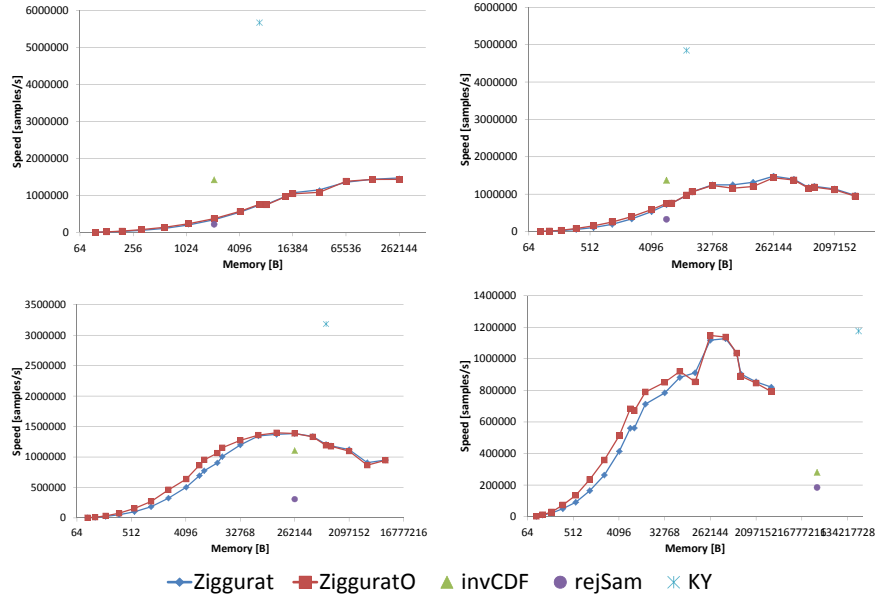
In this section we discuss the performance of our implementation of the discrete Ziggurat algorithm. We first describe the experiments we performed to test the efficiency, then present their results and analyze the gathered data. Furthermore, we compare our sampler to implementations of existing samplers for discrete Gaussians.

### 4.1 Environment and Setup

The experiments were carried out on a Sun XFire 4400 server with 16 Quad-Core AMD Opteron 8356 CPUs running at 2.3GHz (we only used one CPU), having in total 64GB of RAM and running a 64bit Debian 7.1. All implementations use the Number Theory Library (NTL, cf. [Sho]) with precision  $n = 106$  bits in consistency to our choice of parameters in Section 3.2 to assure a statistical distance for Ziggurat of at most  $2^{-100}$ . Furthermore, we used the tailcut  $t = 13$  and the discrete Gaussian parameters  $\sigma \in \{10, 32, 1000, 1.6 \cdot 10^5\}$ . The value  $\sigma = 32$  maintains the worst-to-average-case reduction [Reg05] in several schemes for a certain parameter set, and the rather large value  $\sigma = 1.6 \cdot 10^5$  is chosen according to Galbraith and Dwarakanath [GD12]. The other two values  $\sigma = 10, 1000$  were chosen arbitrarily inbetween and at the lower end to allow a better comparison.

We queried each algorithm iteratively 1 million times to output a single sample per call. These experiments were applied to the discrete Ziggurat with the optimization using the straight line  $s$  (ZigguratO), discrete Ziggurat without optimization (Ziggurat), inverse CDF (invCDF), rejection sampling with precomputed lookup-table (rejSam), and Knuth-Yao (KY). Furthermore we tested both Ziggurat algorithms with a pre-computed lookup-table for the support  $B_0^+$  (ZigguratOP and ZigguratP, respectively).

For each algorithm we measured the running time using the (Linux-internal) function `clock_gettime` with clock `CLOCK_PROCESS_CPUTIME_ID`.



**Fig. 3.** Results for inverse CDF, rejection sampling, Knuth-Yao, and discrete Ziggurat with and without optimization for parameters  $\sigma = 10, 32, 1000, 1.6 \cdot 10^5$ , respectively.

In order to have non-distorted results we excluded all pre- and post-computations (e.g. setting up lookup-tables) from the measurements. Regarding the memory, we did not perform per-runtime analyses, but computed the amount of memory by adding up the number of fixed variables in regard to their types in NTL. For our choice of parameters, in Ziggurat(O) the values on the  $x$ -axis need 8 bytes and on the  $y$ -axis 24 bytes of memory. With  $m$  rectangles the total amount of memory is thus  $32(m + 2)$  bytes (including  $\sigma, t, \omega, m$ ). For both invCDF and rejSam we need to store a lookup-table of  $t\sigma$  values à 16 bytes, resulting in 2080 bytes for  $\sigma = 10$ . The same amount of memory is used by Ziggurat(O) with  $m = 63$  rectangles. The size of Knuth-Yao is approximated by  $(\#\text{intermediates} + \#\text{leaves})/2$  bits, where  $\#\text{intermediates} = n \cdot 2^{\lceil \log \log(n \cdot t\sigma) \rceil}$  and  $\#\text{leaves} = n \cdot 2^{\lceil \log \log(t\sigma) \rceil}$  for precision  $n = 106$  bits as above.

