



Secure Noise Sampling for DP in MPC with Finite Precision

Hannah Keller
hkeller@cs.au.dk
Aarhus University
Denmark

Helen Möllering
moellering@crypto.cs.tu-
darmstadt.de
McKinsey & Company
Germany

Thomas Schneider
schneider@crypto.cs.tu-
darmstadt.de
Technical University of Darmstadt
Germany

Oleksandr Tkachenko
oleksandr.tkachenko1@gmail.com
DFINITY Foundation
Germany

Liang Zhao
liang.zhao@tu-darmstadt.de
Technical University of Darmstadt
Germany

ABSTRACT

While secure multi-party computation (MPC) protects the privacy of inputs and intermediate values of a computation, differential privacy (DP) ensures that the output itself does not reveal too much about individual inputs. For this purpose, MPC can be used to generate noise and add this noise to the output. However, securely generating and adding this noise is a challenge considering real-world implementations on finite-precision computers, since many DP mechanisms guarantee privacy only when noise is sampled from continuous distributions requiring infinite precision.

We introduce efficient MPC protocols that securely realize noise sampling for several plaintext DP mechanisms that are secure against existing precision-based attacks: the discrete Laplace and Gaussian mechanisms, the snapping mechanism, and the integer-scaling Laplace and Gaussian mechanisms. Due to their inherent trade-offs, the favorable mechanism for a specific application depends on the available computation resources, type of function evaluated, and desired (ϵ, δ) -DP guarantee.

The benchmarks of our protocols implemented in the state-of-the-art MPC framework MOTION (Braun et al., TOPS'22) demonstrate highly efficient online runtimes of less than 32 ms/query and down to about 1ms/query with batching in the two-party setting. Also the respective offline phases are practical, requiring only 51 ms to 5.6 seconds/query depending on the batch size.

KEYWORDS

Secure Multi-party Computation, Differential Privacy, Noise Sampling, Secure Implementations, Finite-Precision Computing

ACM Reference Format:

Hannah Keller, Helen Möllering, Thomas Schneider, Oleksandr Tkachenko, and Liang Zhao. 2024. Secure Noise Sampling for DP in MPC with Finite Precision. In *The 19th International Conference on Availability, Reliability and Security (ARES 2024), July 30–August 02, 2024, Vienna, Austria*. ACM, New York, NY, USA, 12 pages. <https://doi.org/10.1145/3664476.3664490>

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

ARES 2024, July 30–August 02, 2024, Vienna, Austria

© 2024 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 979-8-4007-1718-5/24/07

<https://doi.org/10.1145/3664476.3664490>

1 MOTIVATION

The aggregation and statistical analysis of many individuals' data became common across multiple industries, e.g., for the detection of financial frauds [9], the improvement of disease diagnostics [26], or the matching of organ donors with patients [16]. Such analyses may, however, conflict with data privacy. To complicate the matter, data is often stored at different locations, e.g., hospitals or banks, and those parties do not wish or are legally not permitted to share their information with the other parties. However, collaboration is imperative for effectively gaining new insights from distributed data. Secure multi-party computation (MPC) [5, 19, 33] offers a cryptographic approach to execute an analysis on data in a privacy-preserving manner, even when the data is distributed among multiple parties. MPC guarantees that no party learns more than what is revealed by the published analysis output.

MPC protocols usually yield exact results, but provide no protection against attacks that use these outputs to make inferences about the individuals that contributed their data to the analysis. Without additional output privacy measures, published results may be vulnerable to multiple attacks ranging from simple linkage and reconstruction attacks [74, 84] to inference and model inversion attacks [45, 54, 83] in machine learning applications. In many of these scenarios, an adversary may use publicly available information to fully reconstruct the input data used for the analysis. Differential privacy (DP) [39] has emerged as the de facto standard for guaranteeing that the published output of a function does not reveal too much information about its input. In fact, companies such as Google [88] and Apple [85], as well as the US Census Bureau [34], already use DP in practice.

In many cases, DP is guaranteed by mechanisms that sample noise from some distribution and add this noise to some computed function result. When the output of the function to be released is an integer, noise can be sampled from a discrete distribution, like the discrete Laplace [48] or discrete Gaussian distribution [23]. If the function output is not an integer, many DP mechanisms require noise sampled from a continuous distribution. Representing real numbers from continuous distributions with only finite precision, as is necessary on computers, leads to a violation of the DP guarantee in general [46]. In fact, textbook Laplace and Gaussian mechanism implementations allow an attacker to recover the entire database [59, 71]. Therefore, we consider five DP mechanisms in our work that are known to guarantee DP when implemented

on finite-precision computers, two that sample from discrete distributions, and three that sample from approximated versions of continuous distributions.

Those five DP mechanisms are, however, designed for a scenario with a central database which is often not available for real-world applications where data is held by multiple parties. This scenario is commonly known as the central DP model [38]. As an alternative model, local DP [63] drops this assumption, instead requiring each individual to add noise necessary for preserving DP to their own data and publish the result, after which the function is computed on this noisy published data. This approach adds significantly more noise, resulting in published results with more error and less utility.

Using MPC, the trusted aggregator in the central DP model can be replaced by a distributed protocol run by multiple non-colluding parties. We focus on an outsourcing scenario, where a large number of data owners secret share their data to a small number of non-colluding parties. We realize the central DP model where the non-colluding parties jointly sample noise from a distribution under MPC with limited precision, such that DP is securely guaranteed and good utility is preserved. We choose floating-point arithmetic to implement the above five secure DP mechanisms for two reasons: (1) it satisfies the accuracy requirement of these DP mechanisms which is not the case for fixed-point arithmetic; (2) it is more efficient than the noise sampling methods of [27, 38, 42] that are based on processing a large sequence of Boolean bits when the required sampling interval is fairly large, e.g., sampling a geometric random value in $[0, 2^{128})$. Besides, existing plaintext sampling algorithms for DP mechanisms cannot simply be “translated” into an MPC protocol, so we add multiple optimizations to yield a practically efficient solution.

Outline and Contributions. After introducing the baseline techniques used in our work in §2 and providing an overview of related work in §3, we present the following contributions:

- (1) *MPC-based Noise Sampling with Finite Precision.* We introduce MPC protocols for sampling random integers and floating-point values from the Laplace and Gaussian distributions in §4. Notably, our protocols are not susceptible to precision-based attacks (cf. §3), which was not considered by the MPC community so far.
- (2) *MPC-Protocols for Discrete Mechanisms.* Based on our secure random sampling protocols, we design MPC protocols for the discrete Laplace and Gaussian mechanisms [23] in §4.1, which satisfy DP for queries returning integer values. We introduce several protocol-level optimizations, that can halve the circuit size and reduce its depth by 36–88% depending on the employed MPC technique in §4. We reduce runtimes by up to $8.6\times$ (cf. §5.2) compared to a naive implementation.
- (3) *MPC-Protocols for Continuous Mechanisms.* Similarly, we design the first MPC protocols for the snapping and integer-scaling Laplace and Gaussian mechanisms with floating-point arithmetic. Our MPC protocols for the snapping mechanism by Mironov [71] in §4.2 and Google’s integer-scaling Laplace and Gaussian mechanisms [51] in §4.3 are applicable for floating-point arithmetic.
- (4) *Open-source implementation and benchmarks.* We implement our five MPC-protocols for DP mechanisms with both integer and floating-point arithmetic in the state-of-the-art MPC framework MOTION [19]. The implementation will be open-sourced upon

acceptance of our work. Furthermore, we extensively benchmark all algorithms and experimentally compare them to related work in §5. All but one of our MPC-based sampling protocols are independent of the input data, so they can be entirely pre-computed in an offline phase. The online phases of our MPC-based DP mechanisms are highly efficient using batching, e.g., 1.4–9.5 ms per query with a batch size of 30 or 40 (cf. §5.2).

2 PRELIMINARIES

In this section, we summarize the concept of differential privacy and the cryptographic building blocks used in our work. We consider security from two different perspectives: First, we achieve *computational privacy* via MPC, i.e., protecting the input data X of different data owners during the computation. Second, we achieve *output privacy* via DP, i.e., the information leakage about the input data X that can be inferred from the output $f(X)$ is limited by the DP guarantees.

2.1 Output Privacy

Differential Privacy (DP). Intuitively, DP guarantees that including an individual’s data record in an analysis has only a limited impact on the result of the analysis. Let $X \in D^n$ be a database represented as a vector of n entries from some domain D . Typically a domain will be of types $\{0, 1\}^d$ or \mathbb{R}^d . The size of a database X is measured by its ℓ_1 norm: $\|X\|_1 = n$. X' is a neighboring database of X if it differs by one record, such that X' is constructed by replacing one record of X with a different record. For DP, randomized mechanisms M are used to ensure the similarity of the outputs of a statistical analysis on two neighboring databases X and X' . Differential privacy is formally defined as follows:

Definition 2.1 (Differential Privacy [38, 39]). A randomized algorithm $M : D^n \rightarrow \mathcal{Y}$ is (ϵ, δ) -differentially private if for any two neighboring databases $X, X' \in D^n$, and all sets $T \subseteq \mathcal{Y}$,

$$\Pr [M(X) \in T] \leq e^\epsilon \cdot \Pr [M(X') \in T] + \delta. \quad (1)$$

Smaller ϵ values provide a stronger privacy guarantee. ϵ is chosen to be a small, non-negligible constant. δ should be much smaller than $\frac{1}{\|X\|_1}$ to prevent the leakage of complete data records.

Discrete Outputs. When functions have discrete outputs, it is possible to add noise sampled from a distribution of discrete values, avoiding precision-based attacks on some implementations of continuous distributions. Take integer outputs for example, the discrete Laplace or discrete Gaussian distributions can be used for noise addition.

Laplace and Gaussian Mechanisms. Additive Laplace [39] and Gaussian [40] noise satisfy DP by adding noise to the output of a function $f : D^n \rightarrow \mathbb{R}^k$, that maps a database $X \in D^n$ to k real numbers. Thereby, the noise’s magnitude must be chosen based on the ℓ_1 - and ℓ_2 -sensitivities of function f : $\Delta_1 f = \max\|f(X) - f(X')\|_1$ and $\Delta_2(f) = \max\|f(X) - f(X')\|_2$ for any two neighboring databases X and X' . I.e., it measures the maximum difference in the output that modifying a single record in a database can cause.

While the Laplace mechanism guarantees ϵ -DP, the Gaussian mechanism can only offer (ϵ, δ) -DP, a weaker guarantee. The Laplace distribution overall samples less noise than the Gaussian distribution for one-dimensional single outputs; however, the

Gaussian mechanism allows for better utility when releasing high-dimensional outputs, which is relevant for machine learning applications [1, 57, 70]. Furthermore, Gaussian noise has the same distribution as much naturally occurring noise in data measurements, and several sources of Gaussian noise add nicely.

LEMMA 2.2 (LAPLACE MECHANISM [39]). *Given any function $f : D^n \rightarrow \mathbb{R}^k$ with ℓ_1 -sensitivity $\Delta_1 f$ and a privacy parameter ϵ , the Laplace mechanism satisfies $(\epsilon, 0)$ -DP and is defined as:*

$$M_{\text{Lap}}(X) = f(X) + (Y_1, \dots, Y_k), \quad (2)$$

where Y_i are i.i.d. random values drawn from a Laplace distribution $\text{Lap}(b)$ with probability density function $P(x|b) = \frac{1}{2b}e^{-\frac{|x|}{b}}$ and $b = \Delta_1 f / \epsilon$.

