

# Secure Random Sampling in Differential Privacy

Naoise Holohan and Stefano Braghin

IBM Research Europe  
Dublin – Ireland  
{naoise@,stefanob@ie.}ibm.com

**Abstract.** Differential privacy is among the most prominent techniques for preserving privacy of sensitive data, owing to its robust mathematical guarantees and general applicability to a vast array of computations on data, including statistical analysis and machine learning. Previous work demonstrated that concrete implementations of differential privacy mechanisms are vulnerable to statistical attacks. This vulnerability is caused by the approximation of real values to floating point numbers. This paper presents a practical solution to the finite-precision floating point vulnerability, where the inverse transform sampling of the Laplace distribution can itself be inverted, thus enabling an attack where the original value can be retrieved with non-negligible advantage.

The proposed solution has the advantages of being (i) mathematically sound, (ii) generalisable to any infinitely divisible probability distribution, and (iii) of simple implementation in modern architectures. Finally, the solution has been designed to make side channel attack infeasible, because of inherently exponential, in the size of the domain, brute force attacks.

**Keywords:** Differential Privacy · Random Numbers · Computational Complexity

## 1 Introduction

The generation of random samples from probability distributions is a well-studied field, and is an ever-present feature in programming languages and random number generators alike. Random sampling is especially important in the field of differential privacy, where specially-calibrated random noise is used to protect published data from unwanted inference. However, while the definition and analysis of differential privacy is typically viewed in the theoretical, real-valued world of mathematics, translating its rigorous guarantees to the floating-point world of computers requires special attention.

Work by Mironov in 2012 [22] was the first to point to significant vulnerabilities in floating point implementations of differential privacy, which allow for catastrophic destruction of the much-vaunted privacy guarantees. Using naïve sampling of the Laplace distribution, Mironov was able to exploit holes in the output space to reconstruct the original input from a finite list of candidate inputs. Being able to reconstruct a single value with certainty is a blatant breach of differential privacy, which should guarantee uncertainty on any query. In this paper, we further extend Mironov’s approach to present a novel attack on the Gaussian mechanism (Section 2.4).

It is well-known that the cardinality of the real line  $\mathbb{R}$  is the same as that of the unit interval on the real line,  $[0, 1]$ , which allows for precise sampling using the inverse transform method. For floating-point numbers however, there are many more numbers on the real line than there are in the unit interval, which results in the holes that Mironov used in the attack. Typical defences against this vulnerability eliminate these holes by limiting the output space, using such techniques as the “snapping mechanism” [22], and sampling from the discrete analogue of the distribution [14,6], but these require much more complex code to implement.

In this paper, we seek an alternative defence – that of computational complexity. The attack presented by Mironov relies on an attacker inverting the sampling procedure to test the feasibility of candidate inputs. Although random samples generated from a single random number (e.g., the Laplace distribution) are vulnerable to this attack, samples generated by two or more random numbers (e.g., the Gaussian distribution) are less vulnerable. By generating samples from multiple random numbers, we can render the attack sufficiently costly so as to be impractical on today’s computers.

We demonstrate how this approach can be implemented with ease for the Laplace and Gaussian distributions (both popular in differential privacy), and a large family of probability distributions. The approach can be extended to systems of reduced precision (e.g., single- and half- precision floating point), and allows for any desired level of complexity to be achieved.

Given that data is commonly stored as floating point numbers, we believe there is value in developing tools that reflect this reality. Equally, we are cognisant of the need to develop simple solutions to these complex tasks, to allow for easy implementation, adoption, generalisation and understanding of the techniques. As described in [8]:

Fast programs are seldom short, and short programs are likely to be slow. But it is also true that long programs are often not elegant and more error-prone. Short smooth programs survive longer and are understood by a larger audience.

The rest of the paper is organised as follows. Section 2 describe the required background information. The guiding principles to this paper are

then given in Section 3, with the specifics presented in Section 4. Example implementations are given in Section 5 and the protection provided is explored in Section 6. After reviewing the state of the art in Section 7, the contribution is summarised in Section 8.

## 2 Background

This section will give a brief outline of the vulnerability presented in [22], alongside an example attack implementation and its extension to attacking Gaussian sampling, and the current state-of-the-art in mitigating against the attack.

### 2.1 Floating point numbers

To begin, we give a very brief outline of floating point numbers. Floating point numbers take their inspiration from scientific notation to store a wide range of numbers in binary format. Double-precision floating points (also known as *doubles*) occupy 64 bits of storage, comprising 1 bit for the sign, 11 bits for the exponent and the remaining 52 bits for the fraction or mantissa.<sup>1</sup> The corresponding real number for a given double is given by

$$(-1)^{\text{sign}}(1.b_{51}b_{50}\dots b_0)_2 \times 2^{e-1023},$$

where  $b_0, \dots, b_{51} \in \{0, 1\}$  are the bits of the mantissa, and  $e \in \mathbb{N}$  is the non-negative integer exponent. This format allows for the representation of numbers between  $10^{-308}$  and  $10^{308}$ , with varying degrees of granularity.

The IEEE standard<sup>1</sup> gives an algorithm for addition, subtraction, multiplication, division, and square root and requires that implementations produce the same result as that algorithm.

The standard specifies that floating-point numbers are represented as base 2 fractions. For example, the value  $0.001_2$  represents the decimal value  $\frac{0}{2} + \frac{0}{4} + \frac{1}{8}$ . Nonetheless, not all decimal fractions can be represented exactly as binary fractions. Thus, the decimal floating-point numbers are approximated to the binary floating point. For example, in base-2,  $\frac{1}{10}$  is represented as the infinitely-repeating value  $0.0001100110011\dots$

The standard provides information about the error introduced by both the representation and operations that can be performed. For example, the addition (and subtraction) operation can be performed by converting the operands to the same exponent, and then summing (or subtracting) the mantissa. On the other hand, multiplications can be computed by multiplying the mantissa of both operands and adding the exponents.

