Practical Two-party Computational Differential Privacy with Active Security

Fredrik Meisingseth[1⋆], Christian Rechberger[1], and Fabian Schmid[1]

Graz University of Technology, Graz, Austria
firstname.lastname@iaik.tugraz.at

Abstract. In this work we revisit the problem of using general-purpose MPC schemes to emulate the trusted dataholder in differential privacy (DP), to achieve the same accuracy but without the need to trust one single dataholder. In particular, we consider the two-party model where two computational parties (or dataholders), each with their own dataset, wish to compute a canonical DP mechanism on their combined data and to do so with active security. We start by remarking that available definitions of computational DP (CDP) for protocols are somewhat ill-suited for such a use-case, due to them either poorly capturing some strong security guarantees commonly given by general-purpose MPC protocols, or having too strict requirements in the sense that they need significant adjustment in order to be satisfiable by using common DP and MPC techniques. With this in mind, we propose a new version of simulation-based CDP, called SIM*-CDP, and prove it to be stronger than the IND-CDP and SIM-CDP and incomparable to SIM$^+$-CDP. We demonstrate the usability of the SIM*-CDP definition by showing how to satisfy it by the use of an available distributed protocol for sampling truncated geometric noise. Further, we use the protocol to compute two-party inner-products with CDP and active security, and with accuracy equal to that of the central model, being the first to do so. Finally, we provide an open-sourced implementation and benchmark its practical performance. Our implementation generates a truncated geometric sample in between about 0.035 and 3.5 seconds (amortized), depending on network and parameter settings, comparing favourably to existing implementations.

Keywords: Differential privacy, Multiparty computation, UC-security

1 Introduction

The study of differential privacy in various distributed settings has given rise to a plethora of new definitions of DP, such as DP in the *local model (LDP)* [50], the *shuffle model* [6, 17] and definitions with a computationally bounded adversary, giving guarantees of *computational DP (CDP)* [27, 4, 65]. Each of the definitions is subject to its own restrictions in the adversarial model and in the accuracy that can be achieved within them. For instance, it is well-studied that

⋆ Parts of this work was performed whilst at Know-Center, Graz, Austria.

in LDP, which is a computationally efficient model with very few trust assumptions, one must in some settings add much more noise than the standard central model of statistical DP (SDP)[1] [50, 33, 16, 4]. One recurring idea is that one can use general-purpose *multiparty computation (MPC)* techniques to *emulate* a trusted central dataholder and thus one may get the same accuracy as is in the central model without having to trust a central computational party [30, 17]. The troubles in realising this idea, which we can call *generic emulation of the dataholder (GED)*, are firstly that one must accept the, potentially, large computational costs of MPC and secondly that it is not necessarily clear how one should define DP in this new distributed and computational setting. In order to avoid or reduce the computational costs of using MPC, up until now, most of the works in this area have opted for considering passive adversaries [4, 31, 70], only allowing aggregate functions [20, 51] and/or requiring honest majorities [27]. We focus on the case of two parties[2], active (static) corruptions, and require efficient protocols[3] for non-aggregate functionalities that achieve the same accuracy as in the central model. This work consists of two main parts; In the first one, we consider existing definitions of CDP for the setting above, conclude that they leave some things to be wished for and we thus propose an adjusted definition of CDP. In the second part we then turn to implementing an existing protocol for noise sampling [31], prove that it fulfills our new definition (but not some previous ones) and show that when augmented to use a mixed-circuit approach, it is efficient also in practice.

Definitions of CDP. In order to design practical protocols for GED, we would want a CDP notion that is directly compatible with the security notions of state-of-the-art MPC schemes and that allows the emulated dataholder to compute common SDP mechanisms. Importantly, however, key SDP mechanisms such as the Laplace [29], geometric [38], Gaussian [30] and discrete Gaussian [14] mechanisms cannot be computed exactly in strict probabilistic polynomial time (PPT) on a finite computer, due to having probability densities not being an inverse polynomial power of 2.[4] This means that, since general-purpose MPC requires PPT computable functionalities, the used definition needs to allow either that the protocol does not exactly emulate the dataholder (imperfect correctness) or that the emulated dataholder does not exactly compute the SDP mechanism, or

[1] Throughout this work we use 'DP' to refer to definitions that are both statistical (information-theoretic) and computational. When distinguishing between them we use 'SDP' and 'CDP' respectively.

[2] We consider both in the discussion about definition and in that of protocols only the case of two-party computation, although since all the tools we use are also applicable to settings with more parties (and all definitions can trivially be extended to those settings), we will continue to speak of MPC at times. At all times, the reader can suitably think of the special case of two parties whenever MPC is mentioned.

[3] In particular, we require that the protocols are computable in strict polynomial time in a finite computational model, as suggested in [2].

[4] For more details, see Appendix A.

both. Further, since we consider the case of two parties and active corruptions, for which information-theoretic general-purpose MPC is impossible [19, 39, 34], the only candidates of a suitable DP definition are computational [65, 4]. Since we will refer to it recurrently, let us call the paper [65] *MPRV*, after its authors.

In Section 3, we revisit the CDP definitions for two-party protocols by MPRV. They are all applicable to the setting of active corruptions however we find that they all fit unnaturally to the task of GED. For IND-CDP (Definition 6) and SIM-CDP (Definition 7), the inconvenience lies in that by using MPC to compute an SDP mechanism, one gets a protocol with much stronger guarantees than is captured in those CDP definitions, such as guarantees of security and correctness. This creates the need to analyse the desired properties of the protocol (now correctness, accuracy, security and CDP) separately, contrary to the custom when it comes to general-purpose MPC, which is typically analysed in the ideal/real paradigm where all such properties are formulated and asserted simultaneously. Intuitively, this ill-fitting is due to there not being a separation in the definitions between the protocol, which we want to be efficient, and the ideal DP mechanism, which we want to allow to be inefficient. For SIM$^+$-CDP (Definition 8) there is no such dissonance in the modeling of the protocol since it is already formulated using the ideal/real paradigm. Here the troubles lie rather in the details of the definition, which we will see are too restrictive to allow the notion to be fulfilled by emulating most common SDP mechanisms. This is fundamentally due to SIM$^+$-CDP requiring perfect correctness in the MPC protocol, which together with a demand for protocols running in strict polynomial time rules out any SDP mechanism that uses noise that is not samplable exactly in strict polynomial time. Whereas SIM$^+$-CDP could be achieved by using a finite version of standard SDP mechanisms, for instance using the mechanisms introduced in [2], it does mean a less direct realisation of GED, since the intuition is still to, say, 'use MPC to run the geometric mechanism'. Therefore, in Section 4, we propose an adapted version of SIM$^+$-CDP, calling it SIM*-CDP, which indeed can be satisfied by emulating standard CDP mechanisms due to a relaxation to computational correctness. Other large changes from SIM$^+$-CDP include using the UC (Universal Composability) security framework [11] instead of standalone security [10, 39] and allowing other ideal functionalities than secure function evaluation.[5] We prove that SIM*-CDP is of incomparable strength to SIM$^+$-CDP, meaning that there are in both ways computational tasks that can be solved with one but not the other, and that (like SIM$^+$-CDP) SIM*-CDP is strictly stronger than SIM-CDP and IND-CDP.

[5] We underscore that the merit of our new definition is not that it allows studying new scenarios or is to be preferred over previous definitions in all cases, indeed there are many cryptographic tasks for which UC-secure protocols are missing or for which it is not the most desirable framework to use. Rather the merit is that for settings where UC-secure protocols are readily available, then we have a formulation that takes advantage of that to give results that are both stronger and easier to obtain.

Implementing a Protocol Satisfying the New Definition. To demonstrate the advantages of SIM*-CDP, we implement a generic protocol for satisfying SIM*-CDP for the ideal functionality computing the truncated geometric mechanism. In particular, we analyse the noise sampling protocol of [31], adjust it to use mixed circuits for improved efficiency and give a very direct proof that the resulting protocol satisfies SIM*-CDP. Further, we implement the protocol and thereby present the first implementation of the protocol of [31] and simultaneously the first implementation of the truncated geometric mechanism with active security. Finally, we show how to use the protocol for computing integer inner-products with CDP and accuracy equal to that of the central model and benchmark the implementation, showing its practical efficiency. This treatment might be of independent interest, perhaps primarily due to our considerations relating to that the function sensitivity of the inner-product is dependent on the input domain of the corrupted party, thus creating a need for input validation. We note that whilst the definitions of CDP remain relatively unchanged when going from passive to active corruptions, the concrete privacy proof of a given protocol often changes significantly (as does the practical efficiency of its implementation) thereby the simplicity of our analysis in this more complicated setting showcases the usability of SIM*-CDP.

Contributions:

- We identify aspects of existing CDP definitions that make them an unnatural fit to the approach of generic emulation of a central trusted dataholder that computes an inefficient SDP mechanism. Therefore, we present a new version of SIM$^+$-CDP, which we call SIM*-CDP, and formally relate it to previous definitions (Sections 3 and 4).
- We demonstrate the usability of the SIM*-CDP definition by showing how it can be satisfied with the truncated geometric mechanism by proving that the efficient MPC protocol by [31] for sampling geometric noise satisfies our definition (Sections 5 and 6).
- We improve the efficiency of the protocol by using mixed circuits and use the protocol to compute two-party inner-products with CDP and active security, to the best of our knowledge being the first to do so with accuracy equal to that in the central model. Our open-sourced implementation is the first implementation of the noise sampling protocol of [31]. We provide benchmarks of the implementation and thereby show that it is efficient in practice (Section 7).

Related works. The first work that aims to emulate a central trusted party for DP by use of MPC is *Our data, ourselves* [27], where a protocol is proposed for computing sums with security against active adversaries corrupting less than a third of the parties, a part of which is a method for distributed noise generation. Following [27], other works have also proposed noise sampling protocols for DP in an MPC setting [1, 15, 31, 73] and the work most related to ours is

EIKN [31, 32]. EIKN gives an efficient MPC protocol for sampling an approximate truncated geometric distribution, which we use in this work. The privacy proofs given there are however only for honest majorities and thus do not apply to the two-party case [31, 32]. In a recent preprint [52] efficient noise sampling protocols for passive corruptions and dishonest majorities are provided. It is noted in passing that the protocols can easily be made secure against active adversaries by implementing them in a framework with active security but the type of CDP this would result in is not discussed. Our proposed SIM*-CDP definition offers an immediate answer to that. The work of [15] proposes a method for performing Bernoulli trials that is asymptotically superior to the one we use however their method relies on implementing oblivious data structures hence making it unsuitable for direct combination with the secret-sharing-based MPC schemes that we use.

Another line of work that is of relevance to ours due to it dealing with combining definitions of security for MPC schemes and DP is the string of papers considering MPC with differentially private leakage [47, 61, 44], where the idea is to improve the efficiency of an MPC protocol by allowing the protocol execution itself (not the result) to leak some extra information, but to restrict this leakage to be differentially private. Whilst the task solved in this line of work is quite different from the one we study, there are similarities in the formalities, which we discuss after having introduced SIM*-CDP.