## 4.2 Results

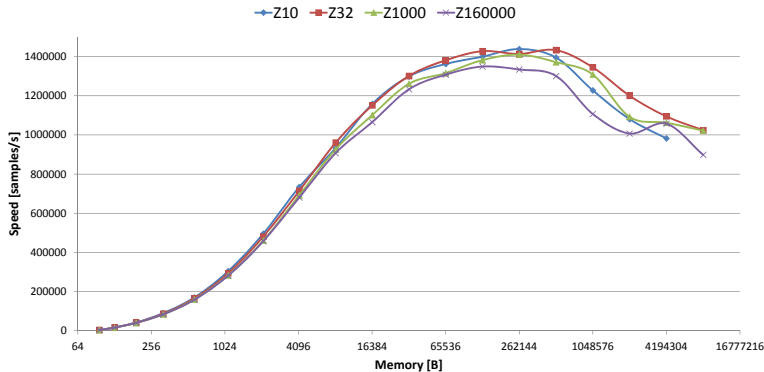
Figure 3 shows results for inverse CDF, rejection sampling, Knuth-Yao, and discrete Ziggurat with and without optimizations for different numbers of rectangles. It shows four different graphs for different values of  $\sigma$ . For small values of  $\sigma$ , the inverse CDF method outperforms both discrete

Ziggurat and rejection sampling for the same fixed amount of memory. For example, our implementation `invCDF` samples about 1.37 million samples per second for  $\sigma = 32$ . On the other hand, rejection sampling is quite slow due to a large rejection area. Even with a precomputed lookup-table, `rejSam` only achieves about 327,000 samples per second, which is a factor 4.2 slower than `invCDF`. The naïve approach without lookup-table solely achieves 2,500 samples per second, being a factor 558 slower than `invCDF`. For the same amount of memory, `ZigguratO` achieves an overall number of about 753,000 samples per second, while `Ziggurat` outputs 718,000 samples per second. Compared to the other two methods, `Ziggurat` is 1.91 times slower than `invCDF` and a factor 2.19 faster than `rejSam`. Our implementation of Knuth-Yao outperforms all the other methods by at least a factor of 3.53, outputting 4.85 million samples per second. This comes at the cost of nearly doubled memory usage.

In the extreme case  $\sigma = 1.6 \cdot 10^5$ , the fastest instantiation of `Ziggurat` outputs 1.13 million samples per second with a memory usage of 524 kilobytes. Inverse CDF creates 281,000 samples per second while using 33.55 megabytes, thus being about a factor 4 slower than `Ziggurat`. For `rejSam` the situation is even worse: Using the same amount of memory as `invCDF`, it only outputs 185,000 samples per second – a factor 6.1 slower than `Ziggurat`. The Knuth-Yao algorithm still performs better than `Ziggurat`, but only by 4.26%. On the other hand, KY needs more than 424 times the memory storage of `Ziggurat`. Concluding we state that for larger values of  $\sigma$  the `Ziggurat` algorithm beats both inverse CDF and rejection sampling. Compared to Knuth-Yao, `Ziggurat` achieves almost the same speed but reduces the memory consumption by a factor of more than 400.

Figure 3 shows that we can beat `invCDF` in terms of speed and compete with Knuth-Yao. The reason for this is the simplicity of the algorithm. If many rectangles are stored, the rejection-probability gets very small. Likewise, the probability to sample an  $x$  in a right rectangle  $R_i^r$  gets very small. Therefore, the algorithm only samples a rectangle and afterwards samples a value within this rectangle, which can be done very fast.

As one can furthermore see in Figure 3, the discrete `Ziggurat` algorithm shows a large flexibility in regard to the speed-memory trade-off. For a small amount of memory (i.e. number of rectangles) it is quite slow, e.g. for  $\sigma = 32$  and 8 rectangles it obtains about 57,000 samples per second. For increasing memory allocation the speed of `Ziggurat(O)` increases. This statement holds for all values of  $\sigma$  we tested. As can be seen by the



**Fig. 4.** Discrete Ziggurat for different parameters  $\sigma$  (ZX denotes Ziggurat for  $\sigma = X$ ).

graphs, the speed of Ziggurat decreases for increasing number of rectangles. This was first surprising to us. Further analysis showed that this is due to the fact that with increasing number of rectangles (i.e. amount of allocated memory) the processor cannot keep the partition table in the fast caches, but has to obtain requested memory addresses from slower caches on demand. In addition, the large number of rectangles requires more bits to be sampled in a single step of the algorithm.

The trade-off provided by the Ziggurat-algorithms is indeed a property the other approaches do not share. InvCDF assigns every possible value to an interval on the  $y$ -axis. Consequently, one has to store at least the borders of the intervals. Decreasing the precision of the borders will decrease the memory consumption, but as well decrease the quality of the sampler. Increasing the precision or storing intermediate values, on the other hand, will not decrease the running time. The same happens to rejection sampling if the precision of the precomputed values is changed. Knuth-Yao stores for every element in the support the probability to sample this element. Decreasing the precision of the stored probabilities would (like for invCDF) decrease the quality of the sampler. While there might be efficient ways to store those values, there is a minimal amount of space required to store this information. Knuth-Yao as well as invCDF and rejection sampling therefore only provide a trade-off between quality and speed/memory consumption.

In Figure 4 we draw the time-memory trade-off for the Ziggurat algorithm for different values of  $\sigma$ . One can see that the performance of the Ziggurat algorithm decreases for larger  $\sigma$ . What is interesting in the graph is that the Ziggurat algorithm for  $\sigma = 10$  is slower for a large amount of rectangles than for  $\sigma = 32$ . This is puzzling as we cannot directly explain

the described behaviour. We want to state that during our experiments we saw quite large fluctuations for several runs of the algorithm. Maybe this could explain the better performance for  $\sigma = 32$  in comparison to  $\sigma = 10$ .

We also compared ZigguratO and Ziggurat in regard to speed.<sup>2</sup> The improvement rate increases up to 30% for a total memory of 320 bytes, then decreases to around 6% for 2080 bytes, until finally for 130KB and bigger there is no improvement. Overall, the described behaviour is not surprising since for increased memory the number of rectangles gets larger, so that the rejection area is very small. This leads to nearly no evaluations in the right sub-rectangles  $R_i^r$  and therefore to no computation of the straight line  $s$  (or even  $\bar{\rho}_\sigma$ ).