LEMMA 2.3 (GAUSSIAN MECHANISM [40]). *Given any function $f : D^n \rightarrow \mathbb{R}^k$ with ℓ_2 -sensitivity $\Delta_2(f)$ and privacy parameters (ϵ, δ) , the Gaussian mechanism is defined as:*

$$M_{\text{Gauss}}(X) = f(X) + (Y_1, \dots, Y_k), \quad (3)$$

where Y_i are i.i.d. random values drawn from a Gaussian distribution $\mathcal{N}(\mu = 0, \sigma^2)$ with probability density function $P(x|\mu = 0, \sigma^2) = \frac{1}{\sigma\sqrt{2\pi}} \cdot e^{-(x-\mu)^2/(2\sigma^2)}$ and $\sigma^2 > \frac{2 \ln(\frac{1.25}{\delta}) \cdot (\Delta_2(f))^2}{\epsilon^2}$.

Distributed Computational Differential Privacy. The original definition of differential privacy holds against computationally unbounded adversaries. However, efficient MPC protocols are often only secure against computationally bounded adversaries. Hence, we use the definition of distributed or computational differential privacy introduced in [12] that considers computationally bounded adversaries.

Definition 2.4 (Distributed Computational Differential Privacy (CDP) [12]). A randomized protocol Π implemented among N computation parties $\mathcal{P} = \{\mathcal{P}_1, \dots, \mathcal{P}_N\}$ satisfies $(\epsilon, \delta = \text{negl}(\kappa))$ computational distributed differential privacy w.r.t. a coalition $C \subset \mathcal{P}$ of semi-honest computation parties of size t , if for every probabilistic polynomial-time distinguisher D and for any neighboring databases X, X' and any possible set of views (e.g., internal state and exchanged messages of party \mathcal{P}_i) for protocol Π ,

$$\Pr \left[D(\text{VIEW}_{\Pi}^C(X)) = 1 \right] \leq e^\epsilon \cdot \Pr \left[D(\text{VIEW}_{\Pi}^C(X')) = 1 \right] + \text{negl}(\kappa),$$

where κ is a security parameter for computationally bounded adversaries.

2.2 Computational Privacy

Secure Multi-Party Computation (MPC). MPC enables $N \geq 2$ parties $\mathcal{P}_1, \dots, \mathcal{P}_N$ to securely compute any function $y = f(x_1, \dots, x_N)$ while revealing nothing but the output y . There are two typical MPC deployment scenarios: Either, the N data owners jointly run the computation taking their data as input which is, however, kept private from each other. Alternatively, MPC can also be used in an *outsourcing* scenario [62], where M data owners generate shares of their private inputs and send the shares to the N non-colluding computing parties. Those evaluate f on the received secret shared inputs using MPC and obtain the secret shared output while learning no information. Our benchmarks in §5 use the MPC framework

MOTION [19] that implements mixed-protocol MPC tolerating up to $N - 1$ passively corrupted parties. MOTION combines the GMW protocol [49] in its Boolean and arithmetic versions with the constant round BMR protocol [14].

We can use the result from [12, 38, 87] to prove that a mechanism remains computationally differentially private when implemented using a computationally secure MPC protocol.

THEOREM 2.5. *Define some function $\mathcal{M} : D^n \rightarrow \mathcal{Y}$, with input domain D and output domain \mathcal{Y} . Let $\Pi_{\mathcal{M}} : (D^{n/m})^m \rightarrow \mathcal{Y}$ be an m -party protocol that computationally securely computes \mathcal{M} with security parameter κ and security against up to t corruptions. If \mathcal{M} is (ϵ, δ) -DP, and $\delta_\kappa = \text{negl}(\kappa)$, then $\Pi_{\mathcal{M}}$ is $(\epsilon, (\exp(\epsilon) + 1)\delta_\kappa + \delta)$ -CDP with security against up to t corruptions.*

A proof sketch is given in the full version.

Number Representations. While the MPC protocols in MOTION [19] operate over rings (\mathbb{Z}_{2^ℓ}) , where $\ell \geq 1$, DP mechanisms require drawing random values from probability distributions [40]. To represent those, we use integer or floating-point representations. *Floating-Point Representation.* According to the IEEE floating-point standard [67], a floating-point number u is represented with a sign bit S , an exponent E and significant bits $d_1, \dots, d_p \in \{0, 1\}^p$, where $u = (-1)^S \times (1.d_1 \dots d_p)_2 \times 2^E$. Essentially, floating-point numbers can only represent a subset of real numbers leading to limited precision for the others. A series of works [24, 53, 59, 71] present attacks against insecure implementations of plaintext DP mechanisms using floating-point representations (cf. §3).

Since we have to sample from a distribution with finite domain for MPC, we must account for an error probability when operating over rings \mathbb{Z}_{2^ℓ} as it is not possible to sample values greater than 2^ℓ . Therefore, with some small probability δ' , a sample will lie outside the range $[0, 2^\ell)$, which we have to account for in the failure probability. The following lemma follows directly from the privacy of \mathcal{M} and Definition 2.1.

LEMMA 2.6. *Consider an additive noise mechanism $\mathcal{M} = f(x) + n$, where additive noise $n \leftarrow \mathcal{D}$ is sampled from a distribution that satisfies (ϵ, δ) -DP. Then a mechanism \mathcal{M}' satisfies $(\epsilon, \delta + \delta')$ -DP, where δ' is the probability with which $n \leftarrow \mathcal{D}$ is out of the range $[0, 2^\ell)$:*

$$\mathcal{M}'(x) = \begin{cases} f(x) + n, n \leftarrow \mathcal{D} & \text{w.p. } 1 - \delta' \\ f(x) & \text{w.p. } \delta'. \end{cases}$$

3 RELATED WORK

In this section, we discuss related works on combining DP and MPC, as well as on the privacy issues of finite-precision implementations of DP mechanisms.

Combination of DP and MPC. When multiple parties own data, DP protocols need to either be interactive, assume computationally bounded adversaries or both [30], as we consider in our work, to avoid large errors in the released output. Multiple works, e.g., [18, 27, 29, 31, 38, 41, 42, 58, 65, 76, 89] combine MPC techniques with DP to achieve a similar utility as in the central DP model while removing the requirement of a trusted aggregator. Dwork et al. [38] use secret sharing-based MPC for noise generation. Their protocols generate noise from a Gaussian distribution (approximated with a

binomial distribution) and a discrete Laplace distribution (approximated with a Poisson distribution) by processing a sequence of unbiased/biased Boolean bits in MPC. Eriguchi et al. [42] improve the noise generation of [38] by reducing the communication and round complexity of the MPC protocols as well as removing the failure probability of the sampling algorithms. However, the protocols of Dwork et al. [38] and Eriguchi et al. [42] rely on Shamir’s secret sharing [80], which is defined on a field. In contrast, the MPC techniques used in our work are defined on a ring, significantly improving the efficiency of our MPC-based DP protocols. Champion et al. [27] propose secure computation methods for sampling biased bits improving previous work by Dwork et al. [38]. However, their sampling domain is not large enough to support our MPC-based DP mechanism. For example, for geometric sampling, we require the sampling domain to be $[0, 2^{128})$, while a solution based on the biased bits protocols of [27] can only efficiently support a domain of size $[0, 85)$. An extension of their protocol to our required domain size would have exponential time overhead in MPC. Eigner et al. [41] propose an architecture called PrivaDA to realize the Laplace, discrete Laplace, and exponential mechanisms in MPC using a combination of floating and fixed-point arithmetic operations for efficiency reasons. Knott et al. [65] propose a machine learning framework that implements the Gaussian mechanism in MPC. Wu et al. [89] use Shamir’s secret sharing [80]-based protocols that approximate the Laplace and Gaussian distributions using the central limit theorem [8]. However, Eigner et al. [41], Wu et al. [89] and Knott et al. [65] do not provide an analysis regarding whether the implementation of arithmetic operations affects the DP guarantee. In an orthogonal line of work, Pettai et al. [76] and Choquette-Choo et al. [31] present frameworks where a trusted third party generates and adds DP noise to the query result which is computed under MPC. In contrast, our protocols do not rely on a trusted third party.

Although we focus on the central DP model due to its favorable properties with respect to the accuracy, the local DP model has also been combined with cryptographic techniques such as homomorphic encryption [2, 3, 15, 28, 50, 60, 77, 81, 82, 86], functional encryption [90, 92], authenticated encryption [13, 22], secret sharing [55], or other MPC techniques [61] to enhance data privacy.

Attacks against DP Mechanisms. The privacy of many DP mechanisms is based on two implicit assumptions [40]: (1) computations are performed on real numbers with infinite precision; (2) the noise is accurately sampled from a probability distribution (e.g., Laplace or Gaussian distribution). However, an implementation of DP mechanism is typically done with floating-point or fixed-point arithmetic that only provides finite precision. Mironov [71] demonstrates that the Laplace mechanism implementations using Laplace noise $Y \sim \text{Lap}(\lambda)$ sampled with the textbook algorithms (i.e., compute $Y = (2Z - 1) \cdot \lambda \ln(U)$ using double-precision floating-point arithmetic, where $U \in (0, 1)$ and $Z \in \{0, 1\}$) enables an attacker to recover the entire database. Concretely, Mironov [71] shows that outputs of such implementations concentrate on a small subset of floating-point values which are correlated for inputs, leading to a breach of the DP guarantee. Jin et al. [59] extend the floating-point attack of Mironov [71] to the implementations of Gaussian mechanisms that generate Gaussian noise using the Marsaglia polar method [68], Box-Muller method [66], Ziggurat method [69],

etc. Haney et al. [53] present precision-based floating-point attacks against implementations of several DP mechanisms (e.g., Laplace, Gaussian, Staircase [47], etc.) that enable an adversary to infer if the query result $f(D)$ equals 1. Casacuberta et al. [24] introduce another type of attack on DP mechanisms that fail to add a sufficient amount of noise. They exploit the properties (e.g., overflow or rounding) of integer and floating-point arithmetic operations that can lead to an underestimation of the sensitivity (cf. §2.1) in DP mechanisms.

The discussed MPC-based Laplace or Gaussian mechanisms either do not specify how to correctly/securely sample random noise and perturb the input [41, 58, 65], assume infinite precision [76, 89], or rely on a trusted third party that generates and adds DP noise [31, 76]. We implement existing DP mechanisms that are known not to be vulnerable to precision based attacks [53, 59, 71] in MPC protocols.

4 PROTOCOLS

In this section, we present our tailored MPC protocols for three types of DP mechanisms not vulnerable to the precision-based attacks [53, 59, 71], namely the discrete Laplace and discrete Gaussian mechanisms [23], the snapping mechanism [71], and the integer-scaling mechanism [51].

Notation. Here, we introduce the notation used in the remainder of this work. The shares of a secret value x held by N parties $\mathcal{P}_1, \dots, \mathcal{P}_N$ are denoted by $\langle x^d \rangle^s = (\langle x^d \rangle_1^s, \dots, \langle x^d \rangle_N^s)$, where $\langle x^d \rangle_i^s$ is held by party \mathcal{P}_i , $i \in [N]$. The superscript $s \in \{A, B, Y\}$ is the secret sharing type. Here, A is Arithmetic sharing and B is Boolean sharing; both are based on the GMW protocol [49]. Y is the BMR protocol [14], an extension of Yao’s Garbled Circuits protocol [91] from two to multiple parties. The second superscript $d \in \{\mathbb{N}, \mathbb{Z}, \mathbb{L}\}$ indicates the data type in finite domains: \mathbb{N} are unsigned integers, \mathbb{Z} are signed integers, and \mathbb{L} are floating-point values. We omit the superscript or subscript of share $\langle x^d \rangle_i^s$ when it is clear from the context. $\langle x \rangle^D \leftarrow \text{S2D}(\langle x \rangle^S)$ is the conversion from one secret sharing technique S to another D , $S \neq D$ and $S, D \in \{A, B, Y\}$. Conversions between different data representations are presented in the same style. A multiplexer gate $\text{MUX}(a, b, c)$ returns b if a is true or otherwise c . An AND gate is denoted by \wedge , an OR gate by \vee , a NOT gate by \neg , and an XOR gate by \oplus .