Finally, we note that the large line of work on *differentially private data collection*, arguably centered around protocols related to the non-DP aggregation protocol Prio [20, 68, 5, 51], is conceptually related due to being an important topic on distributed CDP protocols. On the practical side, however, most of the techniques used there are not applicable to our work, due to them requiring the simpler setting of passive corruptions and/or aggregate functionalities. This includes techniques for distributed noise generation via infinitely divisible distributions[6], since it restricts the system to only compute aggregate functions, operate under the assumption of passive corruptions or have the data subjects (also called *clients*) encode their data points with respect to the function that is then to be computed. As expanded on in the next section, we focus primarily on *joint computation* of a DP mechanism where each computational party has a specific part of the input database to the mechanism in the clear, instead of having secret shares of the input as is typically the case for private data collection.

2 Preliminaries and Notation

For a natural number n, let $[n] := \{1, \ldots, n\}$. Let $\mathbb{N}^{-1}$ denote $\{1/n : n \in \mathbb{N}\}$.

[6] See, for instance, Section 2 in the phd thesis of Böhler [8].

2.1 Secure computation

We now briefly introduce necessary terminology regarding secure multiparty computation, for a slightly more thorough introduction, see Appendix B. A protocol is described as a set of interactive machines. For our purposes, it is not important exactly how those machines are formalised but for concreteness, we will think of interactive Turing Machines, which are non-uniform unless otherwise stated. We say that an algorithm (or machine or protocol) is *efficient* if it is PPT, meaning it is probabilistic and runs in strict polynomial time. We quantify the security by a protocol by a computational security parameter κ.[7] We consider both active and passive corruptions but assume they are static. For a protocol $\pi = \{P_1, ..., P_n\}$ and a set $C \subset [n]$, let $\{P_C\}$ denote $\{P_i : i \in C\}$ and let $\{P_{-C}\}$ denote $\{P_i : i \notin C\}$. The information available to the coalition C of parties in the protocol is formalised in their view, as defined below. The reason for exchanging $\{P_C\}$ with $\{\tilde{P}_C\}$ is to model active corruptions.

Definition 1 (VIEW, reformulation from [4]). *Let $\pi = \{P_1, ..., P_n\}$ be a protocol and $\mathcal{A}$ be an adversary corrupting a set $C \subset [n]$ of parties. For fixed inputs $D = (D_1, ..., D_n) \in \mathcal{D}$, the view in π of the corrupted parties $\{\tilde{P}_C\}$, denoted $\mathsf{VIEW}^{\mathcal{A}}_{\pi,C}(D)$, is defined as the random variable containing the inputs of the parties in C, their random coins and the messages that they receive during the execution of the protocol $\{\tilde{P}_C\} \cup \{P_{-C}\}$ on inputs D. The randomness is over the random coins of the honest parties $\{P_{-C}\}$.*

Often it is clear from context what parties the adversary corrupts (for instance in a symmetric two-party protocol) and then we omit C from notation. For defining secure computation of protocols we use the standard definitions in the ideal/real-paradigm, in both the standalone [10, 39, 58] and UC frameworks [11, 22]. Very shortly one can say that security is defined by formulating an ideal world in which an incorruptible trusted central party, an *ideal functionality*, performs all computations (and is secure by definition) and then a real-world protocol is deemed secure if no efficient distinguisher can distinguish it from the ideal world.

2.2 Differential Privacy

The notion of *differential privacy (DP)* [29, 26] considers a probabilistic algorithm, or *mechanism*, that maps *databases*, i.e. sets of elements from some data universe χ, to some output range R. We think of databases as ordered sets of some fixed (public) size N, and thus a database D is an element of $\mathcal{D} := \chi^N$. We say that two databases D, D' are *adjacent* if they differ in at most one element. There are however many other adjacency notions that may be more suitable to a given use case, perhaps especially when considering different types of distributed

[7] Many of our results will be quantified by κ even if they also hold with respect to an equal *statistical* security parameter since statistical security implies computational security.

settings, so it is important to note that the definition of (S)DP in itself is agnostic to the choice of adjacency notion, just like all CDP definitions considered in this paper.[8]

Definition 2 (Adjacency notion). *An adjacency notion* ADJ *on the dataset domain* $\mathcal{D}$ *is a set in* $\mathcal{D} \times \mathcal{D}$ *that is symmetric, i.e. if* $(D, D') \in \mathsf{ADJ}$ *then so is* (D', D)*, and* $\forall D \in \mathcal{D}, (D, D) \in \mathsf{ADJ}$*. If* $(D, D') \in \mathsf{ADJ}$ *the we say that* D *and* D' *are adjacent with respect to* ADJ*.*

We recall the standard definition of SDP (reformulation of [26]):

Definition 3 ((ε, δ)-SDP [29, 26]). *A probabilistic algorithm* $\mathcal{M} : \mathcal{D} \to R$ *is* (ε, δ)*-differentially private (SDP) if for all pairs* (D, D') *of adjacent databases in* $\mathcal{D}$ *and all subsets* S *of* R*,*

$$\mathbb{P}(\mathcal{M}(D) \in S) \leq e^{\varepsilon}\mathbb{P}(\mathcal{M}(D') \in S) + \delta, \tag{1}$$

where the probability is over $\mathcal{M}$*'s internal coin tosses.*

We often allow DP parameters to depend on the security parameter κ, letting $\varepsilon_\kappa = \varepsilon(\kappa), \delta_\kappa = \delta(\kappa)$ denote sets of parameters. Then we abuse notation by saying (for instance) that the ensemble $\mathcal{M} = \{\mathcal{M}_\kappa\}_{\kappa \in \mathbb{N}}$ is $(\varepsilon_\kappa, \delta_\kappa)$-SDP if for all large enough κ, $\mathcal{M}_\kappa$ is $(\varepsilon_\kappa, \delta_\kappa)$-SDP. The formulation of SDP above is often called *approximate SDP*, whereas it is called *pure SDP* if δ is fixed to 0. DP is typically studied in what is called the *central model*, of which an illustration can be found in Figure 1. In the central model, the database is simply a set of rows, each of which consists of information about one individual, called a *data subject*. These data subjects send their data to a trusted *dataholder* (without noise) that then computes a mechanism on the accumulated data and then releases the result to an untrusted *data analyst*. In this work, we rather consider DP in the *two-party model* [65, 62] where each data subject holds two database rows (x_i, y_i), each of which is sent to one of two computational parties (or *servers*) that then store their respective row into their database ($\mathbf{x}$ and $\mathbf{y}$ respectively) in the clear. Then these two servers together wish to compute the query f on the concatenation of their databases $D := \mathbf{x}||\mathbf{y}$, both learning the result, and they wish to do this in a differentially private manner with respect to their database. An illustration of this model can be seen in Figure 2.[9] When discussing DP mechanisms, it is critical to consider the usefulness of the mechanism for approximating the query

[8] In particular, since the focus of this paper is on the CDP definition with respect to GED and the choice of adjacency notion does not influence the merit of a CDP definition over another, it is for our discussions not relevant what adjacency notion one chooses to work with. On the practical side, there is however a potentially large difference in how well an adjacency notion fits together with a given MPC technique or setting, but these matters concern the realisation of GED rather than what the approach of GED means for the definition of CDP.

[9] We note that the two-party model is slightly but significantly different from the *two-server/multi-server models* [5, 18], primarily in that those models do not allow any server to have any part of the input dataset in the clear. This difference is of practical

function f. We do this by using the following notion of usefulness, as defined via a utility function.

Definition 4 (Utility function [7, 37]). *A utility function is an efficiently computable deterministic function $u : \mathcal{D} \times \mathcal{R} \to \{0,1\}^*$. A mechanism $\mathcal{M} : \mathcal{D} \to \mathcal{R}$ is α-useful for u if for all $D \in \mathcal{D}$:*

$$\mathop{\mathbb{P}}_{z \leftarrow \mathcal{M}(D)}(u(D,z) = 1) \geq \alpha. \tag{2}$$

A mechanism α-useful for u is said to solve the task (α, u).

A specific utility function we will consider is that which induces the notion of $(s,\alpha)-$*additive-usefulness* for a query function f, namely $u(D,z) = 1$ iff $|f(D) - z| \leq s$. Many popular DP mechanisms (such as the Gaussian, Laplace and geometric mechanisms) work computing the query function and then add noise of a specific distribution calibrated after the *sensitivity* of f (how much any single database entry can change the function evaluation). In this work, we consider this change only in the sense of l_1-distance.

Definition 5 (l_1-sensitivity). *Let $f : \mathcal{D} \to \mathcal{R}$ be a deterministic function, where $\mathcal{R}$ is a vector space on which the l_1-norm $||\mathbf{v}||_1 := \sum_i |v_i|$ is defined, and* ADJ *be an adjacency notion on $\mathcal{D}$. The l_1-sensitivity of f with respect to* ADJ*, denoted Δf, is defined as*

$$\Delta f := \max_{(D,D') \in \mathsf{ADJ}} ||f(D) - f(D')||_1. \tag{3}$$

2.3 Mixed Binary-arithmetic MPC Schemes

In our definitions, we rely on general-purpose MPC schemes with active security. In particular, we work with MPC protocols with restricted computation domain, either in $\mathbb{F}_p$ for arithmetic or $\mathbb{F}_{2^k}$ for binary circuits. For a discussion of active security in these schemes, we refer to Appendix E. In general, MPC schemes in $\mathbb{F}_p$ provide fast algorithms for addition and multiplication. In contrast, in $\mathbb{F}_{2^k}$, comparisons, bit-wise operations, and non-linear functions can be evaluated cheaply. However, storing larger integers results in substantial overhead, and evaluating arithmetic circuits in the binary domain incurs costs depending on the encoded values' bit size.

---

relevance because it means the models are suitable for different scenarios. The two-party model is mostly meant for *joint computation* between two entities each holding their own dataset (which may have been collected over time and without respect to the function evaluation in question) whereas the two-server model is rather tailored towards *data collection*, where one or more entities are collecting the data specifically for the purpose of performing the computation but wish to do so in a way that they never see any part of the dataset in the clear.

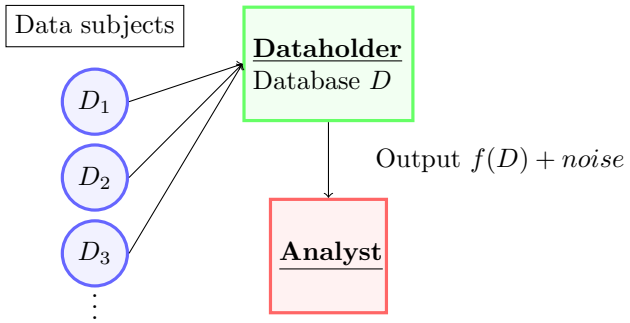

Fig. 1: In the *central model*, the data subjects trust the data holder with their data (D_i) but wish to keep it secret from an (possibly adversarial) analyst learning the (possibly noisy) function evaluation.