Additionally, we compared ZigguratO to ZigguratOP, which operates with at least 2176 bytes of memory. ZigguratOP is slower until about 2.5KB of memory, but then it beats ZigguratO with a speedup of up to 40%, until for memory larger than 262KB there seems to be no speedup at all. This behaviour is reasonable since the lookup-table requires more storage, but simultaneously affects the speed due to replacing  $\bar{\rho}_\sigma$  by table-lookups.

At last, we give insights on the time-split for our implementations Ziggurat and ZigguratO. We used the tool suite `Valgrind` with the tool `Callgrind` to obtain the measurements and analyzed them using the Callee Graph in the `Callgrind-GUI KCachegrind`. Figure 5 shows the percentages for both algorithms. We chose the most interesting sub-routines, i.e. the exponential function (called inside  $\bar{\rho}_\sigma$  in  $R_i^r$ ), the generation of random bits, the computation of the straight line  $s$  (in ZigguratO), and ‘other’ sub-routines. One can see that for a small amount of memory the computation of the exponential function takes most part of the running time, e.g. for 104 bytes (two rectangles) its computation consumes 80-90% of the total running time. As the memory increases, the rejection area gets smaller, i.e. the percentage of the right sub-rectangles  $R_i^r$  compared to their super-rectangles  $R_i$ . Thus, the number of integers sampled inside the  $R_i^r$ ’s decreases. Additionally, the exponential function has to be called less often. Nevertheless, the graphs show that the use of the straight line  $s$  decreases the use of the exponential function (or call to  $\bar{\rho}_\sigma$ ) in ZigguratO in comparison to Ziggurat considerably, while at the same time the computational complexity of  $s$  is not high (at most 6.77%).

---

<sup>2</sup> For additional Figures see Appendix A.2.

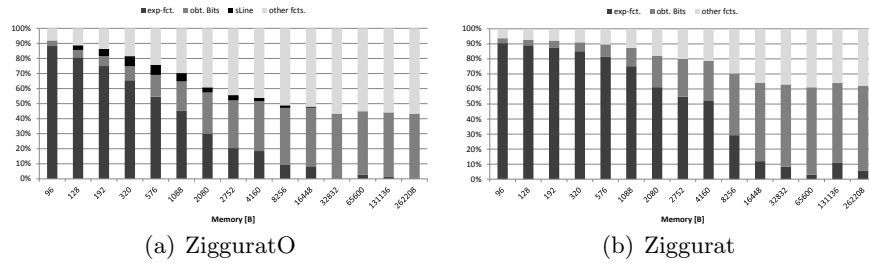


Fig. 5. Time-split of discrete Ziggurat with and without optimization

## References

- [AGHS12] Shweta Agrawal, Craig Gentry, Shai Halevi, and Amit Sahai. Discrete Gaussian Leftover Hash Lemma over infinite domains. Cryptology ePrint Archive, Report 2012/714, 2012. <http://eprint.iacr.org/>.
- [Ban93] W. Banaszczyk. New bounds in some transference theorems in the geometry of numbers. *Mathematische Annalen*, 296(1):625–635, 1993.
- [BGV12] Zvika Brakerski, Craig Gentry, and Vinod Vaikuntanathan. (Leveled) Fully homomorphic encryption without bootstrapping. In *Proceedings of the 3rd Innovations in Theoretical Computer Science Conference, ITCS '12*, pages 309–325, New York, NY, USA, 2012. ACM.
- [DN12] Léo Ducas and Phong Q. Nguyen. Faster Gaussian lattice sampling using lazy floating-point arithmetic. In Xiaoyun Wang and Kazue Sako, editors, *Advances in Cryptology – ASIACRYPT 2012*, volume 7658 of *Lecture Notes in Computer Science*, pages 415–432. Springer Berlin Heidelberg, 2012.
- [GD12] Steven D. Galbraith and Nagarjun C. Dwarakanath. Efficient sampling from discrete Gaussians for lattice-based cryptography on a constrained device. *Preprint*, 2012.
- [GGH12] Sanjam Garg, Craig Gentry, and Shai Halevi. Candidate multilinear maps from ideal lattices. Cryptology ePrint Archive, Report 2012/610, 2012. <http://eprint.iacr.org/>.
- [GPV08] Craig Gentry, Chris Peikert, and Vinod Vaikuntanathan. Trapdoors for hard lattices and new cryptographic constructions. In Richard E. Ladner and Cynthia Dwork, editors, *40th Annual ACM Symposium on Theory of Computing*, pages 197–206, Victoria, British Columbia, Canada, May 17–20, 2008. ACM Press.
- [LP11] Richard Lindner and Chris Peikert. Better key sizes (and attacks) for LWE-based encryption. In Aggelos Kiayias, editor, *Topics in Cryptology – CT-RSA 2011*, volume 6558 of *Lecture Notes in Computer Science*, pages 319–339, San Francisco, CA, USA, February 14–18, 2011. Springer, Berlin, Germany.
- [Lyu12] Vadim Lyubashevsky. Lattice signatures without trapdoors. In David Pointcheval and Thomas Johansson, editors, *Advances in Cryptology – EUROCRYPT 2012*, volume 7237 of *Lecture Notes in Computer Science*, pages 738–755, Cambridge, UK, April 15–19, 2012. Springer, Berlin, Germany.
- [MT00] George Marsaglia and Wai Wan Tsang. The Ziggurat method for generating random variables. *Journal of Statistical Software*, 5(8):1–7, 10 2000.
- [Pei10] Chris Peikert. An efficient and parallel Gaussian sampler for lattices. In Tal Rabin, editor, *Advances in Cryptology – CRYPTO 2010*, volume 6223