<sup>1</sup> ANSI/IEEE Std 754–2019 <http://754r.ucbtest.org>

These operation introduce a rounding error, as the values need to be converted and as multiplications are performed. Moreover, the standard specifies various rounding modes, including round to nearest (where ties round to the nearest even digit in the required position, the most commonly used), round away from zero, round up (toward  $+\infty$ ), round down (toward  $-\infty$ ), and round toward zero (i.e., truncation).

The amount of error can be quantified [13] to a *machine epsilon* (generally denoted as  $E_{\text{mach}}$ ), that depends on the actual precision used in the representation. For example, single precision floating point (32 bits) have a machine epsilon equal to  $2^{-23} \approx 1.19 \times 10^{-7}$ , while double precision floating point numbers (64 bits) have a machine epsilon equal to  $2^{-52} \approx 2.22 \times 10^{-16}$ .

## 2.2 Random number sampling

When sampling random numbers from a probability distribution, the common starting point is to take a (many) random number(s) from the unit interval and transform them to sample from the given distribution. Although the unit interval can be sampled with increasing granularity closer to zero, the typical practice is to sample from a multiple of a small number. For example, the Python programming language returns a multiple of  $2^{-53}$  when sampling from the unit interval.

Inverse transform sampling is a popular method to sample from certain probability distributions, by taking the inverse of the Cumulative Distribution Function (CDF) or Cumulative Mass Function (CMF). For example, given a CDF  $F(x) : \mathbb{R} \rightarrow [0, 1]$  from which we want to sample, and whose inverse  $F^{-1}(x) : [0, 1] \rightarrow \mathbb{R}$  is known, then given  $U \sim \mathcal{U}(0, 1)$ , we have

$$\mathbb{P}(F^{-1}(U) \leq x) = \mathbb{P}(U \leq F(x)) = F(x),$$

for  $x \in \mathbb{R}$ , noting that  $\mathbb{P}(U \leq u) = u$  for a uniform variate  $U \sim \mathcal{U}(0, 1)$  and  $u \in [0, 1]$ .

## 2.3 Mironov attack

Noting that there are only  $2^{53}$  possible uniform variates to be generated and  $2^{64}$  floating point numbers, there is an immediate difficulty in covering the output space with the image of a transformation. Without additional randomness, it is impossible for the inverse transform sampling method to fully cover the real line. This was demonstrated in [22], where the author confirmed that gaps appear in the output space between the outputs of consecutive uniform variates. Mironov formulated an attack that could successfully reconstruct an entire database, with certainty, under specific conditions.

The CDF of the standard Laplace distribution is given by

$$F_{\text{Lap}}(x) = \begin{cases} \frac{1}{2}e^x & \text{if } x \leq 0, \\ 1 - \frac{1}{2}e^{-x} & \text{if } x > 0, \end{cases} \quad (1)$$

with which we can calculate its inverse:

$$F_{\text{Lap}}^{-1}(u) = (-1)^{\lfloor u \rfloor_1} \log(1 - 2|u - 0.5|), \quad (2)$$

where  $\lfloor \cdot \rfloor_k$  denotes rounding to the nearest multiple of  $k \in \mathbb{R}$ . A Laplace variate can then be generated using  $F_{\text{Lap}}^{-1}(U) \sim \text{Lap}(0, 1)$ , by first sampling a uniform variate  $U \sim U(0, 1)$ . We present an example algorithm for attacking a single value in a dataset in Algorithm 1.

**Input:** Attack target  $v$ , DP query  $Q$  using the Laplace mechanism with implementation (2), finite candidate set  $\mathcal{C}$

**Output:** Attack result  $c \in \mathcal{C}$

```

1 while  $|\mathcal{C}| > 1$  do
2    $q = Q(v)$ 
3   for  $c \in \mathcal{C}$  do
4      $u = \lfloor F_{\text{Lap}}(q - c) \rfloor_{2^{-p}}$ 
5     if  $F_{\text{Lap}}^{-1}(u) + c \neq q$  then
6       remove  $c$  from  $\mathcal{C}$ 
7     end
8   end
9 end
10 return remaining element  $c \in \mathcal{C}$ 

```

**Algorithm 1:** Example algorithm implementing the Mironov attack

Critical to this attack is the ability to compute the inverse of the sampling procedure (i.e., the CDF, given in (2), in the case of inverse transform sampling). The success of this attack is therefore independent of the precision of the uniform variate, and can be executed whenever the sampling function can be inverted.

## 2.4 Gaussian attack

We now show an extension of the Mironov attack to the Gaussian distribution, one which we believe to be novel. This attack can be performed on Gaussian variates sampled using the popular Box-Muller transform [4]. Given two uniform variates  $U_1, U_2 \sim U(0, 1)$ , the Box-Muller transform returns two independent Gaussian samples  $N_1, N_2 \sim N(0, 1)$  as follows:

$$N_1 = \sqrt{-2 \log(1 - U_1)} \cos(2\pi U_2), \quad (3a)$$

$$N_2 = \sqrt{-2 \log(1 - U_1)} \sin(2\pi U_2). \quad (3b)$$

Knowing both  $N_1$  and  $N_2$ , we can recover  $U_1$  and  $U_2$  as follows:

$$U_1 = 1 - e^{-\frac{N_1^2 + N_2^2}{2}}, \quad (4a)$$

$$U_2 = \frac{1}{2\pi} \left( \arctan \left( \frac{N_2}{N_1} \right) + \pi \mathbb{1}_{\{N_1 < 0\}} \right). \quad (4b)$$

The same inversion technique as described in Algorithm 1 can be used to eliminate candidates and determine the true input, noting that  $q$  and  $u$  in Lines 2, 4 and 5 will be vectors of two values each.