Several works have proposed solutions to convert shares between computation domains. First, in ABY [25], the authors propose a semi-honest two-party MPC scheme that allows switching between the binary, arithmetic, and garbled circuit domains (Garbled Circuits allow computation of binary circuits with low communication rounds). More recently, Rotaru and Wood introduced *doubly-authenticated bits* [67] and an efficient procedure to securely sample secret bits in the arithmetic and binary domain in malicious settings. Given the shares of an unknown random bit $(\llbracket b \rrbracket_2, \llbracket b \rrbracket_p)$ we can transfer shared bits from the binary to the arithmetic domain by computing the mask $m \leftarrow \mathsf{Reconstruct}(\llbracket x \rrbracket_2 \oplus \llbracket b \rrbracket_2)$ and setting $\llbracket x \rrbracket_p \leftarrow m + \llbracket b \rrbracket_p - 2 \cdot m \llbracket b \rrbracket_p$. Similarly, converting from arithmetic to binary masks the value by addition and evaluates subtraction in the binary domain. The conversion from the arithmetic to the binary domain gets more expensive, depending on the field size. Subsequent work introduced *extended doubly-authenticated bits (eda-bits)* [35], where masking values are shared along with their binary decomposition in the respective domains. The eda-bits represent an improvement in efficiency when converting larger values, and [35] presents dedicated protocols to speed up comparisons in $\mathbb{F}_p$.

3 CDP in the Two-party Model

We now briefly overview the literature on CDP in the two-party model and argue why it is desired to look for new definitions.

3.1 Existing CDP Definitions for Protocols

The formal study of both SDP and CDP in the two-party and multi-party models is initiated in [4, 65], where three definitions of two-party CDP are proposed. These are formulated for the two-party case but the definitions trivially extend to the multi-party case. We also follow this convention. The notion of SDP in

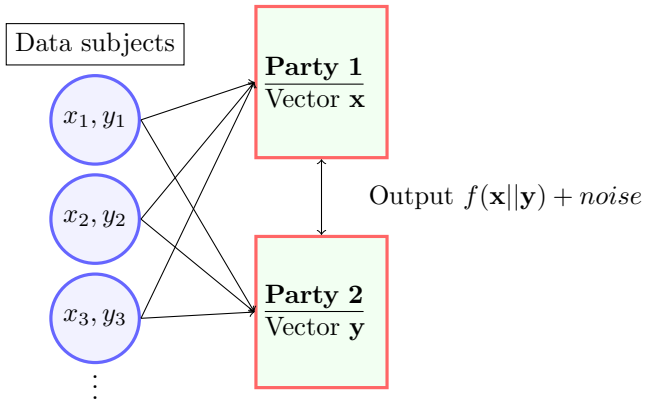

Fig. 2: In the *two-party model*, the data subjects trust two different data holders, which we call *parties*, with a different part of their data, but not with the part of the data that they send to the other data holder. In the end both parties learn the noisy function evaluation. Thus, in a sense, each party plays both the role of a data holder and a data analyst.

the central model is extended to interactive protocols by requiring that the view of the adversary is an SDP mechanism with respect to the input of the honest party. In [4] it is established that there are computational tasks for which the maximum utility in the two-party SDP model is strictly lower than in the central model and therefore there is a need to relax SDP to CDP. The CDP definitions come in two distinct variations, based on how they formalise a protocol execution 'looking SDP' to a computationally bounded party. The first variation is called *indistinguishability-based* and changes the demand that the output distributions of the mechanism are close on adjacent inputs to that this must only hold for all PPT distinguishers acting on the mechanism output. The second variation of CDP is called *simulation-based* and here a mechanism is deemed CDP if there exist an SDP mechanism from which it is computationally indistinguishable. Below we include a reformulation of the definition of indistinguishability-based CDP.[10]

Definition 6 (IND-CDP for protocols, reformulation from [4, 70]). *We say that a 2-party protocol π is $(\varepsilon_\kappa, \delta_\kappa)$-IND-CDP if for all efficient adversaries $\mathcal{A}$ corrupting at most one party, for all efficient distinguishers T, every sufficiently large κ and for all D, D' adjacent with respect to the inputs of the honest party,*

[10] Its original formulations [4, 65] differ slightly from one another, for instance in that [4] allows only passive corruptions and that MPRV lets the distinguisher be non-uniform. We consider these differences however to be of the sort making the definitions more two different instantiations of the same definition rather than two different ones. Note also that they both fix δ_κ as negligible (but non-zero) in κ.

we have

$$\mathbb{P}\left(T\left(\mathsf{VIEW}_{\pi}^{\mathcal{A}}(D)=1\right)\right) \leq e^{\varepsilon_\kappa}\mathbb{P}\left(T\left(\mathsf{VIEW}_{\pi}^{\mathcal{A}}(D')=1\right)\right)+\delta_\kappa. \tag{4}$$

The probabilities are taken over the randomness in π, $\mathcal{A}$ and T.

It is noted in MPRV that if $\delta_\kappa = 0$ then IND-CDP is equivalent to pure SDP. In the two-party model there are two different main formulations of simulation-based CDP, which we include below.[11]

Definition 7 (SIM-CDP for protocols, reformulation from MPRV). *We say that a 2-party protocol π is $(\varepsilon_\kappa, \delta_\kappa)$-SIM-CDP if for all efficient adversaries $\mathcal{A}$ corrupting at most one party, for all efficient distinguishers T and D, D' adjacent with respect to the inputs of the honest party, there exists an ensemble $\{\mathcal{M}_\kappa(\cdot)\}_{\kappa\in\mathbb{N}}$ of $(\varepsilon_\kappa, \delta_\kappa)$-SDP mechanisms $\mathcal{M}_\kappa : \mathcal{D} \to \mathcal{R}_\kappa$ such that for every sufficiently large κ and every $D \in \mathcal{D}$ of size polynomial in κ, it holds that $\mathsf{VIEW}_{\pi}^{\mathcal{A}}(D)$ and $\mathcal{M}_\kappa(D)$ are indistinguishable to T.*

Definition 8 (SIM$^+$-CDP, Reformulation of MPRV). *Let u be a utility function. A 2-party protocol π is $(\alpha, \varepsilon_\kappa, \delta_\kappa)$-SIM$^+$-CDP for u if there exists an $(\varepsilon_\kappa, \delta_\kappa)$-SDP mechanism $\mathcal{M}$ such that:*

- *the mechanism $\mathcal{M}$ is α-useful for u;*
- *π is a secure protocol for the functionality $\mathcal{M}$ as per Definition 17 (standalone security with perfect correctness, efficient protocols and a potentially inefficient simulator).*

3.2 Relations Between CDP Definitions

There is substantial literature on how the CDP definitions relate to each other and although the relations are far from tightly characterised, the rough picture is quite clear. For the parameter regimes the definitions were originally proposed ($\delta_\kappa = \mathsf{negl}(\kappa)$ in IND-CDP and $\delta_\kappa = 0$ in the others), it was shown in MPRV that any protocol that is $(\varepsilon_\kappa, 0)$-SIM$^+$-CDP is also $(\varepsilon_\kappa, 0)$-SIM-CDP and similarly $(\varepsilon_\kappa, 0)$-SIM-CDP implies $(\varepsilon_\kappa, \mathsf{negl}(\kappa))$-IND-CDP. On the other hand, there are tasks that can be solved with $(\varepsilon_\kappa, 0)$-SIM-CDP that cannot be solved with $(\varepsilon_\kappa, 0)$-SIM$^+$-CDP. It was long unknown if there is a similar separation between $(\varepsilon_\kappa, \mathsf{negl}(\kappa))$-IND-CDP and $(\varepsilon_\kappa, 0)$-SIM-CDP but in 2023 such a task was found [37]. The definitions have mostly been related to each other by either considering a fixed task and showing that there are no complexity assumptions under which that task can lead to a separation or by considering a fixed complexity assumption and showing that there are no tasks that lead to a separation under that assumption alone [43, 7, 63]. There is also a line of work about finding minimal complexity assumptions under which (various types of) CDP can be separated from SDP via a specific task, like computing boolean functions or

[11] These were introduced originally with $\delta_\kappa = 0$ and with non-uniform distinguishers.

integer inner-products with a given accuracy [42, 41, 45, 46, 55]. Whereas the relationship between the involved DP definitions is quite well understood in these cases, one should note that the same does not hold generally for other classes of tasks or for more relaxed parameter regimes.

3.3 Using Existing Definitions For GED

Each of the CDP notions above can be satisfied by a protocol that uses general-purpose MPC techniques to realise a functionality that computes an SDP mechanism, i.e. a protocol for GED. This has been shown for IND-CDP in [4, 70] and for SIM-CDP in MPRV. For SIM$^+$-CDP it is immediate, since there are general-purpose MPC schemes for the notion of secure computation used in SIM$^+$-CDP. We argue that GED results in guarantees that are fundamentally stronger than those in IND-CDP and SIM-CDP, therefore warranting a definition that captures them more closely, and that there are details in the SIM$^+$-CDP definition that make it inconvenient to work with for GED, although intuitively it is very suitable.

On using IND-CDP or SIM-CDP. One strength of GED for constructing CDP protocols is that one has guarantees about the behaviour of the protocol which exceed that of the adversarial view appearing DP. More precisely, the use of MPC allows guaranteeing *security* (dictating the influence an adversary may have on the protocol) and *correctness* (specifying the accuracy requirement of an honest execution). These properties are proven in the ideal/real paradigm (see Appendix B), that is, by specifying an ideal functionality that defines all of the desired properties of the protocol and then proving that the real protocol behaves almost the same. Here there arises a dissonance in intuition to the perspective in IND-CDP, since that notion considers only the real-world protocol. Therefore, when one uses IND-CDP together with GED, one has CDP as a property of the protocol in the real world, rather than in the ideal world with all other desired properties. This type of dissonance is smaller when it comes to SIM-CDP since the mechanism $\mathcal{M}$ in SIM-CDP is also in a way a description of the simulator and ideal functionality. The dissonance here, however, is that the formalisation of simulation is vastly relaxed in SIM-CDP, for instance in that the simulator has access to all private inputs.

On using SIM$^+$-CDP. SIM$^+$-CDP does not suffer the above-described modeling-wise dissonance and neither does it poorly capture the guarantees granted by secure computation. The problem with using SIM$^+$-CDP in the context of GED lies rather in that some of the details in the definition are too restrictive, meaning that they rule our realising GED for many of the most fundamental SDP mechanisms. In particular, the SIM$^+$-CDP definition requires the real-world protocol π to run in strict polynomial time and simultaneously have perfect correctness, meaning that its output distribution in an honest execution is *identical* to that of the SDP mechanism in the ideal functionality. This implies that any protocol

satisfying SIM$^+$-CDP must do so with respect to an SDP mechanism that can be computed *exactly* in *strict* polynomial time. Unfortunately, this rules out several commonly used SDP mechanisms, such as the Laplace or Gaussian mechanisms, which we now showcase with the example of the Laplace mechanism.