- of *Lecture Notes in Computer Science*, pages 80–97, Santa Barbara, CA, USA, August 15–19, 2010. Springer, Berlin, Germany.
- [Reg05] Oded Regev. On lattices, learning with errors, random linear codes, and cryptography. In *Proceedings of the thirty-seventh annual ACM symposium on Theory of computing*, STOC '05, pages 84–93, New York, NY, USA, 2005. ACM.
- [Sho] Victor Shoup. Number Theory Library (NTL) for C++. <http://www.shoup.net/ntl/>.
- [TLLV07] David B. Thomas, Wayne Luk, Philip H.W. Leong, and John D. Villasenor. Gaussian random number generators. *ACM Comput. Surv.*, 39(4), November 2007.
- [WHCB13] Patrick Weiden, Andreas Hülsing, Daniel Cabarcas, and Johannes Buchmann. Instantiating treeless signature schemes. Cryptology ePrint Archive, Report 2013/065, 2013. <http://eprint.iacr.org/>.

## A Appendix

In this Appendix we present the pseudocode for the discrete Ziggurat algorithm, give additional Figures in regard to our experimental results, and give the proof for Theorem 1 from Section 3.

### A.1 Pseudocode for Discrete Ziggurat

In Figure 6 we present the pseudocode for our implementation of the discrete Ziggurat algorithm of Section 2. In particular, we give the pseudocode for ZigguratO. From this, one obtains pseudocode for Ziggurat by removing lines 11–17, 19 and 20.

### A.2 Additional Figures regarding Results

In Figure 7 we present the rate of improvement of Ziggurat with optimization (ZigguratO) over Ziggurat without the straight line approach. For a small amount of memory, the improvement using the straight line approach is quite good (around 20-30% for memory usage between 128 and 576 bytes), while for larger memory, i.e. higher number of rectangles, the improvement vanishes due to nearly no rejection area.

Figure 8 shows the speed of ZigguratO and its corresponding variant ZigguratOP with precomputed lookup-table. ZigguratOP can perform only with memory larger or equal to 2176 bytes due to the size of the lookup-table. Thus, given a small amount of memory, it is not possible to apply ZigguratOP. But for available memory larger than 2.5KB ZigguratOP outperforms ZigguratO up to 40%.



**Algorithm 1:** ZigguratO**Input:**  $m, \sigma, \lfloor x_1 \rfloor, \dots, \lfloor x_m \rfloor, \bar{y}_0, \bar{y}_1, \dots, \bar{y}_m, \omega$ **Output:** number distributed according to a discrete Gaussian distribution

```

1 while true do
2    $i \stackrel{\$}{\leftarrow} \{1, \dots, m\}, s \stackrel{\$}{\leftarrow} \{-1, 1\}, x \stackrel{\$}{\leftarrow} \{0, \dots, \lfloor x_i \rfloor\};$ 
   // choose rectangle, sign and value
3   if  $0 < x \leq \lfloor x_{i-1} \rfloor$  then return  $sx$ ;
4   else
5     if  $x = 0$  then
6        $b \stackrel{\$}{\leftarrow} \{0, 1\};$ 
7       if  $b = 0$  then return  $sx$ ;
8       else continue;
9     else
10      // in rejection area  $R_i^r$  now
11       $y' \stackrel{\$}{\leftarrow} \{0, \dots, 2^\omega - 1\}, \bar{y} = y' \cdot (\bar{y}_{i-1} - \bar{y}_i);$ 
12      if  $\lfloor x_i \rfloor + 1 \leq \sigma$  then
13        // in concave-down case
14        if
15           $\bar{y} \leq 2^\omega \cdot sLine(\lfloor x_{i-1} \rfloor, \lfloor x_i \rfloor, \bar{y}_{i-1}, \bar{y}_i; x) \vee \bar{y} \leq 2^\omega \cdot (\bar{\rho}_\sigma(x) - \bar{y}_i)$ 
16          then return  $sx$ ;
17          else continue;
18        else if  $\sigma \leq \lfloor x_{i-1} \rfloor$  then
19          // in concave-up case
20          if
21             $\bar{y} \geq 2^\omega \cdot sLine(\lfloor x_{i-1} \rfloor, \lfloor x_i \rfloor, \bar{y}_{i-1}, \bar{y}_i; x-1) \vee \bar{y} > 2^\omega \cdot (\bar{\rho}_\sigma(x) - \bar{y}_i)$ 
22            then continue;
23            else return  $sx$ ;
24          else
25            if  $\bar{y} \leq 2^\omega \cdot (\bar{\rho}_\sigma(x) - \bar{y}_i)$  then return  $sx$ ;
26            else continue;
27          end
28        end
29      end
30    end
31  end
32 end

```

**Algorithm 2:**  $sLine(\lfloor x_{i-1} \rfloor, \lfloor x_i \rfloor, \bar{y}_{i-1}, \bar{y}_i; x)$ 

```

1 if  $\lfloor x_i \rfloor = \lfloor x_{i-1} \rfloor$  then return  $-1$ ;
2 Set  $\hat{y}_i = \bar{y}_i$  and  $\hat{y}_{i-1} = \begin{cases} \bar{y}_{i-1} & i > 1 \\ 1 & i = 1 \end{cases}$ 
3 return  $\frac{\hat{y}_i - \hat{y}_{i-1}}{\lfloor x_i \rfloor - \lfloor x_{i-1} \rfloor} \cdot (x - \lfloor x_i \rfloor)$ 