Getting two values from a mechanism utilising the Box-Muller transform can be done by executing the same query twice in succession. Ensuring both values are from the same Box-Muller operation, and determining which is associated with the  $\cos$  and which is associated with the  $\sin$  can be done by examining the source code and executing a simple timing attack. It is a common implementation to return the first variate to the user, while caching the second variate to be returned the next time the function is called [24]. The calculation overhead in the first step can be measured to determine the phasing to implement the attack.

This attack can be mitigated against by discarding one of the variates, or by using both variates in an output of  $\frac{1}{\sqrt{2}}(N_1 + N_2)$ .

## 2.5 Existing defences

To mitigate against the attack, Mironov proposed the *snapping mechanism* to sample a noisy output with Laplace-like additive noise. The snapping mechanism involves snapping the noisy output to the nearest factor of the scale of the noise, resulting in a significant reduction in granularity of the output. For example, given a privacy budget  $\epsilon = 2^{-5}$ , the only outputs will be multiples of  $2^5 = 32$ . Alongside the reduction in granularity, the snapping mechanism is also cumbersome to implement, requiring a custom approach to sampling from the unit interval, as well as clipping and rounding operations.

Another approach is to sample from the discrete analogue of the desired distribution, thereby giving more control to the user on the discretisation of the output [14, 6]. [14] in particular offers solutions for the Laplace and Gaussian distributions, but requires complex sampling procedures. It is not clear if it is simple, or possible, to adapt their solution to other probability distributions.

While there are many algorithms to sample from the discrete Gaussian distribution [6, 18], all common random number generator packages produce floating-point Gaussian variates, making the implementation of such solutions more complex.

### 3 General principles

In this section we present the general principles of the proposed solution. We first make the following assumptions in analysing the protections provided by our approach:

**Assumption 1** *Mechanism source code is public.*

**Assumption 2** *Uniform variates are generated using a cryptographically-secure pseudorandom number generator (CSPRNG).*

**Assumption 3** *Uniform variates are generated as multiples of  $2^{-p}$  for some fixed  $p \in \mathbb{N}_{>0}$ .*

Hence, for  $U \sim \mathcal{U}(0, 1)$ , we have

$$U \in [0, 1) \cap (2^{-p}\mathbb{N}). \quad (5)$$

Assumption 1 ensures that we cannot achieve security of random number generation through obscurity, and that the protections are an inherent characteristic of the system. Assumption 2 ensures that an attacker cannot predict forthcoming uniform variates, as can be done with some standard Pseudo-Random Number Generators (PRNGs) (e.g., the Mersenne-Twister PRNG). Assumption 3 simplifies some of the analysis of this paper, and aligns with random number generation in many programming languages.

For clarity, we define the set of uniform variates as follows for a given precision  $p$ .

**Definition 1 (Uniform variates).** *Given  $p \in \mathbb{N}$ , we define*

$$\mathbb{U} = [0, 1) \cap (2^{-p}\mathbb{N}),$$

*the set of all multiples of  $2^{-p}$  in the half-open unit interval  $[0, 1)$ .*

*This denotes the set of uniform variates in many programming languages. For example, in the Python programming language,  $p = 53$ .*

Although the holes in the output space of the random number generation provide the key to the attack as shown in Section 2.3, it relies on the invertibility of the sampling procedure to execute efficiently. If the sampling procedure can be made non-injective (and consequently, non-invertible), then any sampled variate could be the result of more than one combination of uniform variate(s). Therefore, satisfying the non-equality condition of Line 5 in Algorithm 1 would require searching through all possible uniform variate combinations that approximately give equality in Line 5 (up to floating point rounding errors).

We can flesh out the details of this using the Gaussian distribution as an example.




Fig. 1: Level curves of (6) for given values of  $N_1$ .

*Example 1 (Gaussian distribution).* Consider samples from the Gaussian distribution, as given in Section 2.4. Although it is possible to reconstruct  $U_1$  and  $U_2$  knowing both outputs from the Box-Muller transform, we can show that knowing a single Box-Muller output is not sufficient to execute the Mironov attack in a single step. Suppose  $N_1$  is known to have originated from (3a), and we wish to determine  $U_1$  and  $U_2$  from which it came. We can write

$$U_1(U_2) = 1 - e^{-\frac{N_1^2}{2 \cos^2(2\pi U_2)}}, \quad (6)$$

allowing us to solve for  $U_1$  given  $U_2$ . We can similarly write  $U_2$  as a function of  $U_1$ .

As shown in Algorithm 1, we need only find a single pair  $U_1, U_2 \in \mathbb{U}^2$  that gives equality in Line 5 in order to retain that candidate as feasible. On the real line, any pair  $U_1, U_2 \in [0, 1)$  satisfying (6) will give equality for Line 5. However, due to floating point rounding errors and the limited precision of values  $U \in \mathbb{U}$ , we need to search through  $\mathbb{U}^2$  to find values that satisfy the equality down to the bit of least precision. Figure 1 shows the level curves of (6), illustrating the paths through the grid of  $(U_1, U_2) \in \mathbb{U}^2$  that would be traversed in a typical execution of the attack.

This example prompts two questions:

1. Can we extend the protection offered to Gaussian variates, beyond that of searching through  $2^p$  uniform variates?
2. Can this approach be adopted for other probability distributions? Particularly, ones with injective sampling procedures (e.g., the Laplace distribution)?