Impossibility of GED with the Laplace mechanism in SIM$^+$-CDP. The main question to ponder is whether there exists an efficient protocol that can realise the Laplace mechanism in SIM$^+$-CDP. Unfortunately, there is not.[12] To begin with, the support of the Laplace mechanism is the reals, meaning the output cannot even be written in strictly finite time. Thus we can note that any mechanism in the SIM$^+$-CDP definition must have a finite support. Further, even the (arguably) most Laplace-like such distribution, the geometric distribution [38] truncated to the output domain, cannot be realised in SIM$^+$-CDP in general, since it requires sampling probabilities that are not multiples of $2^{-poly(\kappa)}$ (for details on the impossibility of sampling certain distributions in strict polynomial time, see Appendix A). This means that in order to realise GED with distributions that cannot be sampled exactly in strict polynomial time (as is the case for the Laplace, geometric, Gaussian, discrete Gaussian distributions and truncated versions of them), there needs to be some slack introduced. This could be either in the shape of allowing a small distance between the output of the ideal functionality and that of the protocol (relaxing correctness) or relaxing the demand for strict polynomial time to expected polynomial time, as is argued in [14].

4 A New Version of Simulation-based CDP in the Ideal/real Paradigm

4.1 Our New Definition, SIM*-CDP.

We now propose a new version of SIM$^+$-CDP, which we call SIM*-CDP and then discuss its relationship to previous definitions further.

Definition 9 ($(\varepsilon_\kappa, \delta_\kappa)$−SIM*-CDP). *We say that the two-party protocol π is $(\varepsilon_\kappa, \delta_\kappa)$-SIM*-CDP for the ideal functionality $\mathcal{F}$ and adjacency notion* ADJ *if π UC-realises $\mathcal{F}$ and for all ideal-world adversaries $\mathcal{S}$, the view of $\mathcal{S}$ is $(\varepsilon_\kappa, \delta_\kappa)$-SDP with respect to* ADJ*.*

The main differences between SIM*-CDP and SIM$^+$-CDP are:

- *UC-security is used as security notion.*
- *The ideal functionality is variable (and can be reactive).*

[12] We note that this invalidates the claims in [1] of achieving SIM$^+$-CDP for the (continuous, untruncated) Laplace mechanism. The protocol there does however seem to satisfy a relaxation of SIM$^+$-CDP, in line with the contents of Section 4, although that remains to be formally shown.

- *Correctness is computational rather than perfect.*
- *The ideal-world adversary (simulator) must be efficient (strict PPT).*
- *The requirement of usefulness is removed from the CDP definition.*

We now expand on the motivation behind these changes.

Using UC-security. Although the standalone security framework is heavily used, in the last two decades the security analyses of many popular schemes have taken place in the more expressive UC framework [11]. The main merit of this framework is that the security can be proven to be preserved under arbitrary composition of protocols, leading to a stronger notion of security and an increased modularity in security proofs. Thus, using UC security in the CDP notion is natural for cases where this (stronger) type of security is already achieved by the MPC scheme one intends to use. Further, as we will see below, this change of security framework also directly leads to many other benefits.

A variable ideal functionality. The ideal functionality used in SIM$^+$-CDP is fixed to be that of *secure function evaluation (SFE)*, i.e. the parties jointly compute an SDP mechanism (with abort). With regards to capturing what it means for a protocol to be CDP generally, this is a significant restriction as compared to IND-CDP and SIM-CDP, where CDP is defined without dependence on the functionality of the protocol. In particular, both IND-CDP and SIM-CDP allow direct modeling of reactive functionalities, and as such our new definition arguably lies closer to those definitions conceptually than SIM$^+$-CDP does, in that it is applicable to more functionalities than those that can be expressed as SFE [44, 49, 48]. On the practical side, one relevant reactive functionality is that of SFE with differentially private leakage as in, for instance, [44]. More details about the setting of SFE with DP leakage are found in Appendix C.

Computational correctness. Another positive consequence of changing security framework is that the correctness of the protocol is now (as is standard in UC-security) computational rather than perfect. As explained in the previous section, this relaxation allows the ideal functionality to sample inefficiently samplable distributions and still have an efficient protocol that realises it.

Efficient simulators. As one main goal of our new definition is to have it align closely to common practice in MPC, we choose to require efficient simulation. Whereas this does make fulfilling the definition harder, it also makes the definition stronger.

Not including usefulness in the definition. A final difference between SIM*-CDP and SIM$^+$-CDP is that we choose not to include the requirement for usefulness in the definition of CDP itself. This is done primarily to more closely correspond to how the matter of usefulness is handled for IND-CDP and SIM-CDP

in MPRV, namely that the CDP definition is agnostic to the notion of usefulness (Definition 6 in MPRV [65]) and that usefulness is then added later (Definition 7 in MPRV). Another advantage of not having the usefulness as a part of the CDP definition is that one can choose to consider the usefulness simply of the ideal functionality (as is done in SIM$^+$-CDP) or to consider the usefulness of the protocol directly (as with IND-CDP and SIM-CDP in Definition 7 of MPRV) and then take, for instance, failure probabilities of the protocol into account.[13] To round this subsection off, we re-iterate the standard ideal functionality for SFE with abort, see Figure 3. In Section 6 we propose a protocol for realising this ideal functionality with the geometric mechanism as the functions f_1 and f_2 and prove it is SIM*-CDP in the presence of active corruptions.

Functionality $\mathcal{F}^f_{SFE}$

Parameters:

- A function $f = (f_1, f_2) : (\{0,1\}^*)^2 \to (\{0,1\}^*)^2$.

No corruptions:

- Upon $\mathbf{x}_1$ from P_1 and $\mathbf{x}_2$ from P_2, deliver $f_1(\mathbf{x}_1, \mathbf{x}_2)$ to P_1 and $f_2(\mathbf{x}_1, \mathbf{x}_2)$ to P_2.

Party P_c corrupted (P_h is honest):

- Upon $(\mathsf{Input}, \mathbf{x}_h)$ from P_h and $(\mathsf{Input}, \mathbf{x}_c)$ from P_c, send $f_c(\mathbf{x}_1, \mathbf{x}_2)$ to P_c.
- Upon $(\mathsf{Deliver}, b)$ from P_c, if $b = 1$ then send $f_h(\mathbf{x}_1, \mathbf{x}_2)$ to P_h, otherwise send $\perp$.

Fig. 3: The ideal functionality for SFE with abort.

4.2 Relating SIM*-CDP to Other Definitions

We now relate our new definition to existing ones. For all of the propositions, the proofs are delegated to Appendix D. We prove separations only when $\delta_\kappa = 0$ (or $\delta_\kappa = \mathsf{negl}(\kappa)$, depending on the CDP notion), as is common in the literature, and leave extending the separations to other settings for future work.

Relation to SIM$^+$-CDP. There is no general hierarchy between SIM$^+$-CDP and SIM*-CDP, in the sense that there are both tasks that can be solved with SIM$^+$-CDP but not SIM*-CDP and the other way around. In one direction this is due to SIM*-CDP being more restrictive in that it demands UC-security instead of standalone security since there are well known results of functionalities that

[13] For SIM$^+$-CDP one should note that the usefulness of the protocol is always the same as that of the ideal functionality unless there are active corruptions, due to the requirement of perfect correctness.

can be realised with standalone security but not UC-security unless certain setup assumptions are made [12]. In the other direction, SIM*-CDP is more relaxed than SIM+-CDP with regard to the correctness of the protocol. In more formal terms, see the propositions below.

Proposition 1. *Using the plain UC model, i.e. without setup assumptions, and assuming that enhanced trapdoor permutations (see [39]) exist, there exists ε_κ for which there exists a task that is solvable with $(\varepsilon_\kappa, 0)$-SIM+-CDP but not with $(\varepsilon_\kappa, 0)$-SIM*-CDP. This holds regardless of whether the utility requirement is placed on the real or the ideal protocol with respect to SIM*-CDP.*

Proposition 2. *Using the UC model with the setup assumption of a common reference string (CRS) (see, for instance, [12]) and with the utility in SIM*-CDP being considered in the ideal world (i.e. with regards to the utility of $\mathcal{F}$), there exists ε_κ for which there exists a task that is solvable with $(\varepsilon_\kappa, 0)$-SIM*-CDP but not with $(\varepsilon_\kappa, 0)$-SIM+-CDP.*

Relation to SIM-CDP and IND-CDP. Just as with SIM+-CDP, on the one side if a protocol is SIM*-CDP then it is SIM-CDP (and thus also IND-CDP) but on the other side there are tasks that can be solved with SIM-CDP but not in SIM*-CDP. The second separation is a direct corollary of Proposition 1 due to that all SIM+-CDP protocols also are SIM-CDP protocols with unchanged parameters.

Proposition 3. *For any parameters $\varepsilon_\kappa, \delta_\kappa$, if a two-party protocol π is $(\varepsilon_\kappa, \delta_\kappa)$-SIM*-CDP, then it is also $(\varepsilon_\kappa, \delta_\kappa)$-SIM-CDP and $(\varepsilon_\kappa, \delta_\kappa + \mathsf{negl}(\kappa))$-IND-CDP.*

Corollary 1 (of Proposition 1). *Using the plain UC model, i.e. without setup assumptions, and assuming that enhanced trapdoor permutations (see [39]) exist, there exists ε_κ for which there exists a task that is solvable with $(\varepsilon_\kappa, 0)$-SIM-CDP but not with $(\varepsilon_\kappa, 0)$-SIM*-CDP. This holds regardless of whether the utility requirement is placed on the real or the ideal protocol with respect to SIM*-CDP.*

Relation to MPC-with-DP-leakage definition in [44]. Within the literature on relaxing definitions of secure computation by allowing there to be non-negligible information leakage during protocol execution as long as this leakage is SDP (for a longer discussion on such protocols, see Appendix C), there is a definition (Definition 19) that, like SIM*-CDP, uses UC-security and defines a CDP property of a protocol. The fundamental difference between SIM*-CDP and that definition is that their definition is fixed for a given ideal functionality and only a specific part of the view of the ideal-world adversary $\mathcal{S}$ is required to be SDP, whereas SIM*-CDP is defined for arbitrary ideal functionalities and the entire view of $\mathcal{S}$ must be SDP. Thus SIM*-CDP can be seen as both a restriction of the definition in [44] (where one requires the remaining parts of $\mathcal{S}$'s view to be SDP also) and as a generalisation of it since the ideal functionality is left variable. From another point of view, the definitions try to solve two distinct

problems, but one can suitably consider the need we see to propose an alternative CDP definition to those of (say) IND-CDP and SIM$^+$-CDP as being the analog in general CDP to the motivation in [44] for giving a definition separate to those of [47, 61].

4.3 A More General Definition, SIM$^\circ$-CDP

The core idea of SIM$^+$-CDP and SIM*-CDP is the same (requiring the protocol to realise an SDP ideal functionality) and this opens up a wide space of such definitions since there is an abundance of different notions of secure computation in the MPC literature. One can for instance vary correctness, robustness or efficiency requirements for the different involved entities. This suggests a generalised definition of which SIM$^+$-CDP, SIM*-CDP and other natural variations are instantiations of. Below we formulate such a generalised definition and call it SIM$^\circ$-CDP.