Random Number Generation. In our secure DP mechanisms in §4.1-4.3, we rely on three types of random number generators:

- $(\langle b_0 \rangle^B, \dots, \langle b_{\ell-1} \rangle^B) \leftarrow \text{RandBits}(\ell)$ generates secret-shares of an ℓ -bit uniformly random Boolean string $b \in \{0, 1\}^\ell$.
- $\langle x \rangle^B \leftarrow \text{RandInt}(m)$ (cf. §4.1) generates secret-shares of a uniformly random unsigned integer $x \in [0, m - 1]$ for $m \in \mathbb{Z}$.
- $\langle u^{\mathbb{L}} \rangle^B \leftarrow \text{RandFloat}(l, k)$ generates uniformly random floating-point values $u \in [0, 1)$, where u has a l -bit mantissa and a k -bit exponent.

Computational Privacy. Our protocols have to fulfill both computational as well as output privacy. With respect to *computational privacy*, our protocols leak no intermediate values, i.e., all computation

is run in MPC. Thus, computational privacy follows directly from the provable security of the employed MPC techniques, namely the Arithmetic (A), and the Boolean variant (B) of GMW [49] and BMR [14] (cf. §2) as well as private conversions [19]. We discuss the *output privacy* of the protocols in their respective sections.

The best mechanism in practice depends on the type of function, the available computational resources, and the desired utility and DP guarantees. For functions that output integer values, the discrete Laplace M_{DLap} (cf. §4.1) and Gaussian mechanisms M_{DGauss} (cf. §4.1) are a good choice in general. When the outputs are floating-point values, the snapping mechanism M_{Snap} (cf. §4.2) or integer scaling Laplace M_{ISLap} (cf. §4.3) or Gaussian mechanisms M_{ISGauss} (cf. §4.3) can be used. M_{Snap} is the most efficient but it offers a sub-optimal privacy-utility trade-off [51]. M_{ISLap} and M_{ISGauss} offer a better privacy-utility trade-off, but they require more iterations with floating-point exponentiation, making the protocols less efficient.

4.1 MPC-Based Discrete Laplace/Gaussian Mechanisms

The discrete Laplace [48] and Gaussian DP mechanisms [23] are formulated as: $M_{\text{Discrete}}(D) = f(D) + Q$, where $f(D)$ is the output of a query function and Q is additive noise sampled from the discrete Laplace or discrete Gaussian distributions. M_{Discrete} requires that both the function output $f(D) \in \mathbb{Z}^k$ and the noise $Q \in \mathbb{Z}^k$ are integers.

From an MPC point of view, the critical step is to sample a random integer from the discrete Laplace and Gaussian distributions. Afterwards, the secret-shared random value is securely added to the share of query output $f(D)$, which is a single operation in MPC. Consequently, we focus on the MPC-based sampling protocols in the following.

The discrete Laplace and Gaussian distributions turn out to be closely related, such that a value sampled from a discrete Laplace distribution can be transformed into a sample from a discrete Gaussian distribution [23]. First, Prot. 2 generates shares of a randomly sampled discrete Laplace value $\langle Y^{\mathbb{Z}} \rangle^{\text{B}}$, $Y \sim \text{DLap}(\frac{t}{s})$. Next, Prot. 3 converts those discrete Laplace value shares $\langle Y^{\mathbb{Z}} \rangle^{\text{B}}$ into shares of discrete Gaussian value $\langle G^{\mathbb{Z}} \rangle^{\text{B}}$ sampled from a discrete Gaussian distribution $G \sim \text{DGauss}(\mu = 0, \sigma^2)$, where the variance σ^2 of added noise is calculated from the desired (ϵ, δ) -DP guarantee [10].

Our sampling of random discrete Laplace values in Prot. 2 follows the idea of Canonne et al. [23] because it is more efficient than the sampling methods of [27, 38, 42] that are based on processing a sequence of Boolean bits (cf. §1). Besides, we introduce several optimizations on the MPC side. The authors introduce a rejection sampling-based [37] approach that first generates geometric random values and converts them into discrete Laplace random values. Therefore, we also split the protocol for sampling from a discrete Laplace distribution into two sub-protocols. The first one, shown in Prot. 1, generates secret shares $\langle X^{\mathbb{N}} \rangle^{\text{B}}$ of a geometric random value $X \sim \text{Geo}(p = 1 - e^{-\frac{s}{t}})$. Afterwards, Prot. 2 converts $\langle X^{\mathbb{N}} \rangle^{\text{B}}$ into shares of a discrete Laplace random value $\langle Y^{\mathbb{Z}} \rangle^{\text{B}}$, $Y \sim \text{DLap}(\frac{t}{s})$.

In this section, we will first present the MPC protocol for sampling from a geometric distribution including several optimizations enhancing its efficiency. Next, we present the MPC protocol for sampling from a discrete Laplace distribution, which uses the geometric sampling protocol as a sub-protocol. Last, we present the discrete Gaussian sampling protocol, which uses the discrete Laplace sampling protocol as a sub-protocol.

Geometric Sampling. We begin by presenting the challenges and necessary optimizations involved in generating secret shares of geometric random values, the protocol for which is specified in Prot. 1. Our sub-protocols for oblivious selection Sel and Boolean-String multiplication BoolStrMul are presented in the full version. Concretely, $\langle u^{\mathbb{N}} \rangle^{\text{B}}, \langle b \rangle^{\text{B}} \leftarrow \text{Sel}(\langle u_0^{\mathbb{N}} \rangle^{\text{B}}, \dots, \langle u_{\ell-1}^{\mathbb{N}} \rangle^{\text{B}}, \langle b_0 \rangle^{\text{B}}, \dots, \langle b_{\ell-1} \rangle^{\text{B}})$ outputs a bit-string $u^{\mathbb{N}} = u_i^{\mathbb{N}}$ and a bit $b = b_i$, where i is the index of the first non-zero bit b_i for $i \in [0, \ell - 1]$. Protocol BoolStrMul($\langle a \rangle^{\text{B}}, \langle b_0 \rangle^{\text{B}}, \dots, \langle b_{\ell-1} \rangle^{\text{B}}$) [7] computes the multiplication of one Boolean bit a with a set of ℓ Boolean bits $b_0, \dots, b_{\ell-1}$.

Our MPC protocols directly implement the steps of the plaintext sampling algorithms in [23]. We chose algorithms from Canonne et al. [23] for MPC efficiency reasons. For example, Google’s DP library [52] also contains securely implemented plaintext discrete Laplace sampling, but it is based on binary search geometric sampling. A realization in MPC would require computing 52 iterations (each computing floating-point natural logarithm and exponentiations, thus, an impracticable overhead) as integers up to 2^{52} can be represented precisely in floating-point arithmetic.

To generate samples from a geometric distribution, the while-loops in Algorithm 2 in [23] repeatedly sample from Bernoulli distributions until termination conditions are met. Thus, each loop generating a random value can in theory run for an infinite number of iterations with negligible probability. As the number of iterations is not fixed, a key challenge is to realize it in MPC (1) without information leakage and (2) in an efficient manner.

An idea is to limit the number of iterations by fixing it to a pre-defined number κ of iterations. A check in MPC determines whether the random number has been generated; if it has not, κ additional iterations are run. Unfortunately, this strategy is not fully privacy-preserving, as it leaks one bit of information indicating whether to continue the computation. Critically, Jin et al. [59] construct a timing attack against the discrete Laplace mechanism [23] and show that the magnitude of the generated random value exhibits a positive linear relation to the number of required iterations, i.e., a large random value is usually generated in more iterations.

Instead, we omit the check and run the protocol for a fixed number of iterations κ . This, however, means that the output may be zero with some failure probability p_{failure} . We derive p_{failure} from κ for our MPC-protocols that sample random values from the geometric, discrete Laplace, and Gaussian distributions in the full version. We can guarantee a small failure probability, i.e., $p_{\text{failure}} < 2^{-40}$, by choosing an appropriate number of iterations κ . The DP guarantee will have $\delta = p_{\text{failure}} < 2^{-40}$, which protects individual privacy well when noise is added to functions of datasets with size $\|D\|_1 \ll 2^{40}$. Different choices for κ allow δ to be adjusted.

```

Input :  $s, t$  // Parameters of  $\text{Geo}(1 - e^{-\frac{s}{t}})$ 
Output:  $\langle X^{\mathbb{N}} \rangle^{\mathbb{B}}$  //  $X \sim \text{Geo}(1 - e^{-\frac{s}{t}})$  or  $X = 0$ 
// Skip 1-st loop if  $t == 1$ 
1 if  $t == 1$  then
2    $\langle u^{\mathbb{N}} \rangle^{\mathbb{B}} \leftarrow \langle 0^{\mathbb{N}} \rangle^{\mathbb{B}}$ 
3    $\langle b \rangle^{\mathbb{B}} = 1$ 
4 end
5 else
6   // Draw  $\kappa_1$  random integers from interval  $[0, t - 1]$ , and  $\kappa_1$ 
7   // random bits from a Bernoulli distribution
8   for  $i \leftarrow 0$  to  $\kappa_1 - 1$  do
9      $\langle u_i^{\mathbb{N}} \rangle^{\mathbb{B}} \leftarrow \text{RandInt}(t)$ 
10     $\langle b_i \rangle^{\mathbb{B}} \leftarrow \text{Bernoulli}\left(e^{-\frac{\text{UINT2FL}(\langle u_i^{\mathbb{N}} \rangle^{\mathbb{B}})}{t}}\right)$ 
11  end
12  // Set  $u = u_k$  and  $b = b_k$  if  $b_k == 1$  for  $k \in [0, \kappa_1 - 1]$ 
13   $\langle u^{\mathbb{N}} \rangle^{\mathbb{B}}, \langle b \rangle^{\mathbb{B}} \leftarrow$ 
14   $\text{Sel}\left(\langle u_0^{\mathbb{N}} \rangle^{\mathbb{B}}, \dots, \langle u_{\kappa_1 - 1}^{\mathbb{N}} \rangle^{\mathbb{B}}, \langle b_0 \rangle^{\mathbb{B}}, \dots, \langle b_{\kappa_1 - 1} \rangle^{\mathbb{B}}\right)$ 
15 end
16 // Draw  $\kappa_2$  random bits from a Bernoulli distribution
17 for  $j \leftarrow 0$  to  $\kappa_2 - 1$  do
18    $\langle c_j \rangle^{\mathbb{B}} \leftarrow \text{Bernoulli}(e^{-1})$ 
19    $\langle d_j \rangle^{\mathbb{B}} \leftarrow \neg \langle c_j \rangle^{\mathbb{B}}$ 
20 end
21 // Draw  $v = v_k \sim \text{Geo}(1 - e^{-1})$  and set  $d = d_k$  if  $d_k == 1$  for
22 //  $k \in [0, \kappa_2 - 1]$ 
23  $\langle v^{\mathbb{N}} \rangle^{\mathbb{B}}, \langle d \rangle^{\mathbb{B}} \leftarrow$ 
24  $\text{Sel}(v_0 \leftarrow 0, \dots, v_{\kappa_2 - 1} \leftarrow \kappa_2 - 1, \langle d_0 \rangle^{\mathbb{B}}, \dots, \langle d_{\kappa_2 - 1} \rangle^{\mathbb{B}})$ 
25 // Output  $X = \lfloor \frac{u+v \cdot t}{s} \rfloor$  or  $X = 0$  if  $\text{GeometricExp}(s, t)$  fails
26  $\langle X^{\mathbb{N}} \rangle^{\mathbb{B}} \leftarrow$ 
27  $\text{BoolStrMul}\left(\text{Floor}\left(\frac{\langle u^{\mathbb{N}} \rangle^{\mathbb{B}} + \langle v^{\mathbb{N}} \rangle^{\mathbb{B}} \cdot t}{s}\right), \langle b \rangle^{\mathbb{B}} \wedge \langle d \rangle^{\mathbb{B}}\right)$ 

```

Protocol 1: GeometricExp — our MPC protocol realizing Geometric sampling [23].