If we can use our choice of  $n \in \mathbb{N}$  uniform variates to generate a single sample from a given distribution, then a brute-force attack (similar to Example 1) would require on the order of  $2^{p(n-1)}$  ‘checks’. In essence, it would require fixing  $(n - 1)$  of the variates, and running an inversion attack (e.g., using (6) in the case of the Gaussian distribution) on the  $n$ th variate. Importantly, for a linear increase in the time to sample a variate, the cost of executing the attack will increase exponentially.

Additionally, this approach can be extended to systems of reduced precision. Single- and half-precision systems use only 32 and 16 bits respectively to represent a floating point number, compared to the 64 bits allocated for double-precision floating point numbers. For example, Graphical Processing Units (GPUs) typically use lower-precision floating point numbers in their calculations, making attacks such as Example 1 more feasible.

In implementing such an approach, we can take some inspiration from the subset sum problem in computer science. The subset sum problem asks, given a multiset  $A$  of positive integers ( $a \in \mathbb{N}$  for each  $a \in A$ ), and a target  $T \in \mathbb{Z}$ , is there a subset  $A' \subseteq A$  such that  $\sum_{a \in A'} a = T$ ? It is known that the subset sum problem is NP-complete [10], making it intractable to solve when the set  $A$  is large.

Luckily, there is a branch of probability theory that we can exploit in this context – infinite divisibility.

## 4 Divisibility of probability distributions

Our approach relies on simple and effective ways to generate random numbers from a given distribution using many uniform variates. As explained in Section 3, simply increasing the number of random bits used to generate a random sample from a distribution is not sufficient to protect against inversion attacks. Instead, we seek to increase the number of uniform variates used to generate a single random sample.

A probability distribution  $P$  is *divisible*, if there exists a distribution  $Q$  such that, given independent  $X_1, X_2 \sim Q$ , we have  $X_1 + X_2 \sim P$ .  $P$  is *infinitely divisible* if, for any  $n \in \mathbb{N}$ , there exists a distribution  $R_n$  such that independent  $X_i \sim R_n$  for each  $i \in [n]$  satisfy  $\sum_{i=1}^n X_i \sim P$ .

Probability distributions that are infinitely divisible therefore allow the greatest flexibility in increasing the complexity of this defence. Fortunately, the two most popular distributions used in differential privacy are infinitely divisible.

### 4.1 Preliminaries

For every  $i \in \mathbb{N}$ , we adopt the following notation for these common distributions:

- $N_i \sim \mathcal{N}(0, 1)$  is a collection of independent (standard) Gaussian random variables, and,
- $U_i \sim \mathcal{U}(0, 1)$  is a collection of independent uniform random variables on the unit interval.

Formulations of the probability distributions used in this section are given in Appendix A,

For any  $n \in \mathbb{N}$ , we let  $[n] = [1, n] \cap \mathbb{N}$  denote the positive integers up to and including  $n$ .

We note that when sampling uniformly from the unit interval, random number generators typically sample from  $[0, 1)$  (i.e., excluding 1). This range is undesirable when computing  $\log(U_1)$ , and instead we use  $\log(1 - U_1)$ , since  $1 - U_1 \in (0, 1]$  lies within the domain of logarithms.

Finally, we limit our analysis to standard distributions, those with zero location and unit scale parameters. These standard distributions can be scaled and translated as appropriate to get a distribution of a specific mean and variance (see Appendix A).

## 4.2 Gaussian distribution

It is well-known that the Gaussian distribution is infinitely divisible. Given that the sum of two Gaussians is Gaussian ( $N_1 \pm N_2 \sim \mathcal{N}(0, 2)$ ), we can generate a standard Gaussian from  $n > 0$  standard Gaussians:

$$\frac{1}{\sqrt{n}} \sum_{i=1}^n N_i \sim \mathcal{N}(0, 1). \quad (7)$$

This property of divisibility allows for the sampling of Gaussians with ease for any  $n \in \mathbb{N}$ . Additionally, decomposing any other probability distribution as a function of Gaussians will allow for the subsequent decomposition into any finite number of Gaussians.

## 4.3 Laplace distribution

The Laplace distribution is infinite divisible using the gamma distribution.

**Proposition 1 (Laplace divisibility [20]).** *Given an integer  $n \geq 1$ , and  $X_i, Y_i \sim \Gamma(\frac{1}{n}, 1)$  for each  $i \in [n]$ , then*

$$\sum_{i=1}^n (X_i - Y_i) \sim \text{Lap}(0, 1).$$

The case of  $n = 1$  gives the decomposition of the Laplace distribution as the difference of two exponential distributions.

**Corollary 1.** Given  $E_1, E_2 \sim \text{Exp}(1)$ ,

$$E_1 - E_2 \sim \text{Lap}(0, 1). \quad (8)$$

*Proof.* Noting that  $\Gamma(1, 1) \sim \text{Exp}(1)$  completes the proof.

The special case of  $n = 2$  also deserves special attention.

**Corollary 2.** Given independent samples of the standard Gaussian distribution,  $N_1, N_2, N_3, N_4 \sim \mathcal{N}(0, 1)$ ,

$$\frac{1}{2} (N_1^2 - N_2^2 + N_3^2 - N_4^2) \sim \text{Lap}(0, 1). \quad (9)$$

*Proof.* The chi-squared distribution with 1 degree of freedom is related to the gamma distribution (given  $X \sim \chi^2(1)$ , then  $\frac{1}{2}X \sim \Gamma(\frac{1}{2}, 1)$ ). Furthermore, it is known that a squared sample from the Gaussian distribution is chi-squared ( $N_1^2 \sim \chi^2(1)$ ). Combining these two observations completes the proof.