```

**Fig. 6.** The discrete Ziggurat algorithm with optimization (ZigguratO)

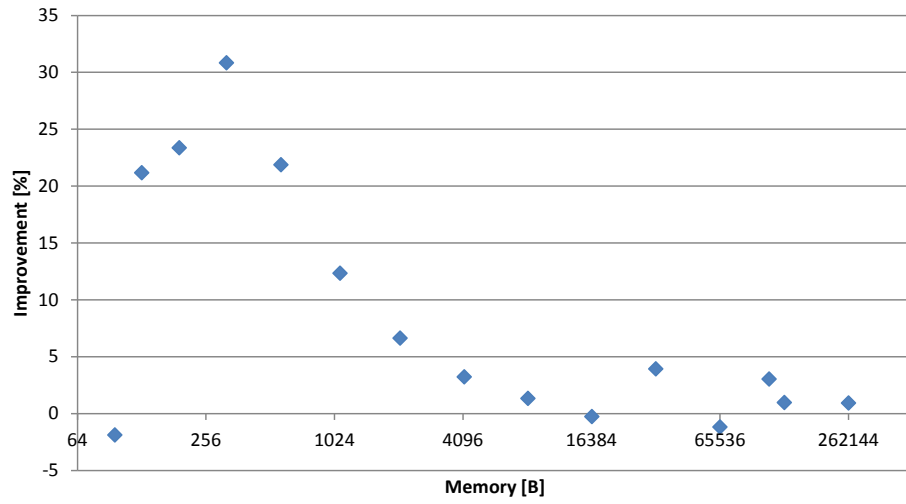


Fig. 7. Improvement rate of ZigguratO over Ziggurat

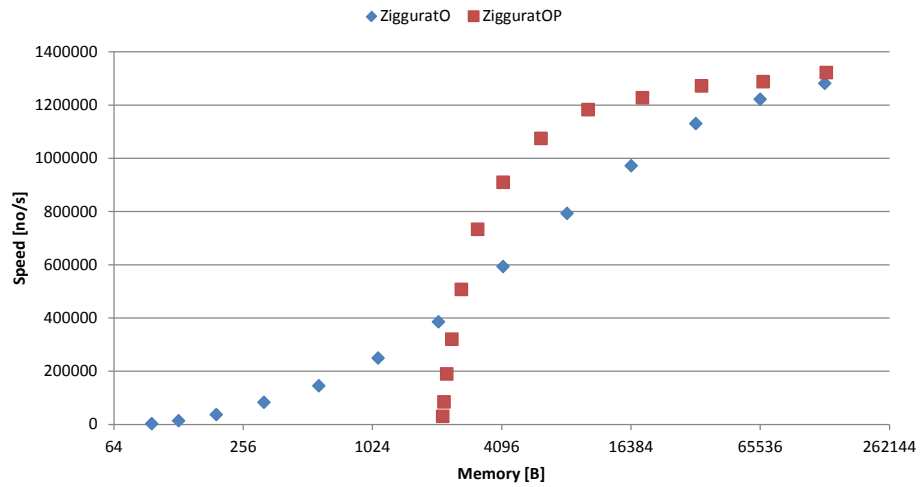


Fig. 8. Comparison of ZigguratO and ZigguratOP

### A.3 Quality of our Sampler: Proof of Theorem 1

In this section we present the proof of Theorem 1 from Section 3. The proof is divided into several parts: First we show that the distribution generated by the discrete Ziggurat algorithm is the target distribution  $\overline{D}_\sigma$  given below. Then, by introducing a number of intermediate distributions, we upper bound the statistical distance of the Gaussian distribution  $D_\sigma$  and the target distribution  $\overline{D}_\sigma$  in a consecutive way. This allows us to make statements about the quality of our sampler for given.

Let  $\overline{y}_i, \overline{h}_i, [x_i], B, B^+$  and  $B_0^+$  be defined for the parameters  $t, n, m, \omega$  as in Sections 2 and 3. For  $x \in B^+$  we define the function

$$\overline{P}(x) := \overline{y}_j + 2^{-\omega} \overline{h}_j \left[ \frac{\overline{\rho}_\sigma(x) - \overline{y}_j}{2^{-\omega} \overline{h}_j} + 1 \right],$$

where  $j$  is the index of the first rectangle containing points with  $x$ -coordinate  $x$ . Furthermore, we define the following distributions.

- Let  $D_\sigma$  be the discrete Gaussian distribution over the integers with mean 0 and parameter  $\sigma$ , i.e.  $\Pr[D_\sigma = x] \sim \rho_\sigma(x)$  for every  $x \in \mathbb{Z}$ .
- Let  $D_\sigma^+$  be the distribution defined by  $\Pr[D_\sigma^+ = x] \sim \begin{cases} \rho_\sigma(x) & \text{for } x \in \mathbb{Z}^+ \\ \frac{1}{2} & \text{for } x = 0. \end{cases}$
- Let  $\widehat{D}_\sigma^+$  be the distribution defined by  $\Pr[\widehat{D}_\sigma^+ = x] \sim \begin{cases} \rho_\sigma(x) & \text{for } x \in B^+ \\ \frac{1}{2} & \text{for } x = 0. \end{cases}$
- Let  $\overline{D}_\sigma^+$  be the distribution defined by  $\Pr[\overline{D}_\sigma^+ = x] \sim \begin{cases} \overline{P}(x) & \text{for } x \in B^+ \\ \frac{1}{2} \overline{P}(x) & \text{for } x = 0. \end{cases}$
- Let  $\overline{D}_\sigma$  be the distribution defined by  $\Pr[\overline{D}_\sigma = x] \sim \overline{P}(x)$  for  $x \in B$ .

Before we start with the calculation of the statistical distances, we show how these distributions correlate with our algorithm.