Definition 10 (SIM$^\circ$-CDP). *We say that a two-party protocol π is $(\varepsilon_\kappa, \delta_\kappa)$-SIM$^\circ$-CDP with respect to ideal/real security notion* SEC *for the ideal functionality $\mathcal{F}$ and adjacency notion* ADJ *if π realises $\mathcal{F}$ in the sense of* SEC *and for all ideal-world adversaries $\mathcal{S}$, the view of $\mathcal{S}$ is $(\varepsilon_\kappa, \delta_\kappa)$-SDP with respect to* ADJ.

In light of this definition, the bulk of the discussion in this section can be seen as concerning the ways in which we regard the specific choice of security notion in SIM$^+$-CDP as being inconvenient with respect to GED. We are aware of only one other used instantiation of SIM$^\circ$-CDP and that is in [5] where SIM$^\circ$-CDP is instantiated using standard standalone security but with computational correctness. That CDP notion is stronger than SIM$^+$-CDP in that it requires efficient simulators but weaker in the sense of having relaxed correctness. The protocol presented in [5] is not SIM$^+$-CDP (imperfect correctness is needed) and neither is it SIM*-CDP (since no UC security proof is given).

5 A SIM*-CDP Version of the Geometric Mechanism

To demonstrate the use of our new definition, we now go through in detail how to satisfy it for the standard SFE ideal functionality with the truncated geometric mechanism as the function. Conceptually, this is very simple; one can simply use any PPT algorithm that samples a distribution with a sufficiently small statistical distance to a truncated geometric distribution and then compute that algorithm in MPC via some general-purpose, active secure, protocol. It is however worth considering hurdles that arise in the details, such as how to handle the mechanism's dependence on the query function, having a query function whose sensitivity depends on the inputs of both parties and the consequences of working over a finite field. One core step is, naturally, to sample a distribution that is close to a truncated geometric distribution. Sampling algorithms for such distributions can be found in [38, 2, 31], however, the truncation is to a range

between 0 and some fixed positive integer. The results and methods however extend to $\mathbb{Z}_q$, and general queries of bounded magnitude.

Definition 11 (Truncated geometric distribution). *Define the truncated geometric distribution $Z \sim Geo_{q,\lambda}(\bar{f})$ centered at $\bar{f} \in \mathbb{Z}_q$, truncated to $\mathbb{Z}_q := [\lceil -q/2 \rceil, \lfloor q/2 \rfloor)$, by its probability mass function:*

$$p_Z(z) = \frac{e^{1/\lambda} - 1}{e^{1/\lambda} + 1} e^{\frac{-|z-\bar{f}|}{\lambda}} \tag{5}$$

for $z \notin \{\lceil -q/2 \rceil, \lceil q/2 - 1 \rceil\}$, and

$$p_Z(z) = \frac{1}{e^{1/\lambda} + 1} e^{\frac{-|z-\bar{f}|}{\lambda}} \tag{6}$$

for $z \in \{\lceil -q/2 \rceil, \lceil q/2 - 1 \rceil\}$.

Definition 12 (Range-truncated geometric mechanism). *Let $\lambda \in \mathbb{N}^{-1}$ and let $f : \mathcal{D} \to \mathbb{Z}_q$ be a deterministic function. The Range-truncated geometric mechanism (RTGeo) over $\mathbb{Z}_q$ for f is defined as*

$$\mathcal{M}^{q,f,\lambda}_{RTGeo}(D) := Geo_{q,\lambda}(f(D)). \tag{7}$$

It is easy to verify that $\mathcal{M}^{q,f,\lambda}_{RTGeo}(D)$ is an $(\varepsilon, 0)$-SDP mechanism as long as $\lambda = \frac{\varepsilon}{\Delta f}$. In line with [2], we only allow $\lambda \in \mathbb{N}^{-1}$, in order to avoid the need to represent real numbers, and this also implies $\varepsilon \in \mathbb{N}^{-1}$. Whereas $\mathcal{M}_{RTGeo}$ gives SDP, it is inconvenient to sample the noise distribution directly, partly because it requires knowledge of $f(D)$ and partly because it may not be efficiently samplable. Therefore we consider the following mechanism.

Definition 13 (Subrange-truncated geometric mech.). *Let $B \in \{1, \ldots, \lceil q/2 \rceil - 1\}$ and $\lambda \in \mathbb{N}^{-1}$. Let the Subrange-truncated geometric (SRTGeo) mechanism over $\mathbb{Z}_q$ with noise truncation to $\mathbb{Z}_{2B}$, for a function $f : \mathcal{D} \to \mathbb{Z}_q$, be defined as $\mathcal{M}^{2B,f,\lambda}_{SRTGeo}(D) := f(D) + Geo_{2B,\lambda}(0)$, with the addition performed over $\mathbb{Z}_q$.*

In the simple lemma below we give a bound on the statistical distance between the two mechanisms we have introduced this far. The proof, as the proofs of all other lemmas, is found in Appendix D. We note that we need to introduce a bound on the absolute value of the query function, so as to not have the sensitivity of the function be affected by the modular arithmetics.

Lemma 1. *Let $f^{max} := \max_{D \in \mathcal{D}} |f(D)|$, $B \in \mathbb{N}$, $\lambda \in \mathbb{N}^{-1}$ and $q > 2f^{max} + 2B$. Then the statistical distance between $\mathcal{M}^{2B,f,\lambda}_{SRTGeo}(D)$ and $\mathcal{M}^{q,f,\lambda}_{RTGeo}(D)$ for all $D \in \mathcal{D}$ is at most $e^{-B/\lambda}$.*

We are now one step closer to a functionality that can be efficiently realised, since the noise sampling is no longer dependent on the function evaluation and the support of the noise is potentially much smaller than the entire $\mathbb{Z}_q$ and

the support of f. The trouble still remains that the probabilities might not be negative polynomial powers of two. In [27, 31] it is presented distributions that can be exactly sampled under this constraint and that have a small statistical distance from a truncated geometric distribution. We use the procedure FDL (*Finite-range Discrete Laplacian*) introduced in EIKN [31].

Definition 14 (FDL function and procedure). *Let* $\mathbf{r} \in \{0,1\}^{Bd+1}$ *be independent fair coins and* $0 < e^{-1/\lambda} < 1$. *Let* $\hat{\alpha}^1 \leftarrow \frac{1-e^{-1/\lambda}}{1+e^{-1/\lambda}}$ *and* $\hat{\alpha}^i \leftarrow 1 - \hat{\alpha}^1$ *for* $i = 2, ..., B$ *be public parameters. Let* $\oplus$ *and* $\wedge$ *denote addition and multiplication over the binary field and let* $\vee$ *be shorthand for computing the OR operation by using binary addition and multiplication. Let all other operands be defined as normally over* $\mathbb{Z}_q$. *Define the function* $\texttt{FDL}_{\lambda,B,d} : \{0,1\}^{Bd+1} \to \mathbb{Z}_{2B} \subseteq \mathbb{Z}_q$ *by the procedure in Algorithm 1. Let* $\alpha = (\alpha_1, \alpha_2, ...)$ *be the bit decomposition of* $\hat{\alpha}$. *The subprocedure* $\texttt{Ber}_{\hat{\alpha}} : \{0,1\}^d \times \{0,1\}^d \to \{0,1\}$ *for generating approximate Bernoulli trials with parameter* $\hat{\alpha}$ *using a randomness seed in* $\{0,1\}^d$ *is defined by the procedure in Algorithm 2.*

Procedure FDL
Input: $\mathbf{r} \in \{0,1\}^{Bd+1}$
1. Sample B approximate Bernoulli trials $\beta_i \leftarrow \texttt{Ber}_{\hat{\alpha}^i}((r_{d(j-1)+1}, ..., r_{dj}))$ for $i = 1, ..., B$.
2. For $i = 1, ..., B$: set $c_i \leftarrow \vee_{j=1}^{i} \beta_j$.
3. Set $l \leftarrow B - \sum_{i=1}^{B} c_i$.
4. Set $\sigma \leftarrow 2 \cdot r_{Bd+1} - 1$.
5. Output $\sigma \cdot l$.

Algorithm 1: The algorithm description for the FDL procedure.

Procedure Ber
Input: $\mathbf{r} \in \{0,1\}^d$, $\alpha \in \{0,1\}^d$
1. For $i = 1, ..., d$, set $c_i \leftarrow \alpha_i \oplus r_i$.
2. For $i = 1, ..., d$, set $e_i \leftarrow \vee_{j=1}^{i} c_j$.
3. For $i = 1, ..., d$, set $v_i \leftarrow e_i \oplus e_{i-1}$, with $e_0 \leftarrow 0$.
4. Set $\beta \leftarrow 1 \oplus_{i=1}^{d} (r_i \wedge v_i)$ and output β.

Algorithm 2: The algorithm description for the Ber procedure.

Note that FDL is an exact method for turning $Bd + 1$ fair coins into a sample of a distribution that is statistically close to a truncated geometric one. It is clear that if the number of fair coins is polynomial in κ then FDL runs in strict polynomial time. With some abuse of notation, we use FDL to denote both the

procedure and the probability distribution it generates upon being given fair coins.[14]

Definition 15 (FDL mechanism). *Let $B \in \{1, \ldots, \lceil q/2 \rceil - 1\}$. Let the Finite-range Discrete Laplace (FDL) mechanism over $\mathbb{Z}_q$ for a function $f : \mathcal{D} \to \mathbb{Z}_q$ be defined as $\mathcal{M}^{\lambda,B,d,f}_{\texttt{FDL}}(D) := f(D) + \texttt{FDL}_{\lambda,B,d}$, with the addition performed over $\mathbb{Z}_q$.*

The following lemma is proven in EIKN [31].

Lemma 2. *Let $f^{max} := \max_{D \in \mathcal{D}} |f(D)|$, $q > 2f^{max} + 2B$ and $B \in \{1, \ldots, \lceil q/2 \rceil - 1\}$. If FDL is given independent fair coins and all the arithmetics are done over $\mathbb{Z}_q$, then the statistical distance between $\mathcal{M}^{\lambda,B,d,f}_{\texttt{FDL}}(D)$ and $\mathcal{M}^{2B,f,\lambda}_{SRTGeo}(D)$ is at most $B \cdot 2^{-d}$.*

Further, we have that $\mathcal{M}^{q,f,\varepsilon/\Delta f}_{RTGeo}(D)$ is a useful approximation of f, as we show in the following lemma.