The Laplace distribution can be decomposed in many other ways, as given in [20, Table 2.3], one of which is of particular interest in our context:

**Proposition 2.** Given independent samples of the standard Gaussian distribution,  $N_1, N_2, N_3, N_4 \sim \mathcal{N}(0, 1)$ ,

$$N_1 N_2 - N_3 N_4 \sim \text{Lap}(0, 1). \quad (10)$$

*Proof.* We note that  $Z = N_1 N_2$  has a characteristic function  $\phi_Z(t) = (1 + t^2)^{-\frac{1}{2}}$ . This gives the characteristic function for (10) of

$$\phi_{N_1 N_2 - N_3 N_4}(t) = \phi_Z(t) \phi_Z(-t) = (1 + t^2)^{-1},$$

which corresponds to the standard Laplace distribution.

We therefore have two methods to sample from the Laplace distribution as a composition of four standard Gaussians. Each Gaussian can be further decomposed using (7).

## 5 Sampling implementations

There are many ways to sample from the Gaussian distribution [24]. We restrict our analysis to the Box-Muller transform [4], owing to its simplicity, compactness, exactness and ubiquity among programming languages.

The polar method [3] of the Box-Muller transform was developed to avoid the costly  $\sin$  and  $\cos$  calculations. Other methods have been proposed to

eliminate the similarly costly  $\log$  and  $\sqrt{\cdot}$  calculations [2,5]. These modifications are typically reserved for low-level languages or machine code, because the added complexity can slow computation time in higher-level languages [1]. Given our focus on high-level languages, such as Python, analysis of these variants is beyond the scope of this paper.

In this section, we offer simple implementation examples in Python, using standard libraries. We hope these examples will be sufficient to enable the reader to port the code to their desired language.

### 5.1 Gaussian sampling

As noted in Section 2.4, a naïve implementation of the Box-Muller transform in sampling from the Gaussian distribution presents vulnerabilities. The question therefore arises of how best to protect Gaussian samples from the Mironov attack, and there are two immediate considerations:

1. Discard one of the pair of Box-Muller samples, or
2. Make use of both samples, noting that  $\frac{1}{\sqrt{2}}(N_1 + N_2) \sim \mathcal{N}(0, 1)$ .

Both of these choices are mathematically identical (albeit with the outputs  $\frac{\pi}{4}$  out of phase<sup>2</sup>), but it may be a more elegant implementation using standard libraries to opt for the second option. Extending this choice to generating Gaussian samples from multiple Box-Muller draws, using the infinite divisibility described in Section 4.2, we use  $2n$  samples<sup>3</sup> from the Box-Muller transform, and divide the result by  $\sqrt{2n}$ . This can be implemented in Python as follows, after importing the `math` and `random` libraries:

---

```
1 sum(gauss(0, 1) for i in range(2 * n)) / sqrt(2 * n)
```

---

Readers are encouraged to discard the first result from Box-Muller before sampling in this way, to ensure no carry-over of uniform variates from a previous draw, which may have already been successfully attacked.

Importantly, the time taken in the generation of a standard Gaussian variate from  $n$  uniform variates increases linearly with  $n$ , whereas the attack complexity as described in Section 3 increases exponentially in  $n$ , given as  $2^{p(n-1)}$ . A small increase in computation overhead to sample the variates gives an exponential increase in attack complexity.

### 5.2 Laplace sampling

We ignore the decomposition of the Laplace distribution into the difference of two exponential variates, and skip straight to the decomposition into Gaussians.

---

<sup>2</sup> Since  $\cos(\theta) + \sin(\theta) = \sqrt{2} \cos(\theta - \frac{\pi}{4})$

<sup>3</sup> We need only use  $n$  samples for sampling procedures that do not share uniform variates between executions, for example the `normalvariate` method in Python's `random` library, which uses the Kinderman-Monahan sampling procedure [19].

As such, we are interested in the forms given in Corollary 2 and Proposition 2. In both these cases, we can safely use two samples from Box-Muller to generate the Laplace sample, since squaring-and-summing or multiplying Box-Muller samples breaks the invertibility given in Section 2.4. Empirical analysis has shown (10) to be the faster option, as (9) requires two additional multiplication operations. In Python, this can be implemented as follows after importing the random library:

---

```
1 gauss(0, 1) * gauss(0, 1) - gauss(0, 1) * gauss(0, 1)
```

---

However, doing so will introduce some redundancy in the uniform variates, which presents an opportunity for those willing to write their own sub-routines.

**Theorem 1.** *Given independent uniform variates  $U_1, U_2, U_3, U_4 \sim \mathcal{U}(0, 1)$ ,*

$$\log(1 - U_1) \cos(\pi U_2) + \log(1 - U_3) \cos(\pi U_4) \sim \text{Lap}(0, 1). \quad (11)$$

*Proof.* Firstly with (9), we can write the difference of squared outputs from the Box-Muller transform as follows,

$$N_1^2 - N_2^2 = -2 \log(1 - U_1) (\cos^2(2\pi U_2) - \sin^2(2\pi U_2)) \quad (12)$$

$$= -2 \log(1 - U_1) \cos(4\pi U_2), \quad (13)$$

using the double-angle formula  $\cos(2\theta) = \cos^2(\theta) - \sin^2(\theta)$ .

Also, since  $-\cos(4\pi U_2) \sim \cos(\pi U_2)$  when  $U_2 \sim \mathcal{U}(0, 1)$ , owing to the periodicity of  $\cos$ , we get the desired result. Analysis of (10) gives a similar representation.