**Theorem 2.**  $\overline{D}_\sigma^+$  is the distribution returned by our algorithm without sampling a sign  $s \in \{-1, 1\}$  and  $\overline{D}_\sigma$  is the distribution returned by our algorithm.

*Proof.* In order to prove the first claim, we determine the probability  $\mathcal{P}_x$  that a given  $x \in B_0^+$  is chosen in the first iteration of the while-loop and not rejected. Since the probability of returning this  $x$  at all is proportional to  $\mathcal{P}_x$ , this shows the result. Let  $j \in \{1, \dots, m\}$  be such that  $x_{j-1} < x \leq x_j$ . If a rectangle  $R_i$  with  $i < j$  is sampled,  $x$  cannot be returned. If  $i > j$  and  $x$  is sampled in the rectangle,  $x$  will be returned

without sampling a  $y$ -coordinate. Finally, if  $i = j$  and  $x$  is sampled,  $x$  will be returned if the transformed  $y$  is at most  $\bar{\rho}_\sigma(x)$ . This leads to

$$\mathcal{P}_x = \frac{1}{m} \left( \sum_{i=j+1}^m \frac{1}{\lfloor x_i \rfloor + 1} + \frac{1}{\lfloor x_j \rfloor + 1} \Pr[2^{-\omega} h_j y + y_j \leq \bar{\rho}_\sigma(x)] \right)$$

with  $y \stackrel{\$}{\leftarrow} \{0, \dots, 2^\omega - 1\}$ . Since  $(\lfloor x_i \rfloor + 1)\bar{h}_i = S$  for every  $i \in \{1, \dots, m\}$ , this leads to

$$\begin{aligned} \mathcal{P}_x &= \frac{1}{m} \left( \sum_{i=j+1}^m \frac{\bar{h}_i}{S} + \frac{\bar{h}_j}{S} \Pr[2^{-\omega} \bar{h}_j y + \bar{y}_j \leq \bar{\rho}_\sigma(x)] \right) \\ &= \frac{1}{mS} (\bar{y}_j + \bar{h}_j \Pr[2^{-\omega} \bar{h}_j y + \bar{y}_j \leq \bar{\rho}_\sigma(x)]). \end{aligned}$$

Since the acceptance probability is given by

$$\begin{aligned} \Pr[2^{-\omega} \bar{h}_j y + \bar{y}_j \leq \bar{\rho}_\sigma(x)] &= \frac{|\{i \in \mathbb{N}_0 \mid 2^{-\omega} \bar{h}_j i + \bar{y}_j \leq \bar{\rho}_\sigma(x)\}|}{2^\omega} \\ &= 2^{-\omega} |\{i \in \mathbb{N}_0 \mid i \leq 2^\omega (\bar{\rho}_\sigma(x) - \bar{y}_j) / \bar{h}_j\}| \\ &= 2^{-\omega} \left\lfloor \frac{\bar{\rho}_\sigma(x) - \bar{y}_j}{2^{-\omega} \bar{h}_j} + 1 \right\rfloor, \end{aligned}$$

we obtain

$$\mathcal{P}_x = \frac{1}{mS} \left( \bar{y}_j + 2^{-\omega} \bar{h}_j \left\lfloor \frac{\bar{\rho}_\sigma(x) - \bar{y}_j}{2^{-\omega} \bar{h}_j} + 1 \right\rfloor \right) \sim \bar{P}(x).$$

The next step is the proof of the second statement. Note that the first statement reveals that one can sample from the distribution defined by our algorithm by sampling a sign and multiplying it to the output of  $\bar{D}_\sigma^+$ . This leads to

$$\Pr[\bar{D}_\sigma = x] = \begin{cases} \frac{1}{2} \Pr[\bar{D}_\sigma^+ = |x|] & \text{for } x \in B \setminus \{0\} \\ \Pr[\bar{D}_\sigma^+ = 0] & \text{for } x = 0 \end{cases} \sim \frac{1}{2} \bar{P}(x)$$

for  $x \in B$ . □

The statistical distance between an arbitrary distribution and the same distribution with finite support is exactly the probability of sampling an element in the cut-off part of the support. We state the special case for the distribution  $D_\sigma^+$  of this general statement.

**Lemma 3.** *With notation as above,*

$$\Delta(D_\sigma^+, \widehat{D}_\sigma^+) = \Pr[D_\sigma^+ > t\sigma].$$

*Proof.* Since for  $x \in B_0^+$ ,  $\Pr[\widehat{D}_\sigma^+ = x] > \Pr[D_\sigma^+ = x]$  and for  $x \in \mathbb{Z}_0^+ \setminus B_0^+$ ,  $\Pr[D_\sigma^+ = x] = 0$ , it follows that

$$\begin{aligned} 2\Delta(D_\sigma^+, \widehat{D}_\sigma^+) &= \sum_{x \in \mathbb{Z}_0^+} \left| \Pr[D_\sigma^+ = x] - \Pr[\widehat{D}_\sigma^+ = x] \right| \\ &= \sum_{x \in B_0^+} \left| \Pr[D_\sigma^+ = x] - \Pr[\widehat{D}_\sigma^+ = x] \right| + \sum_{x \in \mathbb{Z}_0^+ \setminus B_0^+} \left| \Pr[D_\sigma^+ = x] - \Pr[\widehat{D}_\sigma^+ = x] \right| \\ &= \sum_{x \in B_0^+} (\Pr[\widehat{D}_\sigma^+ = x] - \Pr[D_\sigma^+ = x]) + \sum_{x \in \mathbb{Z}_0^+ \setminus B_0^+} \Pr[D_\sigma^+ = x] \\ &= \sum_{x \in B_0^+} \Pr[\widehat{D}_\sigma^+ = x] - \sum_{x \in B_0^+} \Pr[D_\sigma^+ = x] + \Pr[D_\sigma^+ \in \mathbb{Z}_0^+ \setminus B_0^+] \\ &= \Pr[\widehat{D}_\sigma^+ \in B_0^+] - \Pr[D_\sigma^+ \in B_0^+] + \Pr[D_\sigma^+ \in \mathbb{Z}_0^+ \setminus B_0^+] \\ &= 1 - (1 - \Pr[D_\sigma^+ \in \mathbb{Z}_0^+ \setminus B_0^+]) + \Pr[D_\sigma^+ \in \mathbb{Z}_0^+ \setminus B_0^+] \\ &= 2 \Pr[D_\sigma^+ \in \mathbb{Z}_0^+ \setminus B_0^+] \\ &= 2 \Pr[D_\sigma^+ > t\sigma]. \end{aligned}$$