THEOREM 4.1. Consider mechanism \mathcal{M} that is (ϵ, δ) -DP, as well as mechanism \mathcal{M}' defined as follows:

$$\mathcal{M}'(x) = \begin{cases} \mathcal{M}(x) & \text{w.p. } 1 - p_{\text{failure}} \text{ (event } \neg E) \\ f(x) & \text{w.p. } p_{\text{failure}} \text{ (event } E). \end{cases}$$

Then mechanism \mathcal{M}' is $(\epsilon, \delta + p_{\text{failure}})$ -DP. The proof is given in the full version.

We additionally introduce five optimizations on protocol level to further improve the efficiency of our MPC protocol:

(1) *Parallelization.* Canonne et al.'s [23] geometric sampling algorithm uses two nested while-loops. The inner loop is run only if the outer loop was successful. In MPC, we cannot leak if the generation was successful. Hence, we fix the number of iterations to κ_1 and κ_2 iterations, which leads to $\kappa_1 \cdot \kappa_2$ iterations in total. Since the

```

Input :  $t, s$  // Parameters of  $\text{DLap}(\frac{t}{s})$ 
Output:  $\langle Y^{\mathbb{Z}} \rangle^{\mathbb{B}}$  //  $Y \sim \text{DLap}(\frac{t}{s})$  or  $Y = 0$ 
1 for  $i \leftarrow 0$  to  $\kappa_3 - 1$  do
2   // Generate sign  $S_i$  for  $Y \sim \text{DLap}(\frac{t}{s})$ 
3    $\langle S_i \rangle^{\mathbb{B}} \leftarrow \text{RandBits}(1)$ 
4   // Generate the integer part  $m_i \sim \text{Geo}(1 - e^{-\frac{s}{t}})$  for  $Y$ 
5    $\langle m_i^{\mathbb{N}} \rangle^{\mathbb{B}} \leftarrow \text{GeometricExp}(s, t)$ 
6   // Check if  $Y = (1 - 2S_i) \cdot m_i = (-1) \cdot 0 = -0$ 
7    $\langle f_i \rangle^{\mathbb{B}} \leftarrow \neg\left(\left(\langle S_i \rangle^{\mathbb{B}}\right) \wedge \left(\langle m_i^{\mathbb{N}} \rangle^{\mathbb{B}} == 0\right)\right)$ 
8 end
9 // If  $f_k == 1$ , set  $m = m_k$  and  $b = b_k$  for  $k \in [0, \kappa_3 - 1]$ 
10  $\langle m^{\mathbb{N}} \rangle^{\mathbb{B}} \parallel \langle S \rangle^{\mathbb{B}} \leftarrow \text{Sel}\left(\langle m_0^{\mathbb{N}} \rangle^{\mathbb{B}} \parallel \langle S_0 \rangle^{\mathbb{B}}, \dots, \langle f_0 \rangle^{\mathbb{B}}, \dots\right)$ 
11 Set  $\neg \langle S \rangle^{\mathbb{B}}$  as the sign bit of  $\langle m^{\mathbb{N}} \rangle^{\mathbb{B}}$ 
12 // Output  $Y = (1 - 2S) \cdot m$  or  $Y = 0$  if  $\text{DiscreteLaplace}(t, s)$ 
13 // fails
14  $\langle Y^{\mathbb{Z}} \rangle^{\mathbb{B}} \leftarrow \langle m^{\mathbb{N}} \rangle^{\mathbb{B}}$ 

```

Protocol 2: DiscreteLaplace — our MPC protocol realizing discrete Laplace sampling [23].

operations inside the loops are independent, we can run both loops independently and in parallel, leading to only $\kappa_1 + \kappa_2$ iterations. We choose κ_1 and κ_2 to guarantee that both loops output the required random values with high probability. Each loop can be run using Single Instruction Multiple Data (SIMD [17, 36, 79]), which fully parallelizes our protocol, enhances computation efficiency, and reduces memory consumption [36].

(2) *Random Integer Generation.* The random integer generation in line 7 of Prot. 1, indicated by $\text{RandInt}(t)$, generates a uniformly random integer $u \in [0, t - 1]$. It uses the Simple Modular Method [11] to generate secret shares of an t -bit random unsigned integer $x \in \{0, \dots, t - 1\}$ for $t \in \mathbb{Z}$. We observe that when $t = 2^k$, the expensive modular reduction operation can be omitted. Now each party can generate k uniform random bits $(\langle b_0 \rangle^{\mathbb{B}}, \dots, \langle b_{k-1} \rangle^{\mathbb{B}}) \leftarrow \text{RandBits}(k)$ and sets $\langle x \rangle^{\mathbb{B}, \mathbb{N}} = \langle b_0 \rangle^{\mathbb{B}}, \dots, \langle b_{k-1} \rangle^{\mathbb{B}}$ locally to create the secret-shared random integer.

(3) *Floating-Point Division.* We avoid the floating-point division in line 8 of Prot. 1 by first computing $e^{-\frac{1}{t}}$ in plaintext, followed with $e^{-\frac{1}{t}} \cdot e^{\text{UINT2FL}(\langle u_i^{\mathbb{N}} \rangle^{\mathbb{B}})}$. The floating-point multiplication is up to 4× faster than the floating-point division in $\{\mathbb{B}, \mathbb{Y}\}$.

(4) *Integer Division.* Integer division (Line 17 in Prot. 1) is a very expensive operation in MPC [19] (e.g., about 54.40 – 70.15 ms for a single division in B-sharing). Thus, we first convert integers from B-sharing to Y-sharing before dividing and rounding down to the next integer. This approach is up to 20× faster than integer division in B (cf. §5.2).

(5) *Bernoulli Sampling.* Canonne et al. [23] propose a Bernoulli sampling algorithm. However, similar to Prot. 1, it requires a large number of iterations to guarantee a negligible failure probability. Instead, we adopt a protocol from CrypTen [65] for sampling a random value b from a Bernoulli distribution with parameter p

and transfer it from fixed-point to floating-point arithmetic. A random value b drawn from the Bernoulli distribution equals 1 with probability p and 0 with probability $1 - p$. This is equivalent to $b = (x < p)$, where x is a uniformly random value in $(0, 1)$. We further optimize the efficiency of the comparison by purely relying on integer arithmetic. Concretely, each party locally generates κ random bits and interprets those as a random unsigned integer x . Then, it computes $b = x < (p \ll \kappa)$, where $p \ll \kappa$ is the binary representation of p after a left-shift of κ bits. Depending on the MPC technique, our integer comparison-based Bernoulli sampling protocol (Line 8, 13 in Prot. 1) is up to $4.8 - 6.5\times$ faster than the naive floating-point comparison-based protocol.

Discrete Laplace Sampling. Next, we show how to convert the samples from a geometric distribution to samples from a Laplace distribution, as specified in Prot. 2. The protocol generates secret shares of a discrete Laplace random value $Y = (1 - 2S) \cdot m \sim \text{DLap}(\frac{\epsilon}{S})$ using the random values drawn from the geometric distribution $m_i \sim \text{Geo}(1 - e^{-\frac{\epsilon}{i}})$ with Prot. 1, and a random sign bit $S_i \in \{0, 1\}$ (cf. Lines 2 – 3). However, $Y = (1 - 2S) \cdot m$ would equal 0 with twice the probability as it would occur in the discrete Laplace distribution [23]. Hence, we check in line 4 if $S_i == 1$ and $m_i == 0$. If it is the case, we re-sample both values. Otherwise, we either output $Y = (1 - 2S) \cdot m$ for $m = m_i$ and $S = S_i$, if $f_i == 1$ or $Y = 0$ otherwise (cf. Lines 6 – 8). κ_3 is again set such that $p_{\text{fail}}(\text{DiscreteLaplace, Prot. 2}) < 2^{-40}$.

Discrete Gaussian Sampling. We generate shares of a discrete Gaussian random value using Prot. 3 by adapting the plaintext Algorithm 3 of Canonne et al. [23] and set κ_4 such that Prot. 3 fails with a probability less than 2^{-40} .

Correctness. Correctness can be derived from the correctness of outputs from the plaintext algorithms and the MPC techniques employed. Our MPC-based sampling protocols execute all steps as specified by the respective plaintext algorithm. The sampling procedure fails to generate noise from the specified distribution with only a tunable, small probability, which we set to be negligible ($< 2^{-40}$) and incorporate into (ϵ, δ) -DP. This negligible failure probability ensures correctness. Next, we analyze the five optimizations we introduced for the noise sampling w.r.t. their effect on correctness. Due to the obliviousness requirement of MPC, i.e., the control flow of the computation must be independent of the input data, MPC protocols cannot check and terminate after a successful random integer generation. Instead, we fix the number of iterations of both loops generating random integers in Prot. 1, which increases computation overhead, keeping the behavior of the protocol unchanged except for the small failure probability. Our parallelization improves efficiency, but does not affect correctness. The random integer generation as well as the floating-point and integer divisions optimizations are only transformations of the original formulas, i.e., the computation is unchanged. The Bernoulli Sampling protocol was shown to be correct in [65]. Except for fixing the number of iterations (where correctness follows from the same arguments as discussed above¹), our MPC-based discrete Laplace mechanism

¹Note that this is true for all cases where the number of iteration is fixed in our protocols.

Input : σ // Parameter of DGauss ($\mu = 0, \sigma^2$)

Output: $\langle G^Z \rangle^B$ // $G \sim \text{DGauss}(\mu = 0, \sigma^2)$ or $G = 0$

```

1  $t \leftarrow \lfloor \sigma \rfloor + 1$ 
2 for  $j \leftarrow 0$  to  $\kappa_4 - 1$  do
3    $\langle Y_j^Z \rangle^B \leftarrow \text{DiscreteLaplace}(t, s = 1)$ 
4    $\langle a_j^L \rangle^B \leftarrow \frac{\left( \text{UINT2FL} \left( \left\langle Y_j^Z \right\rangle^B \right) - \frac{\sigma^2}{t} \right)^2}{2\sigma^2}$ 
5    $\langle b_j \rangle^B \leftarrow \text{Bernoulli} \left( e^{-\langle a_j^L \rangle^B} \right)$ 
6 end
7  $\langle G^Z \rangle^B, \langle b \rangle^B \leftarrow \text{Sel} \left( \langle Y_0^Z \rangle^B, \dots, \langle b_0 \rangle^B, \dots \right)$ 

```

Protocol 3: DiscreteGaussian — our MPC protocol realizing discrete Gaussian sampling [23].

M_{DLap} and discrete Gaussian mechanism M_{DGauss} in §4.1 are directly derived from the algorithms in [23].