Lemma 3. *Let $q > 2f^{max} + 2B$, $B \in \{1, \ldots, \lceil q/2 \rceil - 1\}$. Let $f : \mathcal{D} \to \mathbb{Z}_q$ be an arbitrary deterministic function with $f^{max} := \max_{D \in \mathcal{D}} |f(D)|$ and let $\hat{f}(D) := \mathcal{M}^{q,f,\lambda}_{RTGeo}(D) : \mathcal{D} \to \mathbb{Z}_q$. Then $\hat{f}$ is $\left(\nu, \frac{2e^{-1/\lambda}}{e^{-1/\lambda}+1} e^{-\nu/\lambda}\right)$-additive-useful for f for any positive integer ν.*

6 A Protocol for the FDL Mechanism

From the previous section, we know that the FDL mechanism is statistically close to the Range-truncated geometric mechanism ($\mathcal{M}_{RTGeo}$), which is pure SDP, and that this holds under some restrictions on the query function and on the parameter choices. At the same time, it is immediate that $\mathcal{M}_{RTGeo}$ is statistically close to the untruncated geometric mechanism (i.e. when the noise is not truncated and that the modular arithmetics thus might cause overflows), as long as the value of the query function is somewhat far away from $\pm q/2$. Therefore, there is a choice to be made regarding which mechanism one chooses to have in the ideal functionality (call this the *ideal mechanism*), given that we will have the protocol compute the FDL mechanism via general-purpose MPC. The trade-off in this choice is that having $\mathcal{M}_{RTGeo}$ as the ideal mechanism will lead to $(\varepsilon_\kappa, 0)$-SIM*-CDP as long as the statistical distances mentioned above are negligible in κ, essentially having the statistical distance be dealt with as part of the correctness slack. On the other hand, this can be avoided by letting $\mathcal{M}_{\texttt{FDL}}$ be the ideal mechanism, thus leading to $(\varepsilon_\kappa, \delta_\kappa)-$SIM*-CDP where the statistical distance is rather incorporated into the δ_κ term. As having an ideal mechanism as close as possible to a standard SDP mechanism is to be seen as a more direct

[14] We also note that the requirement that $e^{-1/\lambda} < 1$ is equivalent to $\lambda > 0$, which is already guaranteed by $\lambda \in \mathbb{N}^{-1}$.

realisation of GED, we opt for having $\mathcal{M}_{RTGeo}$ as the ideal mechanism.

As stated in the preliminaries, we consider two-party computation schemes that operate in $\mathbb{F}_q$ with q being either a prime larger than 2 or a power of 2. We elaborate on active secure schemes for both domains in Appendix E. Implementing the `FDL` algorithm in either domain comes at a significant cost. Note that the `Ber` procedure and the first 2 steps of the `FDL` procedure consist of only binary arithmetics. However, the remainder of the `FDL` procedure consists of integer arithmetics. While there are protocols to evaluate the binary steps in the arithmetic domain, they are usually very costly. On the other hand, evaluating the whole algorithm in the binary domain comes with two problems: the summation and addition in binary would incur a significant cost, and second, the result would be a shared noise in the binary domain. Thus, applying the noise is limited to the binary domain. The mixed circuit approach (see Section 2.3) gives us a well-performing trade-off.

We accept inputs represented in the binary domain, perform all operations until the fourth step through a binary circuit, translate all shares to the arithmetic domain, and perform the rest of the operations through an arithmetic circuit. For each of these "phases", we use protocols introduced before. We use SPDZ_{2^k} [21] for the arithmetic computations, the FKOS protocol [36] for binary circuits and *daBits* (doubly-authenticated bits) [67] for translating between the domains. With correct parametrization, we can achieve the same security guarantees in different computation domains. Thus, the feasibility of the mixed circuit approach is easily tested. The mixed circuit approach is feasible if switching between circuits is cheaper than the computation overhead in either domain. In our application (Section 6.1), we will, as typically for DP applications, focus on arithmetic computations. Evaluating the `FDL` mechanism in the binary domain would, therefore, incur a cost that scales with the underlying application. For the arithmetic case, we have an additional cost of assuring all input ranges (e.g., assert that binary coins $\in \{0, 1\}$) and evaluate binary gates with arithmetic circuits. Section 7 has a longer discussion about input validation.

We describe our protocol using the *Arithmetic Black Box (ABB)*, which is an ideal functionality in the UC framework. Very roughly, the ABB is a functionality that can take inputs from the parties and compute linear combinations and multiplications between stored values and output stored values. We use a flavor of the ABB that can do these operations over $\mathbb{F}_{2^k}$ and $\mathbb{F}_q$. Additionally, the ABB can translate values stored as elements of the binary field to binary values within the larger field. More concretely, we use the formulation of the ABB that can be found in [35] and we include a definition of the ideal functionality in Appendix B.1. Our protocol is presented in Figure 4.
We are now ready to present our main theorem, namely that the protocol we have introduced indeed is $(\varepsilon_\kappa, 0)$-SIM*-CDP. Let $decomp(\lambda, d)$ be short for the bit-decomposition of λ truncated to d bits. The proof is found in Appendix D.7.

Protocol $\pi_{\mathcal{M}_{\texttt{FDL}}}$

Parameters: Natural numbers B, d, q, N, bit decomposition $\hat{\alpha}^1, ..., \hat{\alpha}^d$ and an efficiently computable function $f : \mathbb{Z}_q^{2N} \to \mathbb{Z}_q$, meaning it can be computed using polynomially many multiplications and linear combinations in $\mathbb{Z}_q$. Assume access to $\mathcal{F}_{ABB}$.

Initialisation:

1. Player i locally samples $Bd+1$ fair coins and stores them as $\mathbf{e}^i$.
2. Player i sends random seed vector $\mathbf{e}^i \in \mathbb{Z}_2^{Bd+1}$ as $Bd+1$ consecutive inputs to $\mathcal{F}_{ABB}$ to be stored as elements of the binary field.
3. For $j = 1, ..., Bd+1$ the players compute $r_i \leftarrow e_j^1 \oplus e_j^2$ via $\mathcal{F}_{ABB}$.

Noise sampling:

1. Each operation in the first two steps of the `FDL` specification is performed via $\mathcal{F}_{ABB}$.
2. In $\mathcal{F}_{ABB}$, the values $c_1, ..., c_B$ and r_{Bd+1} are transformed to elements in the arithmetic field.
3. All remaining operations in the `FDL` specification are performed via $\mathcal{F}_{ABB}$.

Finishing:

1. Player 1 sends $\mathbf{x} \in \mathbb{Z}_q^N$ and player 2 sends $\mathbf{y} \in \mathbb{Z}_q^N$ to $\mathcal{F}_{ABB}$ and then f is computed via $\mathcal{F}_{ABB}$ according to its specification. The result is stored as $\bar{f}$.
2. The sum of $\bar{f}$ and the `FDL` sample is computed via $\mathcal{F}_{ABB}$ and the result is output to the players.

Fig. 4: The protocol description for the `FDL` mechanism in the $\mathcal{F}_{ABB}$-hybrid world.

Theorem 1. *Let $q > 2f_\kappa^{max} + 2B_\kappa$, $B_\kappa \in \{1, \ldots, \lceil q/2 \rceil - 1\}$, $\lambda_\kappa = \frac{\varepsilon_\kappa}{\Delta f_\kappa}$ and let $e^{-B_\kappa/\lambda_\kappa}$ and $B_\kappa 2^{-d_\kappa}$ be negligible in κ. Let $\{f_\kappa : \mathbb{Z}_q^{2N} \to \mathbb{Z}_q\}_{\kappa\in\mathbb{N}}$ be an ensemble of efficiently computable deterministic functions with $f_\kappa^{max} := \max_{D\in\mathbb{Z}_q^{2N}} |f_\kappa(D)|$. Let $\{\hat{f}_\kappa(D)\}_{\kappa\in\mathbb{N}}$ be $\{\mathcal{M}_{RTGeo}^{q,f_\kappa,\lambda_\kappa}(D)\}_{\kappa\in\mathbb{N}}$.*

Then $\pi_{\mathcal{M}_{\texttt{FDL}}}(B_\kappa, d_\kappa, q, N, decomp(\lambda_\kappa, d_\kappa), f_\kappa)$ is an $(\varepsilon_\kappa, 0)$-SIM-CDP protocol for the ideal functionality $\mathcal{F}_{SFE}^{\hat{f}_\kappa}$, with respect to the same adjacency notion as in the calculation of Δf_κ, in the $\mathcal{F}_{ABB}$-hybrid world.*

Asymptotic computational cost. We consider the computational cost of $\pi_{\mathcal{M}_{\texttt{FDL}}}$ in terms of calls to the ABB, ignoring the cost of computing f. This rough model for calculating computation cost is reasonable in two ways: Firstly, local operations are canonically negligible in terms of computation cost compared to operations that require interaction. Secondly, in practice, the instantiation of the ABB greatly influences the computation cost in practical terms. As is shown

in EIKN [31], the asymptotic computational complexity of the FDL function is $O(Bd)$. This complexity follows directly from Definition 14 since all steps of the FDL procedure are repeated B times (that is, B Bernoulli trials are sampled and there are B elements in the sum) and within the Bernoulli trial subprocedure, all steps consist of d arithmetic operations.

It is important to note that the cost of sampling the noise is independent of the data query f. Relative DP usefulness intuitively increases as the number of elements in the input dataset grows. However, the performance of the sampling protocol scales with the number of queries and not with the size of the input dataset, thus amortizing its execution time further.

6.1 Application: Integer Inner-products with Bounded Elements

We now compute integer inner-products using the $\pi_{\mathcal{M}_{FDL}}$ protocol. This query type is particularly interesting for a few reasons. First, it is non-linear and cannot be expressed as an aggregate function without knowledge of the other party's inputs. Second, it is a fundamental building block for more complicated queries like matrix multiplications with vast applications in data processing, such as machine learning. To use $\pi_{\mathcal{M}_{FDL}}$, the query needs a bounded maximal absolute value, and for accuracy, we want the sensitivity of the query to be small. Therefore, we consider only inner-products where the input vectors have elements between $a \in \mathbb{Z}_q$ and $b \in \mathbb{Z}_q$. We assume that the difference between a and b is a power of 2, to facilitate inserting an input as a sequence of bits.

We consider DP with the bounded (*'change-one'*) adjacency notion and the data universe is $([a, b])^*$, such that each input D to f (as well as the protocol and the mechanism) is a tuple of $2N$ elements from $[a, b]$. Let $D := \mathbf{x}||\mathbf{y}$. The inner-product $f(D)$ is defined as $\langle \mathbf{x}, \mathbf{y} \rangle := \sum_{i=1}^{N} x_i y_i$ with operations over $\mathbb{Z}_q$. The sensitivity Δf of the inner-product is $\max(|a^2 - ab|, |b^2 - ab|)$, under the assumption that $|f(\mathbf{x}, \mathbf{y})|$ is smaller than $\lfloor q/2 \rfloor$ such that field operations mimic integer behavior. We also have that $f^{max} = N \cdot \max(a^2, b^2)$.

Parameter choices. From the properties above, the following parameter considerations follow: Let the security parameter be the bit-length of a field element, i.e. $\kappa = \lceil \log_2(q) \rceil$, as is canonical. Let both ε_κ and Δf (by choice of a, b) be independent of κ. Further, we can set the FDL specific parameters as $B = d = \kappa$. Finally, we have $q > 2f^{max} + 2B = 2N \cdot \max(a^2, b^2) + 2B$, where the inequality holds for sufficiently large κ.

In practice, one strategy is to choose κ as a canonical value for statistical security in cryptography, e.g., $\kappa = 40$, and then let this also be B and d.[15] The practical choice of ε is highly challenging, and there is a lively discussion in the

[15] Note that we use κ as a *computational* security parameter but that statistical security implies computational security. One appealing alternative is to introduce an

literature on it, although consensus is largely lacking [28, 57, 64, 56]. Luckily, there is no direct dependence on the choice of ε in the other parameters. Finally, this leaves the choices of a, b, and N. Here, we care about the distance $|a-b|$ and the size of N. Both parameters allow for wider usage scenarios when increased. However, increasing N has adverse effects on runtime, and a larger distance causes a higher sensitivity and decreased usefulness (if ε is kept fixed). Finally, there is a trade-off between N and the sizes of a, b due to their dependence on q. In practice, this can be circumvented by increasing the modulus size q in the underlying MPC instantiation.