□

In order to get a concrete bound for the statistical distance, we need to bound the probability of the tail. We can use the special case of [Ban93], Lemma 1.5 for the one-dimensional case and  $L = \mathbb{Z}$ .

**Lemma 4.** *For every  $t \geq 1$ ,*

1.  $\rho_\sigma(\mathbb{Z} \setminus [-\sigma t, \sigma t]) < t e^{(1-t^2)/2} \cdot \rho_\sigma(\mathbb{Z})$  and
2.  $\Pr[D_\sigma^+ > t\sigma] < t e^{(1-t^2)/2}$ .

*Proof.* The first result is proved in [Ban93]. The second one follows since

$$\Pr[D_\sigma^+ > t\sigma] = \frac{\sum_{x > t\sigma} \rho_\sigma(x)}{\sum_{x \in \mathbb{Z}^+} \rho_\sigma(x) + 1/2} = \frac{2 \sum_{x > t\sigma} \rho_\sigma(x)}{2 \sum_{x \in \mathbb{Z}^+} \rho_\sigma(x) + 1} < t e^{(1-t^2)/2}.$$

□

The following result bounds the statistical distance of two distributions induced by functions  $f$  and  $g$ . It is true for any finite support, but for simplicity we will denote the support by  $B_0^+$ .

**Lemma 5.** Let  $f : B_0^+ \rightarrow \mathbb{R}_0^+$  and  $g : B_0^+ \rightarrow \mathbb{R}_0^+$  be non-negative functions and  $F, G$  be distributions with support  $B_0^+$  satisfying  $\Pr[F = x] \sim f(x)$  and  $\Pr[G = x] \sim g(x)$  for all  $x \in B_0^+$ , respectively. Then

$$\Delta(F, G) \leq \frac{1}{g(B_0^+)} \sum_{x \in B_0^+} |f(x) - g(x)|.$$

*Proof.* Note that  $\Pr[F = x] = \frac{f(x)}{f(B_0^+)}$  and  $\Pr[G = x] = \frac{g(x)}{g(B_0^+)}$  for every  $x \in B_0^+$ . The lemma follows since

$$\begin{aligned} 2\Delta(F, G) &= \sum_{x \in B_0^+} \left| \frac{f(x)}{f(B_0^+)} - \frac{g(x)}{g(B_0^+)} \right| \\ &= \frac{1}{g(B_0^+)} \sum_{x \in B_0^+} \left| f(x) - g(x) + \left( \frac{g(B_0^+)}{f(B_0^+)} - 1 \right) f(x) \right| \\ &\stackrel{\triangle}{\leq} \frac{1}{g(B_0^+)} \left( \sum_{x \in B_0^+} |f(x) - g(x)| + \sum_{x \in B_0^+} \left| \left( \frac{g(B_0^+)}{f(B_0^+)} - 1 \right) f(x) \right| \right) \\ &= \frac{1}{g(B_0^+)} \left( \sum_{x \in B_0^+} |f(x) - g(x)| + |g(B_0^+) - f(B_0^+)| \right) \\ &\stackrel{\triangle}{\leq} \frac{2}{g(B_0^+)} \sum_{x \in Z} |f(x) - g(x)|. \end{aligned}$$

□

Lemma 5 can be used to bound the error caused by the limited precision and the discrete sampling of  $y$ .

**Lemma 6.** With the definition of  $D_\sigma$  and  $D_\sigma^+$  from above,

$$\Delta(\widehat{D}_\sigma^+, \overline{D}_\sigma^+) < \frac{|B_0^+|}{\rho_\sigma(B^+) + \frac{1}{2}} (2^{-\omega+1} + 2^{-n}).$$

*Proof.* Let  $j \in \{1, \dots, m\}$  be such that  $x_{j-1} < x \leq x_j$ . Since for every  $x \in B_0^+$ ,

$$\begin{aligned} |\overline{P}(x) - \overline{\rho}_\sigma(x)| &= \left| \overline{y}_j + 2^{-\omega} \overline{h}_j \left[ \frac{\overline{\rho}_\sigma(x) - \overline{y}_j}{2^{-\omega} \overline{h}_j} + 1 \right] - \overline{\rho}_\sigma(x) \right| \\ &\stackrel{\triangle}{\leq} 2^{-\omega} \overline{h}_j \underbrace{\left| \left[ \frac{\overline{\rho}_\sigma(x) - \overline{y}_j}{2^{-\omega} \overline{h}_j} \right] - \frac{\overline{\rho}_\sigma(x) - \overline{y}_j}{2^{-\omega} \overline{h}_j} \right|}_{< 1} + 2^{-\omega} \overline{h}_j, \end{aligned}$$

we conclude that  $|\bar{P}(x) - \bar{\rho}_\sigma(x)| < 2^{-\omega}\bar{h}_j + 2^{-\omega}\bar{h}_j = 2^{-\omega+1}\bar{h}_j \leq 2^{-\omega+1}$ . Furthermore,  $|\bar{\rho}_\sigma(x) - \rho_\sigma(x)| \leq 2^{-n}$  for every  $x \in B_0^+$  and by the triangle inequality it follows that