Output Privacy. Output privacy (quantified by (ϵ, δ) -CDP) follows directly from the correctness and output privacy of the plaintext algorithms, as well as Lemma 2.6 and Theorems 2.5 and 4.1. Concretely, [23] proves that the original mechanisms satisfy (ϵ, δ) -DP, and our protocols correctly and securely generate noise from the same distributions except with a small failure probability. Taking into account the finite domain, our protocols satisfy $(\epsilon, \delta + \delta_\lambda + \delta' + p_{\text{failure}})$ -CDP, where δ_λ is negligible in the security parameter λ from Theorem 2.5, p_{failure} is set to some small failure probability from Theorem 4.1, and δ' is the probability with which the sampled noise exceeds the maximum value from Lemma 2.6. This bound is based on the tail bounds of the discrete Laplace and Gaussian distributions.

4.2 MPC-Based Snapping Mechanism

Mironov [71] proposes the snapping mechanism M_{Snap} as a remedy for the insecure implementations of the Laplace mechanism using *floating-point arithmetic*. The snapping mechanism is parameterized by B and λ , which can be chosen in advance, and satisfies $(\frac{1}{\lambda} + 2^{-49} \cdot \frac{B}{\lambda})$ -DP for $\lambda < B < 2^{46} \cdot \lambda$. The mechanism is defined as follows:

$$M_{\text{Snap}}(f(D), \lambda, \Lambda, B) = \text{clamp}_B \left(\lfloor \text{clamp}_B(f(D)) + Y_{\text{noise}} \rfloor_\Lambda \right),$$

$$Y_{\text{noise}} = S \cdot \lambda \cdot \ln(U).$$
(4)

In Eq. 4, $f(D)$ is the output of a query function that takes database D as input. Y_{noise} is the Laplace noise, S is the sign of the noise and uniformly distributed over $\{-1, 1\}$, and U is a random real number over $(0, 1)$ represented as a floating-point number that is output with probability proportional to its unit in the last place. $\ln(x)$ is the floating-point natural logarithm with exact rounding, i.e., $\ln(x)$ always rounds the output to the closest floating-point number. The addition (+) and multiplication (\cdot) operations in Eq. 4 are floating-point arithmetic operations. Function $\text{clamp}_B(x)$ limits the output to the interval $[-B, B]$ by outputting B if $x > B$, $-B$ if $x < -B$, and x otherwise. Function $\lfloor x \rfloor_\Lambda$ rounds x to the nearest

```

// Secret-shared exact query result, parameters of  $M_{\text{Snap}}$ 
Input :  $\langle f(D)^{\mathbb{L}} \rangle^B, \lambda, \Lambda, B$ 
// Secret-shared DP query result
Output:  $\langle y_{\text{Snap}}^{\mathbb{L}} \rangle^B$ 

// Generate a random floating-point number  $U \in (0,1)$  and a
// random bit
1  $\langle U^{\mathbb{L}} \rangle^B \leftarrow \text{RandFloat}$ 
2  $\langle S \rangle^B \leftarrow \text{RandBits}(1)$ 
   // Bound  $f(D)$  to  $[-B, B]$ 
3  $\langle f(D)_{\text{clamp}}^{\mathbb{L}} \rangle^B \leftarrow \text{Clamp}(B, \langle f(D)^{\mathbb{L}} \rangle^B)$ 
   // Generate Laplace noise  $Y = S \cdot \lambda \cdot \text{LN}(U)$ 
4  $\langle Y_{\text{noise}}^{\mathbb{L}} \rangle^B \leftarrow \lambda \cdot \text{LN}(\langle U^{\mathbb{L}} \rangle^B)$ 
5  $\text{Sign}(\langle S \rangle^B, \langle Y_{\text{noise}}^{\mathbb{L}} \rangle^B)$ 
   // Add noise  $Y_{\text{noise}}$  to the bounded query result  $f(D)_{\text{clamp}}$ 
6  $\langle x^{\mathbb{L}} \rangle^B \leftarrow \langle f(D)_{\text{clamp}}^{\mathbb{L}} \rangle^B + \langle Y_{\text{noise}}^{\mathbb{L}} \rangle^B$ 
   // Round perturbation result  $x$  to multiple of  $\Lambda$ 
7  $\langle x_{\Lambda}^{\mathbb{L}} \rangle^B \leftarrow \lfloor \langle x^{\mathbb{L}} \rangle^B \rfloor_{\Lambda}$ 
   // Bound  $x_{\Lambda}$  to  $[-B, B]$ 
8  $\langle y_{\text{Snap}}^{\mathbb{L}} \rangle^B \leftarrow \text{Clamp}(B, \langle x_{\Lambda}^{\mathbb{L}} \rangle^B)$ 

```

Protocol 4: M_{Snap} — our MPC protocol realizing the snapping mechanism [71].

multiple of Λ , where Λ is the smallest power of 2 greater than or equal to λ . Notice that clamping can be done without introducing additional error based on the choice of parameter B . To create the MPC version for the snapping mechanism [71] in Prot. 4, we first sample Y_{noise} , which can be done independently of the input and pre-computed in the offline phase. In the online phase, we must compute and clamp the function output, add Y_{noise} , and clamp the result $\text{clamp}_B(\lfloor \text{clamp}_B(f(D)) + Y_{\text{noise}} \rfloor_{\Lambda})$. We also encounter some challenges and introduce protocol optimizations to improve efficiency as follows:

(1) *Floating-Point Arithmetic.* The snapping mechanism [71] requires floating-point arithmetic to generate the random noise values added to the query result. However, floating-point arithmetic in MPC is expensive [4, 6], so most MPC frameworks [65, 72, 73] prefer to use fixed-point arithmetic [25]. We extend the MOTION framework [19] to support floating-point arithmetic in $\{A, B, Y\}$ -sharing. Specifically, we convert circuits from [35] that support IEEE 754 compliant floating-point arithmetic operations to the Bristol circuit format used in MOTION for $\{B, Y\}$ -sharing.

We also implement MPC protocols for logical/arithmetic shifting from [43] that are needed for the floating-point operations in A -sharing following ideas of Aliasgari et al. [4]. The difference is that the protocols in [4] are designed for Shamir’s secret sharing [80], whereas the operations are performed over a prime field \mathbb{F}_p (modulo p) that supports the inverse operation, while the A -sharing operations in the MOTION framework [19] are performed over a ring (modulo \mathbb{Z}_{2^t}).

(2) *Input Bounding.* MPC requires an input-independent program flow, i.e., branching depending on input values is not possible. Thus, the MPC-protocol realizing clamp_B must ensure that it does not leak the interval of the input x . We realize this by a multiplexer-based floating-point comparison protocol.

(3) *Rounding to Nearest Multiple of Λ .* $\lfloor x \rfloor_{\Lambda}$ rounds floating-point inputs x to the nearest multiple of Λ . As Λ is a power of two, our protocol can directly compute the rounding on the binary representation of input x . We create new depth- and size-optimized Boolean circuits with the CBMC-GC circuit compiler [21] to realize $\lfloor x \rfloor_{\Lambda}$ in MPC using $\{B, Y\}$ -sharing. The program for CBMC-GC [21] is inspired by Covington’s work [32] that relies on bit manipulation of the binary representation of floating-point numbers x . The bit manipulation operations are equivalent to first computing $x' = \frac{x}{\Lambda}$ using floating-point division and then rounding x' to the nearest integer. The rounded result can then be obtained by a simple multiplication: $\lfloor x \rfloor_{\Lambda} = x' \cdot \Lambda$. Evaluating our circuit for rounding with MPC is 4.5 – 15.0 \times faster than the above floating-point operations, as it only requires bit-level logical operations such as bit-shifting, AND, and XOR.

(4) *Sign Setting.* In Eq. 4, the sign of the Laplace random value Y_{noise} is multiplied by a random value $S \in \{-1, 1\}$. This expensive floating-point multiplication can be avoided by directly manipulating the sign bit of Y_{noise} . Concretely, in $\{B, Y\}$ -sharing (resp. A -sharing) the Boolean sign bit (resp. the arithmetic share containing the sign) of $\langle Y_{\text{noise}} \rangle$ is simply replaced by a freshly generated random secret-shared Boolean bit $\langle s \rangle$ (resp. random arithmetic share), where $s \in \{0, 1\}$. We indicate this operation as $\text{Sign}(\langle s \rangle, \langle Y_{\text{noise}} \rangle)$.

(5) *Secret Sharing.* A difficult open problem in MPC research is how to effectively determine the best mix of MPC-techniques for the most efficient, privacy-preserving instantiation of an algorithm. First attempts to create automatic compilers [20, 21, 35, 44, 56, 75] still exhibit significant shortcomings with respect to efficiency compared to protocols where the MPC-techniques have been carefully combined by hand. Thus, to find the most efficient instantiation for our MPC protocol of the snapping mechanism [71] (as well as for our other protocols), we micro-benchmark all relevant sub-protocols (i.e., floating-point addition, multiplication, and natural logarithm, Clamp and $\lfloor x \rfloor_{\Lambda}$) with each sharing ($\{A, B, Y\}$) to compose the most efficient mix. The results can be found in the full version. Taking conversion cost into account, using Y in the two-party setting and B -sharing in the multi-party setting for all parts in a LAN network (cf. §5.1) leads to the most efficient solution. The benchmark result of Prot. 4 is given in §5.2.

Correctness. Our snapping mechanism M_{Snap} (cf. Prot. 4) directly realizes the plaintext algorithm by Mironov [71] in MPC. Our five optimizations discussed in §4.2 are MPC protocol optimizations, i.e., formula transformations, efficient circuit generations, and efficient combinations of MPC techniques, which do not change the underlying computation.

Output Privacy. [71] proves that the original mechanism satisfies (ϵ, δ) -DP. Taking into account the finite domain, our protocols satisfy $(\epsilon, \delta + \delta_{\lambda} + \delta')$ -CDP, where δ_{λ} is negligible in the security parameter λ from Theorem 2.5, and δ' is the probability with which the sampled noise exceeds the maximum value from Lemma 2.6, which is based on the tail bounds of the Laplace distribution.

4.3 MPC-Based Integer-Scaling Mechanisms

Google [51] introduced a framework to securely realize DP by adding appropriately *scaled* discrete noise using floating-point arithmetic operations. Their integer-scaling Laplace mechanism was built to address some of the challenges associated with the snapping mechanism [71] and the precision-based attacks [53, 59, 71]. Namely, the snapping mechanism adds more noise than would theoretically be necessary for a given DP guarantee, offering a weaker privacy-utility trade-off than the integer-scaling mechanism [23]. The integer-scaling mechanism also offers a Gaussian variant. In the following, we call the two mechanisms introduced in [51] the integer-scaling Laplace mechanism (cf. §4.3) and the integer-scaling Gaussian mechanism (cf. §4.3). Both are defined as follows:

$$M_{\text{IS}}(f(D), r, \epsilon, \delta) = f_r(D) + i \cdot r, \quad (5)$$

where discrete random values i are scaled by a resolution parameter $r = 2^k$ (for $k \in [-1022, 970]^2$) which controls the discretization of the simulated continuous noise. M_{IS} satisfies (ϵ, δ) -DP and the function $f_r(D)$ rounds the output of a query function $f(D)$ to the nearest multiple of r . The scaled discrete noise $i \cdot r$ is used to simulate the continuous noise, e.g., the Laplace noise in Lemma 2.2 or the Gaussian noise Lemma 2.3 with the resolution parameter r . To create MPC protocols for the integer-scaling mechanisms, we first sample $i \cdot r$, which can be done independently of the input and pre-computed in the offline phase. In the online phase, we must compute and clamp the function output $f_r(D)$ and add the noise.