In Python, this can be implemented as follows after importing the math and random libraries:

---

```
1 log(1 - random()) * cos(pi * random()) + log(1 - random()) * cos(pi *
  random())
```

---

**Remark 1:** We observe that  $U_2 \in [0, 1)$  in a floating-point environment, which results in an asymmetric output for  $\cos(\pi U_2) \in (-1, 1]$ . We can eliminate this asymmetry by using the first bit of randomness of  $U_2$  to sample the sign, and then using the remaining bits to sample the magnitude, giving

$$\cos(\pi U_2) \sim (-1)^{\lfloor U_2 \rfloor_1} \cos\left(\pi \left(U_2 \bmod \frac{1}{2}\right)\right)$$

in real number arithmetic, with an output space of  $[-1, 1] \setminus \{0\}$  when  $U_2 \in [0, 1)$ . This can be implemented in Python as follows:

---

```
1 u2 = random()
2 copysign(cos(pi * (u2 % 0.5)), u2 - 0.5)
```

---

**Remark 2:** The representation in Theorem 1 brings two benefits over sampling naïvely with Box-Muller: (i) reduced redundancy in the  $\cos$  argument (using Box-Muller, the first two random bits of  $U_2$  and  $U_4$  are redundant), and, (ii) greater computational efficiency (empirical evidence suggests a halving in the time taken to sample).

### 5.3 Choosing $n$

In choosing the number,  $n$ , of uniform variates to use in sampling a single Laplace or Gaussian variate, we take inspiration from current standards in cryptography. The Advanced Encryption Standard (AES) supports key sizes of 128, 192 and 256 bits, corresponding with search spaces of sizes  $2^{128}$ ,  $2^{192}$  and  $2^{256}$  respectively. Thus, it will take  $O(\text{key size})$  steps for a polynomial adversary to enumerate, and test, all possible keys.

In order to achieve a similar search space size for our application, assuming a precision of  $p = 53$  (as in Python), we require  $n = 4, 5$  and  $6$  respectively. We consider it sufficient to use  $n = 4$  in most applications, corresponding to the implementation of Theorem 1. This standard has been adopted in the implementation of the Laplace, Gaussian and other mechanisms<sup>4</sup> in the `diffprivlib` open source library [16].

The effect of  $n$  on the time it takes to sample from the Laplace distribution is shown in Figure 2. Despite using four times as many uniform variates, the implementation of Theorem 1 only takes twice as long as the naïve sampling with a single uniform variate. Using the native `math` and `random` libraries is costly for larger  $n$ , owing to Python’s slow for loops. Numpy’s C codebase [25] allows for fast computation even for large  $n$ , and can be further leveraged to produce multiple samples in parallel, with superior per-sample computation time than naïve sampling. The code for these simulations is given in Appendix B.

## 6 Gaussian attack complexity

We briefly revisit the brute force attack on the Gaussian distribution in Example 1 to highlight the robustness of the proposed approach. Given a single Gaussian variate, we know from Example 1 that we can write  $U_1$  as a function of  $U_2$ . This allows us to estimate a lower bound on the number of checks required on a given variate  $N_1$ , since we know that  $U_1, U_2 \in \mathbb{U} \subset [0, 1)$ . We stress that this is only a lower bound, as rounding errors in floating point arithmetic would typically require the checking of adjacent uniform variates.

<sup>4</sup> <https://github.com/IBM/differential-privacy-library/tree/main/diffprivlib/mechanisms>




Fig. 2: Mean execution time for sampling a single standard Laplace variate, in microseconds ( $\mu$ s), over 5 000 000 runs in Python.

For example, using (6), given  $U_2 \in [0, 1)$ , we have

$$U_1 \in \left[1 - e^{-\frac{N_1^2}{2}}, 1\right),$$

allowing us to reduce the search space to  $U_1 \in \mathbb{U} \cap \left[1 - e^{-\frac{N_1^2}{2}}, 1\right)$ . We can therefore approximate the number of checks as  $e^{-\frac{N_1^2}{2}} 2^p$ .

Knowing that  $N_1 \sim \mathcal{N}(0, 1)$ , we can get the expected number of checks as follows:

$$\mathbb{E}[\text{checks}] = 2^p \int e^{-\frac{x^2}{2}} f(x) dx = 2^p C \int e^{-\frac{(\sqrt{2}x)^2}{2}} dx = 2^p \frac{C}{\sqrt{2}} \int e^{-\frac{x^2}{2}} dx \quad (14)$$

$$= \frac{2^p}{\sqrt{2}} = 2^{p-\frac{1}{2}}, \quad (15)$$

where  $f(x)$  denotes the probability density function of the standard Gaussian distribution  $\mathcal{N}(0, 1)$  and  $C$  is the normalisation constant (see Appendix A.2).

This demonstrates the robustness of the approach, with (15) confirming that the complexity of the attack only decreases by a small constant from the theoretical limit. While outlying values lead to smaller search spaces, their vanishingly small probabilities of being sampled result in a small aggregate impact.

## 7 Related work

The original attack on the Laplace mechanism, which also inspired this work, was first presented in [22]. Subsequent work on the subject in [12], where the authors analyse the situation of infinite-precision semantic at the implementation level. The authors present results arguing that in general there are violations of the differential privacy property, proposing a variation of the differential privacy definition leading to a degradation of the privacy level.

Thorough examination of randomness required for differential privacy is presented in [11], where the authors analyse various techniques for random number generation, presenting a picture of strengths and limitations of the commonly used sources of randomness.

Other work in [9] and [23] analyse theoretical and practical limitations of the source of randomness used in differential privacy. In particular [9] presents an cryptographic requirements for real randomness, and describe the analogies existing in differential privacy.

In [12] the authors show the violation of differential privacy property caused by the approximation introduced by the finite-precision representation of continuous data. The authors also present the conditions under which limited by acceptable privacy guarantees can be provided, under only a minor degradation of the privacy level.