$$|\bar{P}(x) - \rho_\sigma(x)| < 2^{-\omega+1} + 2^{-n}.$$

Applying Lemma 5 with

$$f : x \mapsto \begin{cases} \rho_\sigma(x) & \text{for } x > 0 \\ 1/2 & \text{for } x = 0 \end{cases}$$

and

$$g : x \mapsto \begin{cases} \bar{P}(x) & \text{for } x > 0 \\ 1/2 \cdot \bar{P}(0) & \text{for } x = 0 \end{cases}$$

leads to

$$\begin{aligned} \Delta(\hat{D}_\sigma^+, \bar{D}_\sigma^+) &\leq \frac{1}{g(B_0^+)} \sum_{x \in B_0^+} |f(x) - g(x)| \\ &\leq \frac{|B_0^+|}{\rho_\sigma(B^+) + \frac{1}{2}} (2^{-\omega+1} + 2^{-n}). \end{aligned}$$

□

The next result shows that sign-sampling does not change the statistical distance. It is quite natural if one thinks of the statistical distance as the size of the area between the density functions of the two distributions. Remember that  $D_\sigma^+ = |D_\sigma|$  and  $\bar{D}_\sigma^+ = |\bar{D}_\sigma|$ . Taking the absolute value means adding the reflected left half of the density function to the right half. Since both density functions are symmetric, the right half is doubled and it is easy to see that the size of the area between the distributions does not change.

**Lemma 7.** *With the definition of  $D_\sigma$  and  $\bar{D}_\sigma$  from above,*

$$\Delta(D_\sigma, \bar{D}_\sigma) = \Delta(D_\sigma^+, \bar{D}_\sigma^+).$$

*Proof.* Recall that one can sample from  $D_\sigma$  by sampling  $x \leftarrow D_\sigma^+$ , sampling a sign  $s \stackrel{\$}{\leftarrow} \{-1, 1\}$  and outputting  $sx$ . It is easy to see that

$$\Pr[D_\sigma = x] = \begin{cases} \frac{1}{2} \Pr[D_\sigma^+ = |x|] & \text{for } x \in \mathbb{Z} \setminus \{0\} \\ \Pr[D_\sigma^+ = 0] & \text{for } x = 0. \end{cases}$$

Likewise,

$$\Pr[\overline{D}_\sigma = x] = \begin{cases} \frac{1}{2} \Pr[\overline{D}_\sigma^+ = |x|] & \text{for } x \in B \setminus \{0\} \\ \Pr[\overline{D}_\sigma^+ = 0] & \text{for } x = 0. \end{cases}$$

Since  $\Pr[\overline{D}_\sigma = x] = 0$  for  $x \in \mathbb{Z} \setminus B$  and  $\Pr[\overline{D}_\sigma^+ = x] = 0$  for  $x \in \mathbb{Z}^+ \setminus B^+$ ,

$$\Pr[\overline{D}_\sigma = x] = \begin{cases} \frac{1}{2} \Pr[\overline{D}_\sigma^+ = |x|] & \text{for } x \in \mathbb{Z} \setminus \{0\} \\ \Pr[\overline{D}_\sigma^+ = 0] & \text{for } x = 0. \end{cases}$$

This leads to

$$\begin{aligned} 2\Delta(D_\sigma, \overline{D}_\sigma) &= \sum_{x \in \mathbb{Z}} |\Pr[D_\sigma = x] - \Pr[\overline{D}_\sigma = x]| \\ &= \sum_{\substack{x \in \mathbb{Z} \\ x \neq 0}} \left| \frac{1}{2} \Pr[D_\sigma^+ = |x|] - \frac{1}{2} \Pr[\overline{D}_\sigma^+ = |x|] \right| + \left| \Pr[D_\sigma^+ = 0] - \Pr[\overline{D}_\sigma^+ = 0] \right| \\ &= 2 \sum_{x \in \mathbb{Z}^+} \left| \frac{1}{2} \Pr[D_\sigma^+ = x] - \frac{1}{2} \Pr[\overline{D}_\sigma^+ = x] \right| + \left| \Pr[D_\sigma^+ = 0] - \Pr[\overline{D}_\sigma^+ = 0] \right| \\ &= \sum_{x \in \mathbb{Z}^+} \left| \Pr[D_\sigma^+ = x] - \Pr[\overline{D}_\sigma^+ = x] \right| \\ &= 2\Delta(D_\sigma^+, \overline{D}_\sigma^+). \end{aligned}$$

□

Now we can finally prove the main result. We add it for completeness.

**Theorem 1.** *The statistical distance between the discrete Gaussian distribution  $D_\sigma$  and the distribution  $\overline{D}_\sigma$  output by our algorithm is bounded by*

$$\Delta(D_\sigma, \overline{D}_\sigma) < te^{(1-t^2)/2} + \frac{|B_0^+|}{\overline{\rho}_\sigma(B^+) + \frac{1}{2}} (2^{-\omega+1} + 2^{-n}).$$

*Proof.* Lemmata 3, 4 and 6 show that

$$\Delta(D_\sigma^+, \overline{D}_\sigma^+) < te^{(1-t^2)/2} + \frac{|B_0^+|}{\overline{\rho}_\sigma(B^+) + \frac{1}{2}} (2^{-\omega+1} + 2^{-n}).$$

The rest follows since Lemma 7 states that  $\Delta(D_\sigma, \overline{D}_\sigma) = \Delta(D_\sigma^+, \overline{D}_\sigma^+)$ . □