Integer-Scaling Laplace Mechanism. Here, we introduce our MPC protocol for the integer-scaling Laplace mechanism M_{ISLap} [51]. Prot. 5 presents our MPC protocol for the integer-scaling Laplace mechanism using the previously presented sub-protocol (cf. Prot. 2). The resolution parameter r is used to re-scale the exact query result $f(D)$ and DP noise i . It is set to the smallest power of 2 that is greater than $\frac{\Delta_1 f}{2^{\gamma} \epsilon}$ for $\gamma \in [10, 45]$, where γ controls accuracy and discretization, ϵ is the DP parameter, and $\Delta_1 f$ is the ℓ_1 -sensitivity of $f(D)$. We first generate shares of a discrete random value $\langle i_{\text{DLap}} \rangle$ using Prot. 2 (cf. Line 1). Then, in line 2, $\langle i_{\text{DLap}} \rangle$ is converted from integer to floating-point representation. To prevent precision-based attacks [53, 59, 71], integer i_{DLap} has to be scaled by r without precision loss, which requires $i_{\text{DLap}} \in [2^{-52}, 2^{52}]$. We use our MulPow2 protocol to re-scaled i_{DLap} by a factor $r = 2^k$, that is 3.1 – 20.5× faster than the direct floating-point multiplication, as it operates on the exponent part of the floating-point numbers. Lastly, we compute the DP query result $\langle f_r(D) \rangle$ after rounding $\langle f(D) \rangle$ to the nearest multiple of r by adding the scaled Laplace noise Y_{Lap} . We use the same rounding operation $\lfloor \cdot \rfloor_r$ we presented in §4.2.

We empirically evaluate the efficiency for each sub-protocol for the different secret sharing techniques. The most efficient approach is to fully run the protocol in Y -sharing in a LAN network (cf. §5.1) in the two-party setting or in B in the multi-party setting. The benchmark results of Prot. 5 are given in §5.2.

Integer-Scaling Gaussian Mechanism. The integer-scaling Gaussian and Laplace mechanisms both use symmetrical binomial DP noise i . We refer the readers to the full version for details.

²[51] sets $k \in [-1022, 1023]$, but if $i = 2^{52}$ and $r = 2^k = 2^{1023}$, $i \cdot r$ cannot be represented correctly as a double-precision floating point number.

```
// Secret-shared exact query result, resolution and DP
// parameters of  $M_{\text{ISLap}}$ ,  $\ell_1$ -sensitivity of query function  $f$ 
Input :  $\langle f(D) \rangle^{\text{L}}, r, \epsilon, \Delta f$ 
// Secret-shared DP query result
Output :  $\langle y_{\text{ISLap}} \rangle^{\text{B}}$ 

// Generate  $i_{\text{DLap}} \sim \text{DLap}(\frac{t}{s})$ , where  $\frac{t}{s} = \frac{\Delta f + r}{r \cdot \epsilon}$ 
1  $\langle i_{\text{DLap}}^{\text{N}} \rangle^{\text{B}} \leftarrow \text{DiscreteLaplace}(t, s)$ 
// Re-scale  $Y_{\text{Lap}} = i_{\text{DLap}} \cdot 2^k$ 
2  $\langle Y_{\text{Lap}}^{\text{L}} \rangle^{\text{B}} \leftarrow \text{MulPow2}(\text{UINT2FL}(\langle i_{\text{DLap}}^{\text{N}} \rangle^{\text{B}}), k = \log_2 r)$ 
// Round  $f(D)$  to nearest multiple of  $r$ 
3  $\langle f_r(D) \rangle^{\text{L}} \leftarrow \lfloor \langle f(D) \rangle^{\text{L}} \rfloor_r$ 
// Perturb  $f_r(D)$  with Laplace noise  $Y_{\text{Lap}}$ 
4  $\langle y_{\text{ISLap}}^{\text{L}} \rangle^{\text{B}} \leftarrow \langle f_r(D) \rangle^{\text{L}} + \langle Y_{\text{Lap}}^{\text{L}} \rangle^{\text{B}}$ 
```

Protocol 5: M_{ISLap} – our MPC protocol realizing the integer-scaling Laplace mechanism [51].

```
// Secret-shared exact query result, resolution parameter of
//  $M_{\text{ISGauss}}$ , parameter of  $\text{SymmBino}(n, p = 0.5)$ 
Input :  $\langle f(D) \rangle^{\text{L}}, r, \sqrt{n}$ 
// Secret-shared DP query result
Output :  $\langle y_{\text{ISGauss}} \rangle^{\text{B}}$ 

1  $\langle i_{\text{SymmBinoNoise}}^{\text{N}} \rangle^{\text{B}} \leftarrow \text{SymmetricalBinomial}(\sqrt{n})$ 
2  $\langle Y_{\text{GaussNoise}}^{\text{L}} \rangle^{\text{B}} \leftarrow \text{MulPow2}(\text{UINT2FL}(\langle i_{\text{N}} \rangle^{\text{B}}), k = \log_2 r)$ 
3  $\langle f_r(D) \rangle^{\text{L}} \leftarrow \lfloor \langle f(D) \rangle^{\text{L}} \rfloor_r$ 
4  $\langle y_{\text{ISGauss}}^{\text{L}} \rangle^{\text{B}} \leftarrow \langle f_r(D) \rangle^{\text{L}} + \langle Y_{\text{GaussNoise}}^{\text{L}} \rangle^{\text{B}}$ 
```

Protocol 6: M_{ISGauss} – our MPC protocol realizing the integer-scaling Gaussian mechanism [51].

Correctness. The integer-scaling Laplace mechanism M_{ISLap} (cf. Prot. 5) and integer-scaling Gaussian mechanism M_{ISGauss} (cf. Prot. 6) are directly based on [51].

Output Privacy. [51] proves that the original mechanisms satisfy (ϵ, δ) -DP. Since our protocols correctly and securely generate noise from the same distributions except with a small failure probability, taking into account the finite domain, our protocols satisfy $(\epsilon, \delta + \delta_\lambda + \delta' + p_{\text{failure}})$ -CDP, where δ_λ is negligible in the security parameter λ from Theorem 2.5, p_{failure} is set to some small failure probability from Theorem 4.1, and δ' is the probability with which the sampled noise exceeds the maximum value from Lemma 2.6, based on the tail bounds of the Laplace and Gaussian distributions.

5 EXPERIMENTAL EVALUATION

To evaluate the efficiency of our MPC protocols for secure DP mechanisms presented in §4, we provide an implementation and extensive benchmarks. We aim at high usability and will release

our code³ as open-source software under the permissive MIT license. In this section, we first present our benchmark setup and give a security and complexity analysis of our MPC protocols before discussing the experimental results. Our benchmarks evaluate computation and communication efficiency, including comparisons to existing MPC protocols for *insecure* DP mechanisms [41, 65].

5.1 Experimental Setup

Server Configuration. The experiments are run on five servers equipped with Intel Core i9-7960X processors and 128GB RAM. We consider three network environments: (1) LAN10: 10-Gbit/s with 1ms RTT, (2) LAN1: 1-Gbit/s with 1ms RTT, and (3) WAN100: 1-Mbit/s with 100ms RTT. Benchmark results for LAN1 and WAN100 are given in the full version. We extend MOTION [19] to 8/16/32/64/128-bit signed integer and 32/64-bit floating-point arithmetic in $\{A, B, Y\}$ as well as conversions between those.

Setting. The implementation of our MPC-based DP mechanisms generate 64-bit random integer/floating-point values used as DP noise. They can be run in an outsourcing [62] or a multi-party computation scenario where N data owners run the computation. In our experiments, we benchmark the outsourcing setting where an arbitrary number of data owners secret share their data (i.e., the query results) and send the shares to the computing parties instantiated by the servers. Those then jointly add the noise for guaranteeing DP. Note that the noise shares can be generated *offline* (i.e., independently of the input data and, thus, *before* receiving the input data shares from the data owners). Moreover, in this scenario, the noise generation is *independent* of the number of data owners. Additionally, pre-computed DP noise shares can also be transferred to N-PC scenarios as long as the magnitude of the noise, the MPC protocols, and data types are correctly configured.

Implementation. In our experiments, we compare with the state-of-the-art, but *insecure* MPC-based (discrete) Laplace mechanisms in PrivaDA [41] and Gaussian mechanisms in CrypTen [65]. Note that although M_{DLap} [41] is not susceptible to the precision-based attacks [53, 59, 71] discussed in §3, its correctness relies on real number arithmetic that is not satisfied when using floating-point arithmetic. We (re-)implement our protocols as well as previous work in the state-of-the-art full threshold passively secure MPC framework MOTION [19] for a fair comparison. Note that our proposed MPC protocols can naturally be translated to frameworks with stronger security such as ABY3 [72] or MP-SDPZ [64]. We provide a complexity analysis of circuit size and depth, as well as experimental communication costs in the full version.

5.2 Performance

We present efficiency benchmarks of protocols in §4.

Sharing Techniques. In the two-party (2PC) setting, Yao’s Garbled Circuit [91] with Three-Halves garbling [78] turned out to be the most efficient technique for our building blocks in the microbenchmarks. With $N \geq 3$ parties, B-sharing was more efficient. All microbenchmarks for our building blocks can be found in the full version.

Optimization Effects. To evaluate the performance of optimizations presented in §4, we implement the *naive* and *optimized* versions of M_{DLap} , M_{DGauss} (§4.1), M_{Snap} (§4.2), and M_{ISLap} (§4.3) and $M_{ISGauss}$ (§4.3). The *naive* version refers to the protocols that are transferred from the plaintext algorithms without optimizations.

Runtimes. The benchmarks of our MPC-based DP mechanisms in LAN10 are given in Tab. 1. We test multiple batch sizes in our experiments, i.e., the number of independent random values of DP noise generated in parallel. As the batch size increases, the overhead per sampled value decreases, so computation cost amortizes. Therefore, we choose the largest possible batch size that does not cause a memory overflow (which is denoted by “-” for no result).

Discrete Laplace/Gaussian Mechanisms. The upper part of Tab. 1 contains the runtime for the naive, optimized, and vulnerable [41] discrete Laplace mechanisms M_{DLap} (§4.1). The results in Tab. 1 show that M_{DLap} (§4.1) runs out of memory when trying to generate more than one random noise value in the 5PC setting. The offline runtime of our secure optimized M_{DLap} (§4.1) is $26.8 - 42.8\times$ slower than the insecure M_{DLap} [41]. Tab. 1 also presents the discrete Gaussian mechanism M_{DGauss} (§4.1). In our experiments, the memory restrictions of our hardware allow us to only generate Gaussian random value when $\sigma \leq 1$.

Snapping and Integer-Scaling Laplace Mechanisms. In the lower half, Tab. 1 contains the runtimes of our MPC protocol for the snapping mechanism [71] M_{Snap} (Prot. 4) and the integer-scaling Laplace mechanism [51] M_{ISLap} (Prot. 5) as well as of the *vulnerable* version of the Laplace mechanism M_{Lap} [41] from PrivaDA. We implement all three mechanisms using floating-point arithmetic in $\{B, Y\}$, i.e., when the exact query result $f(D)$ are 64-bit floating-point numbers. In the 2PC setting, our optimized version of M_{Snap} has the best total runtime. Concretely, it is about 47% faster than our optimized M_{ISLap} and about 19% faster than the insecure M_{Lap} [41]. In the 5PC setting, our optimized M_{Snap} is 28% faster and our optimized M_{ISLap} is $1.9\times$ slower than the insecure M_{Lap} [41].