[15] describes several different kinds of covert-channel attacks for differential privacy frameworks. The authors present possible countermeasures with particular focus on one specific solution based on a new primitive called *predictable transactions*.

[7] presents a systematic study of a fundamental limitation of locally differentially private protocols with respect to their vulnerability to adversarial manipulation. The authors also present a solution to provide increased security via a protocol that deploys local differential privacy and reinforces it with cryptographic techniques.

Finally, [17] presents a study similar to the one first introduced in [22], but concentrate on the exponential mechanism of McSherry and Talwar [21].

## 8 Conclusion

In this paper we have presented an alternative defence to a particular floating point vulnerability in differential privacy. The solution we presented, using the infinite divisibility of probability distributions, is simple to understand, quick and easy to implement, difficult to attack, and generalisable to different probability distributions and system precisions. Of particular interest is the ability to generate samples from the Laplace distribution in a single statement of code (Section 5.2), with strong attack guarantees. We



believe this is an important contribution to the literature on differential privacy in mitigating some of the risks associated with operating in a floating point environment.

**Acknowledgments** The authors wish to thank David Malone (Hamilton Institute, Maynooth University) for useful discussions at the beginning of this project.

## References

1. Ahrens, J.H., Dieter, U.: Computer methods for sampling from the exponential and normal distributions. *Commun. ACM* **15**(10), 873–882 (Oct 1972). <https://doi.org/10.1145/355604.361593>
2. Ahrens, J.H., Dieter, U.: Efficient table-free sampling methods for the exponential, Cauchy, and normal distributions. *Commun. ACM* **31**(11), 1330–1337 (Nov 1988). <https://doi.org/10.1145/50087.50094>
3. Bell, J.R.: Algorithm 334: Normal random deviates. *Commun. ACM* **11**(7), 498 (Jul 1968). <https://doi.org/10.1145/363397.363547>
4. Box, G.E.P., Muller, M.E.: A note on the generation of random normal deviates. *The Annals of Mathematical Statistics* **29**(2), 610 – 611 (1958). <https://doi.org/10.1214/aoms/1177706645>
5. Brent, R.P.: Fast normal random number generators on vector processors. Tec. Rep. TR-CS-93-04, Department of Computer Science, The Australian National University, Canberra, 0200 ACT, Australia (1993)
6. Canonne, C.L., Kamath, G., Steinke, T.: The discrete Gaussian for differential privacy. In: Larochelle, H., Ranzato, M., Hadsell, R., Balcan, M.F., Lin, H. (eds.) *Advances in Neural Information Processing Systems*. vol. 33, pp. 15676–15688. Curran Associates, Inc. (2020), <https://proceedings.neurips.cc/paper/2020/file/b53b3a3d6ab90ce0268229151c9bde11-Paper.pdf>
7. Cheu, A., Smith, A., Ullman, J.: Manipulation attacks in local differential privacy. arXiv preprint arXiv:1909.09630 (2019)
8. Devroye, L.: *Non-uniform random variate generation*. Springer-Verlag, New York (1986)
9. Dodis, Y., López-Alt, A., Mironov, I., Vadhan, S.: Differential privacy with imperfect randomness. In: Safavi-Naini, R., Canetti, R. (eds.) *Advances in Cryptology – CRYPTO 2012*. pp. 497–516. Springer Berlin Heidelberg, Berlin, Heidelberg (2012)
10. Garey, M.R., Johnson, D.S.: *Computers and intractability: A guide to the theory of NP-completeness*. W. H. Freeman and Company, New York (1979)
11. Garfinkel, S.L., Leclerc, P.: Randomness concerns when deploying differential privacy. In: *Proceedings of the 19th Workshop on Privacy in the Electronic Society*. p. 73–86. WPES’20, Association for Computing Machinery, New York, NY, USA (2020). <https://doi.org/10.1145/3411497.3420211>
12. Gazeau, I., Miller, D., Palamidessi, C.: Preserving differential privacy under finite-precision semantics. *Theoretical Computer Science* **655**, 92 – 108 (2016). <https://doi.org/10.1016/j.tcs.2016.01.015>, quantitative Aspects of Programming Languages and Systems (2013-14)

13. Goldberg, D.: What every computer scientist should know about floating-point arithmetic. *ACM computing surveys (CSUR)* **23**(1), 5–48 (1991)
14. Google Differential Privacy Team: Secure noise generation. [github.com/google](https://github.com/google/differential-privacy) **1**(1), 1–14 (Jun 2020)
15. Haeberlen, A., Pierce, B.C., Narayan, A.: Differential privacy under fire. In: *USENIX Security Symposium*. vol. 33 (2011)
16. Holohan, N., Braghin, S., Mac Aonghusa, P., Levacher, K.: Diffprivlib: the IBM differential privacy library. *ArXiv e-prints* **1907.02444 [cs.CR]** (Jul 2019)
17. Ilvento, C.: Implementing the exponential mechanism with base-2 differential privacy. In: *Proceedings of the 2020 ACM SIGSAC Conference on Computer and Communications Security*. pp. 717–742 (2020)
18. Karney, C.F.F.: Sampling exactly from the normal distribution. *ACM Trans. Math. Softw.* **42**(1) (Jan 2016). <https://doi.org/10.1145/2710016>
19. Kinderman, A.J., Monahan, J.F.: Computer generation of random variables using the ratio of uniform deviates. *ACM Trans. Math. Softw.* **3**(3), 257–260 (Sep 1977). <https://doi.org/10.1145/355744.355750>
20. Kotz, S., Kozubowski, T., Podgorski, K.: *The Laplace distribution and generalizations: a revisit with applications to communications, economics, engineering, and finance*. Springer Science & Business Media (2012)
21. McSherry, F., Talwar, K.: Mechanism design via differential privacy. In: *Foundations of Computer Science, 2007. FOCS’07. 48th Annual IEEE Symposium on*. pp. 94–103. IEEE (2007)
22. Mironov, I.: On significance of the least significant bits for differential privacy. In: *Proceedings of the 2012 ACM Conference on Computer and Communications Security*. p. 650–661. CCS ’12, Association for Computing Machinery, New York, NY, USA (2012). <https://doi.org/10.1145/2382196.2382264>
23. Mironov, I., Pandey, O., Reingold, O., Vadhan, S.: Computational differential privacy. In: *Annual International Cryptology Conference*. pp. 126–142. Springer (2009)
24. Thomas, D.B., Luk, W., Leong, P.H., Villasenor, J.D.: Gaussian random number generators. *ACM Comput. Surv.* **39**(4), 11–es (Nov 2007). <https://doi.org/10.1145/1287620.1287622>
25. van der Walt, S., Colbert, S.C., Varoquaux, G.: The NumPy array: A structure for efficient numerical computation. *Computing in Science and Engineering* **13**(2), 22–30 (2011). <https://doi.org/10.1109/MCSE.2011.37>