7 Implementation and Practical Performance

We tested our protocol by implementing it in the multi-protocol SPDZ (MP-SPDZ) [53] library. Among others, it provides efficient implementations of the SPDZ_{2^k} [21] and the FKOS [36] MPC schemes, and da-bit [67] and eda-bit [35] implementations. We implement procedure `Ber` in the FKOS scheme and procedure `FDL` in the mixed-circuit setting with FKOS and SPDZ_{2^k}. We find that only one switch between computation domains is necessary, making mixed-circuit computation highly competitive in performance. More precisely, this approach is faster than previous instantiations if the conversion cost is lower than the additional overhead in the unfit computation domain. Given the protocol in EIKN [31], circuit conversion has to be faster than the overhead of computing the Bernoulli and prefix-or functionality in the arithmetic domain.

In MPC schemes, communication is typically the bottleneck of efficient function evaluation. While some communication is necessary during the computation, much of the data transfer happens in a pre-processing phase. In our setup, we have three main components that require expensive pre-processing: shared randomness for inputs, authenticated multiplication triples, and doubly authenticated bits. In our inner-product use case, we only generate one `FDL` sample. However, most pre-processing operations come in blocks of size B or d. In our implementation, we take special care to minimize the communication rounds and adapt the pre-processing batch sizes to accommodate our protocol execution.

Our setting provides security in the presence of active adversaries. Since these parties can deviate arbitrarily from the protocol, they might send input out of range. It is, therefore, necessary to prove the correctness of the input domain in both the FDL mechanism and the query function. There are different strategies to achieve such a feat. We note that the ABB accepts inputs of two types, either elements in the binary field or the larger finite field. We need to restrict the values to the pre-defined range for inputs in the arithmetic domain. Were we not to perform such an input validation, this would result in an increased sensitivity of the function (in relationship to what is a priori agreed upon by the two parties),

additional statistical parameter separate from κ, let them be proportional and align B, d to the new parameter instead.

thwarting the privacy level of the DP mechanism. In the presence of passive adversaries, however, there is of course no need to validate the inputs since the adversary will per definition not give out-of-range inputs. This requirement of a *proof of function sensitivity* also arises in other scenarios where the sensitivity is directly dependent on the secret data of multiple parties.

To provide such a range-proof of the inputs of each party, we consider two main options: Firstly, one could accept the inputs as elements in the larger field and then perform a zero-knowledge range proof[16] within the MPC domain, and secondly, one could accept the inputs bit-by-bit and re-compose those bits into elements of the larger field. These approaches present a trade-off in input and proof complexity. In the first approach, the cost of inputting a value is constant (i.e., depending on 2^k in our example) while proving the range is linear in the bound. In the bit-by-bit setting, the input and proving cost are logarithmic in the bound. The second approach is thus more efficient for larger bound values depending on the specific scheme.

7.1 Benchmarks

In this section, we present benchmarks of our `FDL` mechanism with $B = d = \kappa$ and measure performance for different settings[17]. Relevant for parameter $\hat{\alpha}$, the bit decomposition of the Bernoulli bias, is the decomposition length d. When setting a value α, the binary decomposition truncates this value to the predefined precision. Although our code can be instantiated with any number of parties, we fixed the number of parties to 2 as to align with the formalities of earlier sections. We provide exemplary data points at 40- and 80-bit, typical statistical security parameters. Next, we evaluate the mechanism at 128-bit, a usual conservative choice as a computational security parameter. Note that, in MP-SPDZ, the underlying security parameters for SPDZ_{2^k} are fixed to 64-bit computational and 64-bit statistical security. We run all benchmarks on a Linux server with an AMD Ryzen 9 7900X CPU (4.7 GHz). Each party only has access to one thread for computations. We separate our results into single sample computation and amortized evaluation for 1000 samples. The single sample evaluation is further split into the pre-processing and online phases of MPC, where the pre-processing step consists of generating necessary multiplication triples and da-bits.

Table 1 presents the runtime metrics for different network settings. In Setting 1, we have an unrestricted LAN setup. Setting 2 simulates a less powerful LAN setup by limiting the network to 1Gbit/s and the round-trip time (RTT) to 1ms. Finally, in Setting 3, we simulate a WAN network with 100Mbit/s and 100ms RTT, reflecting a solid but distant connection (e.g., intercontinental). Given the asymptotic complexity $O(Bd)$, the runtime results reflect the expected quadratic

[16] For instance, such as described in the Bulletproofs paper [9].

[17] The anonymized code can be found at https://anonymous.4open.science/r/laplace_sampler-57CE/.

Table 1: Runtime in milliseconds of benchmarks with different security levels. Total runtime is for a single sample, while amortized runtime assumes 1000 samples.

Protocol	κ	Prep.	Online	Total	Amort.
10 Gbit/s with RTT of 1 ms					
Ours	40	74.7	42	116.6	34.6
	80	94.2	119.9	214.1	118.5
	128	130	276.9	406.9	283.4
[52]	40	1606	37.72	1 643	992†
1 Gbit/s with RTT of 1 ms					
Ours	40	182.9	248.4	431.2	69.7
	80	245.6	650.2	895.7	209.7
	128	345.6	1 362	1 707	520.3
[52]	40	4 707	4.81	–	4 711‡
100 Mbit/s with RTT of 100 ms					
Ours	40	11 256	20 486	31 742	577.9
	80	15 215	51 794	67 009	1 604
	128	20 795	105 350	126 145	3 558
[52]	40	42 352	47.99	–	42 400‡

† Amortized over 40 samples
‡ Amortized over 10 samples, no single sample performance provided.

growth in the security parameter. Regarding the network settings, communication is needed for inputs, binary AND gates, arithmetic multiplication, secret share conversion, and outputs. Since inputs, conversions, and computations depend on one or both parameters B, or d, the negative impact of a reduced network speed and increased RTT is increased. Compared to concurrent work [52], our mechanism outperforms theirs in runtime and memory for the overall computation in the fast network settings.[18] Arguably, their setup heavily optimizes the online phase, making it more efficient if pre-processing can be performed in advance. However, sampling geometric noise in MPC can generally be seen as pre-processing since the sensitivity of a function is known before the data is processed, and the parties can already engage in the noise sampling procedure before their inputs to the query function have been fixed. Further, their geometric mechanism has a low round complexity, showing improved performance in WAN network settings. Comparing with [31] is challenging as only asymptotic complexities are given there and the results are based on arithmetic evaluations of binary computations from [66]. Our approach, on the other hand,

[18] One should further note that [52] is in the more efficient setting of passive adversaries, thus making direct comparisons skewed in their favor.

Table 2: Network cost in MB of different geometric sampling settings. The amortized cost assumes 1000 samples.

Protocol	κ	Prep.	Online	Total	Amort.
Ours	40	14.7	17.9	65.3	23.8
	80	20.9	58.3	158.3	75.3
	128	29.2	143.4	345.2	173.6
[52]	40	–	–	492.7†	–

† Run with single sample, no amortized network cost provided.

is based on mixed circuits [67] and includes substantial performance improvements by dedicated parameter optimizations. In our benchmarks, we adhered to the following principle. We aimed to reduce the communication complexity for low-latency networks, while for the high-latency networks, we reduced the round complexity. This trade-off can be determined with the pre-processing batch size parameter. Given a high minimum batch size for the MPC schemes we use, computing a single geometric sample leads to substantial overhead. Thus, it is crucial to parametrize the implementation according to the number of samples and expected network latency.

In Table 2, we present benchmarks for network costs for each security parameter. We see that the network cost of our implementation is lower than that in [52], further showing that their round complexity is much lower than that of the malicious secure SPDZ_{2^k} protocol. Given the network cost, we could further reduce the network bandwidth before its limiting impact equals a slow RTT. In our amortized costs column, we present the network traffic per sample in a computation of 1 000 samples.

8 Conclusion and Outlooks

In this work, we revisit the idea of generic emulation of the central dataholder (GED) as a method to achieve accuracy equal to that of the central model of DP without the need for a single trusted dataholder. The bulk of our work is spent analysing existing definitions of computational DP (CDP) in the multiparty setting, noting that whereas they are very well-suited for theoretic study and use with special-purpose MPC schemes, they all fit somewhat suboptimally to the task of GED. Since one of them, SIM^{+}-CDP, appears to fit very well conceptually but has some details preventing its use together with canonical statistical DP (SDP) mechanisms, we propose both a generalised version of it, SIM°-CDP, and another instantiation of that generalised definition, SIM^{*}-CDP, that we argue is more fitting to the current state-of-the-art in both general-purpose MPC and SDP. We relate SIM^{*}-CDP to IND-CDP and SIM-CDP, showing that it is a stronger notion in the sense that all SIM^{*}-CDP protocols are also SIM-

CDP (and thus also IND-CDP) with unchanged parameters, whilst there are computational tasks (such as computing a given functionality to within a given absolute error with constant probability) that can be solved with SIM-CDP but are impossible with SIM*-CDP. Further, we show that SIM$^+$-CDP and SIM*-CDP are separated from each other in both directions in the sense that there are tasks solvable for one of them but not the other. Some of these results, however, are established under specific parameter regimes (as is commonplace in the CDP literature) and therefore extending them to wider regimes is an interesting open problem. On the practical side, we show how to achieve SIM*-CDP via the truncated geometric (discrete Laplace) mechanism by using a state-of-the-art protocol for distributed noise sampling and analyse the use of this protocol for computing integer inner-products with active security in the two-party setting. We then provide an open-sourced implementation of the protocol using the MP-SPDZ library and show that it is very efficient in practice.

As always when formulating new definitions in cryptography questions arise, such as whether the definition is intuitive, practically usable, and not overly relaxed or strict. On the usability front, we present evidence that SIM*-CDP is practical since it allows us to design efficient, quite general protocols of natural tasks that fulfill it, and the proof that the definition is satisfied follows essentially directly from the use of general-purpose MPC and an SDP mechanism. Further, the definition appears intuitive due to its closeness to both previous definitions and established formalities in both the DP and MPC domains. There is, however, much need for additional scrutiny, and this is the case also for the question about balance in the definition. Interesting open directions here are to more tightly relate the definition to previous ones and explore whether there is some characteristic trait of SDP that is captured in the previous ones but not in SIM*-CDP.

Acknowledgements. We gratefully thank the authors of MPRV [65] for providing a full version of their paper and for answering our questions. We thank Lea Demelius and Peter Waldert for fruitful discussions on earlier drafts of this paper. We thank anonymous PETS reviewers for helping us vastly improve the disposition and clarity of the work.