Integer-Scaling Gaussian Mechanism. The lower part of Tab. 1 also presents the runtime of our MPC-based integer-scaling mechanism $M_{ISGauss}$ (Prot. 6) and the *insecure* Gaussian Mechanisms M_{Gauss} [65] of CrypTen. We implement these mechanisms using floating-point arithmetic in Y in the 2PC setting and in B in the 3PC and 5PC settings for efficiency reasons. Due to memory restrictions of our hardware, we were not able to run $M_{ISGauss}$ with more than 2 parties and a batch size of 30 or with 5 parties and a batch size of 2. The offline runtime of our secure $M_{ISGauss}$ is $44.5 - 287.1\times$ slower than the insecure M_{Gauss} [65].

Offline vs. Online Phase. Note that our MPC-based noise sampling protocols can fully be run offline before the actual input is available. Thus, in the time-critical online phase, for the discrete DP mechanisms (cf. §4.1) only a secure addition of the noise to the secret-shared function output has to be performed which is a single efficient local MPC operation. For the other three mechanisms (cf. §4.2-4.3), an additional rounding and scaling is needed, but this only needs between 9 and 55 ms in our MPC-based Laplace mechanisms (cf. Tab. 1) and between 8 ms and 1.1 seconds for the Gaussian mechanisms depending on the number of parties.

³<https://github.com/liangzhao2048/Securely-Realizing-Output-Privacy-in-MPC-using-Differential-Privacy.git>

Table 1: Total runtimes (ms) per generated noise value in LAN10 (10-Gbit/s with 1ms RTT) averaged over 10 protocols runs of our MPC-based DP mechanisms using B and Y-sharing among N parties. ✓ are protocols not susceptible to the finite precision attacks discussed in §3, while ✗ are vulnerable solutions. The best secure results are marked in bold. - denotes memory overflow.

Mechanisms	Security	Batch	N = 2			N = 3			N = 5			
			Prot.	Offline	Online	Prot.	Offline	Online	Prot.	Offline	Online	
Discrete	M_{DLap} [41]	✗	40	Y	37.05	1.48	B	225.89	4.24	B	467.88	17.13
	M_{DLap} (cf. §4.1, naive)	✓	1	Y	14051.48	44.46	B	313628.29	376.21	B	472429.32	382.26
	M_{DLap} (cf. §4.1, optimized)	✓	1	Y	1606.00	37.71	B	12266.09	356.29	B	17953.73	378.13
	M_{DLap} (cf. §4.1, naive)	✓	5	Y	8690.86	4.72	B	36109.00	18.07	B	—	—
	M_{DLap} (cf. §4.1, optimized)	✓	5	Y	1273.36	5.55	B	9662.55	15.46	B	—	—
	M_{DLap} (cf. §4.1, naive)	✓	40	Y	7235.65	1.46	B	—	—	B	—	—
	M_{DLap} (cf. §4.1, optimized)	✓	40	Y	991.78	1.41	B	—	—	B	—	—
	M_{DGauss} (cf. §4.1)	✓	1	Y	6222.30	32.09	B	12459.71	306.06	B	21169.71	414.90
M_{DGauss} (cf. §4.1)	✓	5	Y	5584.64	31.48	B	—	—	B	—	—	
Continuous	M_{Lap} [41]	✗	30	Y	62.84	6.57	B	361.54	18.43	B	652.93	20.53
	M_{Snap} (cf. §4.2, naive)	✓	30	Y	72.05	38.64	B	275.51	128.55	B	685.93	167.41
	M_{Snap} (cf. §4.2, optimized)	✓	30	Y	50.63	10.20	B	261.94	33.56	B	468.54	44.28
	M_{ISLap} (cf. §4.3, naive)	✓	30	Y	817.72	55.07	B	12880.42	142.37	B	16836.56	224.10
	M_{ISLap} (cf. §4.3, optimized)	✓	30	Y	95.05	9.52	B	793.43	32.61	B	1239.24	44.17
	M_{Gauss} [65]	✗	30	Y	106.92	7.15	B	397.21	18.07	B	689.55	29.96
	$M_{ISGauss}$ (cf. §4.3)	✓	2	Y	10696.00	408.37	B	97190.00	843.96	B	197842.00	1141.75
	$M_{ISGauss}$ (cf. §4.3)	✓	4	Y	7836.50	166.82	B	82901.92	492.96	B	—	—
$M_{ISGauss}$ (cf. §4.3)	✓	30	Y	4712.24	8.56	B	—	—	B	—	—	

ACKNOWLEDGMENTS

This project received funding from: the Independent Research Fund Denmark under grant ID 10.46540/2032-00122B (YOSO); the ERC under the European Union’s Horizon 2020 research and innovation program (grant agreement No. 850990 PSOTI). It was co-funded by the DFG within SFB 1119 CROSSING/236615297 and GRK 2050 Privacy & Trust/251805230.

REFERENCES

- [1] Martin Abadi, Andy Chu, Ian Goodfellow, H. Brendan McMahan, Ilya Mironov, Kunal Talwar, and Li Zhang. 2016. Deep Learning with Differential Privacy. In *CCS*.
- [2] Abbas Acar, Berkay Celik, Hidayet Aksu, Selcuk Uluagac, and Patrick McDaniel. 2017. Achieving Secure and Differentially Private Computations in Multiparty Settings. In *PAC*.
- [3] Gergely Ács and Claude Castelluccia. 2011. I Have a DREAM! (Differentially privatE smArt Metering). In *Information Hiding (IH)*.
- [4] Mehrdad Aliasgari, Marina Blanton, Yihua Zhang, and Aaron Steele. 2013. Secure Computation on Floating Point Numbers. In *NDSS*.
- [5] Toshinori Araki, Jun Furukawa, Yehuda Lindell, Ariel Nof, and Kazuma Ohara. 2016. High-Throughput Semi-Honest Secure Three-Party Computation with an Honest Majority. In *CCS*.
- [6] David Archer, Shahla Atapoor, and Nigel Smart. 2021. The Cost of IEEE Arithmetic in Secure Computation. In *LATINCRYPT*.
- [7] Gilad Asharov, Shai Halevi, Yehuda Lindell, and Tal Rabin. 2018. Privacy-Preserving Search of Similar Patients in Genomic Data. In *PETS*.
- [8] Krishna Athreya and Soumendra Lahiri. 2006. *Measure theory and probability theory*. Springer Texts in Statistics.
- [9] John Awoyemi, Adebayo Adetunmbi, and Samuel Oluwadare. 2017. Credit card fraud detection using machine learning techniques: A comparative analysis. In *International Conference on Computing Networking and Informatics (ICCN)*.
- [10] Borja Balle and Yu-Xiang Wang. 2018. Improving the Gaussian Mechanism for Differential Privacy: Analytical Calibration and Optimal Denoising. In *ICML*.
- [11] Elaine Barker and John Kelsey. 2015. *Recommendation for Random Number Generation Using Deterministic Random Bit Generators*. Special Publication NIST.
- [12] Amos Beimel, Kobbi Nissim, and Eran Omri. 2008. Distributed Private Data Analysis: Simultaneously Solving How and What. In *Advances in Cryptology – CRYPTO 2008*, David Wagner (Ed.). Springer Berlin Heidelberg.
- [13] James Bell, Kallista Bonawitz, Adrià Gascón, Tancrede Lepoint, and Mariana Raykova. 2020. Secure Single-Server Aggregation with (Poly)Logarithmic Overhead. In *CCS*.
- [14] Aner Ben-Efraim, Yehuda Lindell, and Eran Omri. 2016. Optimizing Semi-Honest Secure Multiparty Computation for the Internet. In *CCS*.
- [15] Igor Bilogrevic, Julien Freudiger, Emiliano De Cristofaro, and Ersin Uzun. 2014. What’s the Gist? Privacy-Preserving Aggregation of User Profiles. In *ESORICS*.
- [16] Timm Birka, Kay Hamacher, Tobias Kussel, Helen Möllering, and Thomas Schneider. 2022. SPIKE: secure and private investigation of the kidney exchange problem. *BMC Medical Informatics Decision Making* (2022).
- [17] Dan Bogdanov, Sven Laur, and Jan Willemson. 2008. Sharemind: A Framework for Fast Privacy-Preserving Computations. In *ESORICS*.
- [18] Jonas Böhrer and Florian Kerschbaum. 2021. Secure Multi-party Computation of Differentially Private Heavy Hitters. In *CCS*.
- [19] Lennart Braun, Daniel Demmler, Thomas Schneider, and Oleksandr Tkachenko. 2022. MOTION—A Framework for Mixed-Protocol Multi-Party Computation. In *TOPS*.
- [20] Niklas Büscher, Daniel Demmler, Stefan Katzenbeisser, David Kretzmer, and Thomas Schneider. 2018. HyCC: Compilation of Hybrid Protocols for Practical Secure Computation. In *CCS*.
- [21] Niklas Büscher, Andreas Holzer, Alina Weber, and Stefan Katzenbeisser. 2016. Compiling Low Depth Circuits for Practical Secure Computation. In *ESORICS*.
- [22] David Byrd and Antigoni Polychroniadou. 2020. Differentially private secure multi-party computation for federated learning in financial applications. In *ICAIF*.
- [23] Clément Canonne, Gautam Kamath, and Thomas Steinke. 2020. The Discrete Gaussian for Differential Privacy. In *NeurIPS*.
- [24] Silvia Casacuberta, Michael Shoemate, Salil Vadhan, and Connor Wagaman. 2022. Widespread Underestimation of Sensitivity in Differentially Private Libraries and How to Fix It. In *CCS*.
- [25] Octavian Patrino and Amitabh Saxena. 2010. Secure Computation with Fixed-Point Numbers. In *FC*.
- [26] Berkay Celik, David Lopez-Paz, and Patrick McDaniel. 2017. Patient-Driven Privacy Control through Generalized Distillation. In *PAC*.
- [27] Jeffrey Champion, Abhi Shelat, and Jonathan R. Ullman. 2019. Securely Sampling Biased Coins with Applications to Differential Privacy. In *CCS*.
- [28] Hubert Chan, Elaine Shi, and Dawn Song. 2012. Privacy-Preserving Stream Aggregation with Fault Tolerance. In *FC*.
- [29] Melissa Chase, Ran Gilad-Bachrach, Kim Laine, Kristin E. Lauter, and Peter Rindal. 2017. Private Collaborative Neural Network Learning. *Cryptology ePrint Archive*, Paper 2017/762, <https://eprint.iacr.org/2017/762>.
- [30] Albert Cheu and Chao Yan. 2022. Necessary Conditions in Multi-Server Differential Privacy. In *ITCS*.
- [31] Christopher A. Choquette-Choo, Natalie Dullerud, Adam Dziedzic, Yunxiang Zhang, Somesh Jha, Nicolas Papernot, and Xiao Wang. 2021. CaPC Learning:

- Confidential and Private Collaborative Learning. In *ICLR*.
- [32] Christian Covington. 2019. Snapping Mechanism Notes. https://github.com/ctcovington/floating_point/blob/master/snapping_mechanism/notes/snapping_implementation_notes.pdf, Accessed 2022-09-29.
- [33] Ronald Cramer, Ivan Damgård, Daniel Escudero, Peter Scholl, and Chaoping Xing. 2018. $SPDZ_k$: Efficient MPC mod 2^k for Dishonest Majority. In *CRYPTO*.
- [34] Aref Dajani, Amy Langer, Phyllis Singer, Daniel Kifer, Jerome Reiter, Ashwin Machanavajhala, Simson Garfinkel, Scot Dahl, Matthew Graham, Vishesh Karwa, Hang Kim, Philip Leclerc, Ian Schmutte, William Sexton, Lars Vilhuber, and John Abowd. 2020. The modernization of statistical disclosure limitation at the U.S. Census Bureau. U.S. Census Bureau. <https://www2.census.gov/cac/sac/meetings/2017-09/statistical-disclosure-limitation.pdf>
- [35] Daniel Demmler, Ghada Dessouky, Farinaz Koushanfar, Ahmad-Reza Sadeghi, Thomas Schneider, and Shaza Zeitouni. 2015. Automated Synthesis of Optimized Circuits for Secure Computation. In *CCS*.
- [36] Daniel Demmler, Thomas Schneider, and Michael Zohner. 2015. ABY - A Framework for Efficient Mixed-Protocol Secure Two-Party Computation. In *NDSS*.
- [37] Luc Devroye. 1986. *Non-Uniform Random Variate Generation*. Springer Book Archive.
- [38] Cynthia Dwork, Krishnamurthy Kenthapadi, Frank McSherry, Ilya Mironov, and Moni Naor. 2006. Our Data, Ourselves: Privacy Via Distributed Noise Generation. In *CRYPTO*.
- [39] Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam D. Smith. 2006. Calibrating Noise to Sensitivity in Private Data Analysis. In *TCC*.
- [40] Cynthia Dwork and Aaron Roth. 2014. The Algorithmic Foundations of Differential Privacy. *Foundations and Trends in Theoretical Computer Science* (2014).
- [41] Fabienne Eigner, Matteo Maffei, Ivan Piryvalov, Francesca Pampaloni, and Aniket Kate. 2014. Differentially private data aggregation with optimal utility. In *ACSAC*.
- [42] Reo Eriguchi, Atsunori Ichikawa, Noboru Kunihiro, and Koji Nuida. 2021. Efficient Noise Generation to Achieve Differential Privacy with Applications to Secure Multiparty Computation. In *FC*.
- [43] Daniel Escudero, Satrajit Ghosh, Marcel Keller, Rahul Rachuri, and Peter Scholl. 2020. Improved Primitives for MPC over Mixed Arithmetic-Binary Circuits. In *CRYPTO*.
- [44] Vivian Fang, Lloyd Brown, William Lin, Wenting Zheng, Aurojit Panda, and Raluca Ada Popa. 2022. CostCO: An automatic cost modeling framework for secure multi-party computation. In *IEEE EuroS&P*.
- [45] Matt Fredrikson, Somesh Jha, and Thomas Ristenpart. 2015. Model Inversion Attacks That Exploit Confidence Information and Basic Countermeasures. In *CCS*.
- [46] Ivan Gazeau, Dale Miller, and Catuscia Palamidessi. 2013. Preserving differential privacy under finite-precision semantics. In *International Workshop on Quantitative Aspects of Programming Languages and Systems (QAPL)*.
- [47] Quan Geng, Peter Kairouz, Sewoong Oh, and Pramod Viswanath. 2015. The Staircase Mechanism in Differential Privacy. *IEEE Journal of Selected Topics in Signal Processing* (2015).
- [48] Arpita Ghosh, Tim Roughgarden, and Mukund Sundararajan. 2009. Universally utility-maximizing privacy mechanisms. In *STOC*.
- [49] Oded Goldreich, Silvio Micali, and Avi Wigderson. 1987. How to Play any Mental Game or A Completeness Theorem for Protocols with Honest Majority. In *STOC*.
- [50] Maoguo Gong, Jialun Feng, and Yu Xie. 2020. Privacy-enhanced multi-party deep learning. *Neural Networks* (2020).
- [51] Google. 2020. Secure noise generation. https://github.com/google/differential-privacy/blob/main/common_docs/Secure_Noise_Generation.pdf, Accessed 2022-09-29.
- [52] Google. 2022. Google's differential privacy libraries. <https://github.com/google/differential-privacy.git>, Accessed 2022-09-29.
- [53] Samuel Haney, Damien Desfontaines, Luke Hartman, Ruchit Shrestha, and Michael Hay. 2022. Precision-based attacks and interval refining: how to break, then fix, differential privacy on finite computers. *Journal of Privacy and Confidentiality* (2022).
- [54] Jamie Hayes, Luca Melis, George Danezis, and Emiliano De Cristofaro. 2019. LOGAN: Membership Inference Attacks Against Generative Models. In *PETS*.
- [55] Mikko Heikkilä, Eemil Lagerspetz, Samuel Kaski, Kana Shimizu, Sasu Tarkoma, and Antti Honkela. 2017. Differentially private Bayesian learning on distributed data. In *NeurIPS*.
- [56] Muhammad Ishaq, Ana L. Milanova, and Vassilis Zikas. 2019. Efficient MPC via Program Analysis: A Framework for Efficient Optimal Mixing. In *CCS*.
- [57] Matthew Jagielski, Jonathan Ullman, and Alina Oprea. 2020. Auditing Differentially Private Machine Learning: How Private is Private SGD?. In *NeurIPS*.
- [58] Bargav Jayaraman, Lingxiao Wang, David Evans, and Quanquan Gu. 2018. Distributed Learning without Distress: Privacy-Preserving Empirical Risk Minimization. In *NeurIPS*.
- [59] Jiankai Jin, Eleanor McMurtry, Benjamin Rubinstein, and Olga Ohrimenko. 2022. Are We There Yet? Timing and Floating-Point Attacks on Differential Privacy Systems. In *IEEE S&P*.
- [60] Marc Joye and Benoît Libert. 2013. A Scalable Scheme for Privacy-Preserving Aggregation of Time-Series Data. In *FC*.
- [61] Peter Kairouz, Ziyu Liu, and Thomas Steinke. 2021. The Distributed Discrete Gaussian Mechanism for Federated Learning with Secure Aggregation. In *ICML*.
- [62] Seny Kamara and Mariana Raykova. 2011. Secure outsourced computation in a multi-tenant cloud. In *IBM Workshop on Cryptography and Security in Clouds*.
- [63] Shiva Prasad Kasiviswanathan, Homin K. Lee, Kobbi Nissim, Sofya Raskhodnikova, and Adam Smith. 2008. What Can We Learn Privately?. In *FOCS*.
- [64] Marcel Keller. 2020. MP-SPDZ: A versatile framework for multi-party computation. In *CCS*.
- [65] Brian Knott, Shobha Venkataraman, Awni Y. Hannun, Shubho Sengupta, Mark Ibrahim, and Laurens van der Maaten. 2021. CryptTen: Secure Multi-Party Computation Meets Machine Learning. In *NeurIPS*.
- [66] D-U Lee, John D Villaseñor, Wayne Luk, and Philip Heng Wai Leong. 2006. A hardware Gaussian noise generator using the Box-Muller method and its error analysis. *IEEE transactions on computers* (2006).
- [67] Peter W. Markstein. 2008. The New IEEE-754 Standard for Floating Point Arithmetic. In *Numerical Validation in Current Hardware Architectures*.
- [68] George Marsaglia and Thomas A Bray. 1964. A convenient method for generating normal variables. In *SIAM Review*.
- [69] George Marsaglia and Wai Wan Tsang. 2000. The ziggurat method for generating random variables. *Journal of statistical software* (2000).
- [70] Brendan McMahan, Galen Andrew, Ulfar Erlingsson, Steve Chien, Ilya Mironov, Nicolas Papernot, and Peter Kairouz. 2018. A General Approach to Adding Differential Privacy to Iterative Training Procedures. arXiv, <http://arxiv.org/abs/1812.06210>.
- [71] Ilya Mironov. 2012. On significance of the least significant bits for differential privacy. In *CCS*.
- [72] Payman Mohassel and Peter Rindal. 2018. ABY³: A Mixed Protocol Framework for Machine Learning. In *CCS*.
- [73] Payman Mohassel and Yupeng Zhang. 2017. SecureML: A System for Scalable Privacy-Preserving Machine Learning. In *IEEE S&P*.
- [74] Arvind Narayanan and Vitaly Shmatikov. 2008. Robust De-anonymization of Large Sparse Datasets. In *IEEE S&P*.
- [75] Arpita Patra, Thomas Schneider, Ajith Suresh, and Hossein Yalame. 2021. SynCirc: Efficient Synthesis of Depth-Optimized Circuits for Secure Computation. In *HOST*.
- [76] Martin Pettai and Peeter Laud. 2015. Combining Differential Privacy and Secure Multiparty Computation. In *CCS*.
- [77] Vibhor Rastogi and Suman Nath. 2010. Differentially private aggregation of distributed time-series with transformation and encryption. In *SIGMOD*.
- [78] Mike Rosulek and Lawrence Roy. 2021. Three Halves Make a Whole? Beating the Half-Gates Lower Bound for Garbled Circuits. In *CRYPTO*.
- [79] Thomas Schneider and Michael Zohner. 2013. GMW vs. Yao? Efficient Secure Two-Party Computation with Low Depth Circuits. In *FC*.
- [80] Adi Shamir. 1979. How to Share a Secret. *Commun. ACM* (1979).
- [81] Elaine Shi, T.-H. Hubert Chan, Eleanor Gilbert Rieffel, Richard Chow, and Dawn Song. 2011. Privacy-Preserving Aggregation of Time-Series Data. In *NDSS*.
- [82] Elaine Shi, T.-H. Hubert Chan, Eleanor Gilbert Rieffel, and Dawn Song. 2017. Distributed Private Data Analysis: Lower Bounds and Practical Constructions. *ACM Transactions on Algorithms* (2017).
- [83] Reza Shokri and Vitaly Shmatikov. 2015. Privacy-preserving Deep Learning. In *CCS*.
- [84] Latanya Sweeney. 1997. Weaving Technology and Policy Together to Maintain Confidentiality. *The Journal of Law, Medicine & Ethics* (1997).
- [85] Jun Tang, Aleksandra Korolova, Xiaolong Bai, Xueqiang Wang, and Xiaofeng Wang. 2017. Privacy Loss in Apple's Implementation of Differential Privacy on MacOS 10.12. arXiv, <http://arxiv.org/abs/1709.02753>.
- [86] Stacey Truex, Nathalie Baracaldo, Ali Anwar, Thomas Steinke, Heiko Ludwig, Rui Zhang, and Yi Zhou. 2019. A Hybrid Approach to Privacy-Preserving Federated Learning. In *CCS*.
- [87] Salil Vadhan. 2017. The Complexity of Differential Privacy. In *Tutorials on the Foundations of Cryptography*. Springer International Publishing, 347–450.
- [88] Royce Wilson, Celia Yuxin Zhang, William Lam, Damien Desfontaines, Daniel Simmons-Marengo, and Bryant Gipson. 2020. Differentially Private SQL with Bounded User Contribution. In *PETS*.
- [89] Genqiang Wu, Yeping He, JingZheng Wu, and Xianyao Xia. 2016. Inherit Differential Privacy in Distributed Setting: Multiparty Randomized Function Computation. In *IEEE Trustcom/BigDataSE/ISPA*.
- [90] Runhua Xu, Nathalie Baracaldo, Yi Zhou, Ali Anwar, and Heiko Ludwig. 2019. HybridAlpha: An Efficient Approach for Privacy-Preserving Federated Learning. In *Workshop on Artificial Intelligence and Security (AISec@CCS)*.
- [91] Andrew Chi-Chih Yao. 1986. How to Generate and Exchange Secrets. In *FOCS*.
- [92] Lihua Yin, Jiyuan Feng, Hao Xun, Zhe Sun, and Xiaochun Cheng. 2021. A Privacy-Preserving Federated Learning for Multiparty Data Sharing in Social IoTs. *IEEE Transactions on Network Science and Engineering* (2021).