## Appendix

### A Probability density functions

The following probability distributions are referenced in Section 4.

### A.1 Uniform distribution

The uniform distribution on the interval  $[a, b] \subset \mathbb{R}$ ,  $-\infty < a < b < \infty$ , is given by the Probability Density Function (PDF)

$$f_{\mathcal{U}(a,b)}(x) = \begin{cases} \frac{1}{b-a} & \text{if } x \in [a, b], \\ 0 & \text{otherwise.} \end{cases}$$

We make use of the uniform distribution  $\mathcal{U}(0, 1)$  on the unit interval  $[0, 1]$ .

### A.2 Gaussian distribution

The Gaussian distribution with mean  $\mu$  and variance  $\sigma^2$  is given by the PDF

$$f_{\mathcal{N}(\mu,\sigma)}(x) = \frac{1}{\sigma\sqrt{2\pi}} e^{-\frac{1}{2}\left(\frac{x-\mu}{\sigma}\right)^2}.$$

We refer to the case when  $\mu = 0$  and  $\sigma = 1$  as the *standard Gaussian distribution*. If  $N \sim \mathcal{N}(0, 1)$ , then  $\sigma N + \mu \sim \mathcal{N}(\mu, \sigma)$ .

### A.3 Laplace distribution

The Laplace distribution with mean  $\mu$  and variance  $2b^2$  is given by the PDF

$$f_{\text{Lap}(\mu,b)}(x) = \frac{1}{2b} e^{-\frac{|x-\mu|}{b}}.$$

We refer to the case when  $\mu = 0$  and  $b = 1$  as the *standard Laplace distribution*. If  $L \sim \text{Lap}(0, 1)$ , then  $bL + \mu \sim \text{Lap}(\mu, b)$ .

### A.4 Exponential distribution

The exponential distribution with mean  $\frac{1}{\lambda}$  and variance  $\frac{1}{\lambda^2}$  is given by the PDF

$$f_{\text{Exp}(\lambda)}(x) = \lambda e^{-\lambda x}.$$

We refer to the case when  $\lambda = 1$  as the *standard exponential distribution*. If  $E \sim \text{Exp}(1)$ , then  $\frac{E}{\lambda} \sim \text{Exp}(\lambda)$ .

### A.5 Gamma distribution

The gamma distribution with mean  $k\theta$  and variance  $k\theta^2$  is given by the PDF

$$f_{\Gamma(k,\theta)}(x) = \frac{1}{\Gamma(k)\theta^k} x^{k-1} e^{-\frac{x}{\theta}}.$$

If  $G \sim \Gamma(k, \theta)$ , then  $cG \sim \Gamma(k, c\theta)$  for any  $c > 0$ .

## A.6 Chi-squared distribution

The chi-squared distribution with  $k \in \mathbb{N}$  degrees of freedom is given by the PDF

$$f_{\chi^2(k)}(x) = \frac{1}{2^{\frac{k}{2}} \Gamma\left(\frac{k}{2}\right)} x^{\frac{k}{2}-1} e^{-\frac{x}{2}}.$$

## B Code samples

The following code samples were used in estimating execution time for different implementations. This code was run using Python 3.8.6.

### B.1 Naïve sampling

The naïve standard Laplace sampling given by (2) was implemented using:

---

```
1 def laplace_naive():
2     u = random()
3     return copysign(log(1 - 2 * abs(u - 0.5)), u - 0.5)
```

---

### B.2 Theorem 1 sampling

The implementation of Theorem 1 was given by:

---

```
1 def laplace_theorem1():
2     return log(1 - random()) * cos(pi * random()) + log(1 - random()) *
        cos(pi * random())
```

---

### B.3 Sampling with math and random

We combine the Gaussian and Laplace sampling procedures from (7) and (2) to generate standard Laplace samples from  $8n$  uniform variates using the math and random libraries as follows:

---

```
1 def gaussian_sum(n=1):
2     return sum(normalvariate(0, 1) for i in range(n))
3
4 def laplace_math_and_random(n=1):
5     return (gaussian_sum(n) * gaussian_sum(n) - gaussian_sum(n) *
        gaussian_sum(n)) / n
```

---

#### B.4 Sampling with Numpy

Finally, we present an implementation of the same procedure using the popular Numpy package, leveraging its C-based code for faster computations with larger  $n$ :

---

```
1 import numpy as np
2
3 def laplace_numpy(n=1):
4     g1, g2, g3, g4 = np.random.standard_normal(size=(4, 2 * n)).sum(
5         axis=1)
6     return (g1 * g2 - g3 * g4) / 2 / n
```

---