This project was supported by the "DDAI" COMET Module within the COMET – Competence Centers for Excellent Technologies Programme, funded by the Austrian Federal Ministry (BMK and BMDW), the Austrian Research Promotion Agency (FFG), the province of Styria (SFG) and partners from industry and academia. The COMET Programme is managed by FFG.

This project was also in part funded by the CONFIDENTIAL-6G EU project (Grant No: 101096435).

References

1. Anandan, B., Clifton, C.: Laplace noise generation for two-party computational differential privacy. In: 2015 13th Annual Conference on Privacy, Security and Trust (PST). pp. 54–61 (2015). https://doi.org/10.1109/PST.2015.7232954
2. Balcer, V., Vadhan, S.: Differential privacy on finite computers. Journal of Privacy and Confidentiality **9**(2) (Sep 2019). https://doi.org/10.29012/jpc.679, https://journalprivacyconfidentiality.org/index.php/jpc/article/view/679
3. Beaver, D.: Correlated pseudorandomness and the complexity of private computations. In: Proceedings of the Twenty-Eighth Annual ACM Symposium on Theory of Computing. p. 479–488. STOC '96, Association for Computing Machinery, New York, NY, USA (1996). https://doi.org/10.1145/237814.237996, https://doi.org/10.1145/237814.237996
4. Beimel, A., Nissim, K., Omri, E.: Distributed private data analysis: Simultaneously solving how and what. In: Wagner, D. (ed.) Advances in Cryptology – CRYPTO 2008. pp. 451–468. Springer Berlin Heidelberg, Berlin, Heidelberg (2008)
5. Bell, J., Gascón, A., Ghazi, B., Kumar, R., Manurangsi, P., Raykova, M., Schoppmann, P.: Distributed, private, sparse histograms in the two-server model. p. 307–321. CCS '22, Association for Computing Machinery, New York, NY, USA (2022). https://doi.org/10.1145/3548606.3559383, https://doi.org/10.1145/3548606.3559383
6. Bittau, A., Erlingsson, Ú., Maniatis, P., Mironov, I., Raghunathan, A., Lie, D., Rudominer, M., Kode, U., Tinnes, J., Seefeld, B.: Prochlo. In: Proceedings of the 26th Symposium on Operating Systems Principles. ACM (oct 2017). https://doi.org/10.1145/3132747.3132769, https://doi.org/10.1145%2F3132747.3132769
7. Bun, M., Chen, Y.H., Vadhan, S.: Separating computational and statistical differential privacy in the client-server model. In: Hirt, M., Smith, A. (eds.) Theory of Cryptography. pp. 607–634. Springer Berlin Heidelberg, Berlin, Heidelberg (2016)
8. Böhler, J.: Input Secrecy & Output Privacy: Efficient Secure Computation of Differential Privacy Mechanisms. Ph.D. thesis, Karlsruher Institut für Technologie (KIT) (2021). https://doi.org/10.5445/IR/1000141098
9. Bünz, B., Bootle, J., Boneh, D., Poelstra, A., Wuille, P., Maxwell, G.: Bulletproofs: Short proofs for confidential transactions and more. In: 2018 IEEE Symposium on Security and Privacy (SP). pp. 315–334 (2018). https://doi.org/10.1109/SP.2018.00020
10. Canetti, R.: Security and composition of multiparty cryptographic protocols. J. Cryptol. **13**(1), 143–202 (jan 2000). https://doi.org/10.1007/s001459910006, https://doi.org/10.1007/s001459910006
11. Canetti, R.: Universally composable security: A new paradigm for cryptographic protocols. Cryptology ePrint Archive, Paper 2000/067 (2000), https://eprint.iacr.org/2000/067, https://eprint.iacr.org/2000/067
12. Canetti, R., Kushilevitz, E., Lindell, Y.: On the limitations of universally composable two-party computation without set-up assumptions. In: Proceedings of the 22nd International Conference on Theory and Applications of Cryptographic Techniques. p. 68–86. EUROCRYPT'03, Springer-Verlag, Berlin, Heidelberg (2003)
13. Canetti, R., Lindell, Y., Ostrovsky, R., Sahai, A.: Universally composable two-party and multi-party secure computation. In: Proceedings of the Thiry-Fourth Annual ACM Symposium on Theory of Computing. p. 494–503. STOC '02, Association for Computing Machinery, New York, NY, USA (2002). https://doi.org/10.1145/509907.509980, https://doi.org/10.1145/509907.509980

14. Canonne, C., Kamath, G., Steinke, T.: The discrete gaussian for differential privacy. Journal of Privacy and Confidentiality **12**(1) (Jul 2022). https://doi.org/10.29012/jpc.784
15. Champion, J., shelat, a., Ullman, J.: Securely sampling biased coins with applications to differential privacy. In: Proceedings of the 2019 ACM SIGSAC Conference on Computer and Communications Security. p. 603–614. CCS '19, Association for Computing Machinery, New York, NY, USA (2019). https://doi.org/10.1145/3319535.3354256, https://doi.org/10.1145/3319535.3354256
16. Chan, T.H.H., Shi, E., Song, D.: Optimal lower bound for differentially private multi-party aggregation. In: Epstein, L., Ferragina, P. (eds.) Algorithms – ESA 2012. pp. 277–288. Springer Berlin Heidelberg, Berlin, Heidelberg (2012)
17. Cheu, A., Smith, A., Ullman, J., Zeber, D., Zhilyaev, M.: Distributed differential privacy via shuffling. In: Ishai, Y., Rijmen, V. (eds.) Advances in Cryptology – EUROCRYPT 2019. pp. 375–403. Springer International Publishing, Cham (2019)
18. Cheu, A., Yan, C.: Necessary Conditions in Multi-Server Differential Privacy. In: Tauman Kalai, Y. (ed.) 14th Innovations in Theoretical Computer Science Conference (ITCS 2023). Leibniz International Proceedings in Informatics (LIPIcs), vol. 251, pp. 36:1–36:21. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl, Germany (2023). https://doi.org/10.4230/LIPIcs.ITCS.2023.36, https://drops.dagstuhl.de/entities/document/10.4230/LIPIcs.ITCS.2023.36
19. Chor, B., Kushilevitz, E.: A zero-one law for boolean privacy. In: Proceedings of the Twenty-First Annual ACM Symposium on Theory of Computing. p. 62–72. STOC '89, Association for Computing Machinery, New York, NY, USA (1989). https://doi.org/10.1145/73007.73013, https://doi.org/10.1145/73007.73013
20. Corrigan-Gibbs, H., Boneh, D.: Prio: Private, robust, and scalable computation of aggregate statistics. In: 14th USENIX Symposium on Networked Systems Design and Implementation (NSDI 17). pp. 259–282. USENIX Association, Boston, MA (2017), https://www.usenix.org/conference/nsdi17/technical-sessions/presentation/corrigan-gibbs
21. Cramer, R., Damgård, I., Escudero, D., Scholl, P., Xing, C.: Spd$\mathbb{F}_{2^k}$: Efficient MPC mod 2^k for dishonest majority. In: Shacham, H., Boldyreva, A. (eds.) Advances in Cryptology - CRYPTO 2018 - 38th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 19-23, 2018, Proceedings, Part II. Lecture Notes in Computer Science, vol. 10992, pp. 769–798. Springer (2018). https://doi.org/10.1007/978-3-319-96881-0_26, https://doi.org/10.1007/978-3-319-96881-0_26
22. Cramer, R., Damgård, I.B., Nielsen, J.B.: Secure Multiparty Computation and Secret Sharing. Cambridge University Press (2015). https://doi.org/10.1017/CBO9781107337756
23. Damgård, I., Nielsen, J.B.: Universally composable efficient multiparty computation from threshold homomorphic encryption. In: Boneh, D. (ed.) Advances in Cryptology - CRYPTO 2003. pp. 247–264. Springer Berlin Heidelberg, Berlin, Heidelberg (2003)
24. Damgård, I., Pastro, V., Smart, N., Zakarias, S.: Multiparty computation from somewhat homomorphic encryption. In: Safavi-Naini, R., Canetti, R. (eds.) Advances in Cryptology – CRYPTO 2012. pp. 643–662. Springer Berlin Heidelberg, Berlin, Heidelberg (2012)
25. Demmler, D., Schneider, T., Zohner, M.: ABY - A framework for efficient mixed-protocol secure two-party computation. In: 22nd Annual

Network and Distributed System Security Symposium, NDSS 2015, San Diego, California, USA, February 8-11, 2015. The Internet Society (2015), https://www.ndss-symposium.org/ndss2015/aby---framework-efficient-mixed-protocol-secure-two-party-computation
26. Dwork, C.: Differential privacy. In: Proceedings of the 33rd International Conference on Automata, Languages and Programming - Volume Part II. p. 1–12. ICALP'06, Springer-Verlag, Berlin, Heidelberg (2006). https://doi.org/10.1007/11787006_1
27. Dwork, C., Kenthapadi, K., McSherry, F., Mironov, I., Naor, M.: Our data, ourselves: Privacy via distributed noise generation. In: Proceedings of the 24th Annual International Conference on The Theory and Applications of Cryptographic Techniques. p. 486–503. EUROCRYPT'06, Springer-Verlag, Berlin, Heidelberg (2006). https://doi.org/10.1007/11761679_29, https://doi.org/10.1007/11761679_29
28. Dwork, C., Kohli, N., Mulligan, D.: Differential privacy in practice: Expose your epsilons! Journal of Privacy and Confidentiality **9** (10 2019). https://doi.org/10.29012/jpc.689
29. Dwork, C., McSherry, F., Nissim, K., Smith, A.: Calibrating noise to sensitivity in private data analysis. vol. Vol. 3876, pp. 265–284 (01 2006). https://doi.org/10.1007/11681878_14
30. Dwork, C., Roth, A.: The algorithmic foundations of differential privacy. Foundations and Trends® in Theoretical Computer Science **9**(3–4), 211–407 (2014)
31. Eriguchi, R., Ichikawa, A., Kunihiro, N., Nuida, K.: Efficient noise generation to achieve differential privacy with applications to secure multiparty computation. In: Financial Cryptography and Data Security: 25th International Conference, FC 2021, Virtual Event, March 1–5, 2021, Revised Selected Papers, Part I. p. 271–290. Springer-Verlag, Berlin, Heidelberg (2021). https://doi.org/10.1007/978-3-662-64322-8_13, https://doi.org/10.1007/978-3-662-64322-8_13
32. Eriguchi, R., Ichikawa, A., Kunihiro, N., Nuida, K.: Efficient noise generation protocols for differentially private multiparty computation. IEEE Transactions on Dependable and Secure Computing **20**(6), 4486–4501 (2023). https://doi.org/10.1109/TDSC.2022.3227568
33. Úlfar Erlingsson, Feldman, V., Mironov, I., Raghunathan, A., Talwar, K., Thakurta, A.: Amplification by Shuffling: From Local to Central Differential Privacy via Anonymity (2019)
34. Escudero, D.: An introduction to secret-sharing-based secure multiparty computation. Cryptology ePrint Archive, Paper 2022/062 (2022), https://eprint.iacr.org/2022/062, https://eprint.iacr.org/2022/062
35. Escudero, D., Ghosh, S., Keller, M., Rachuri, R., Scholl, P.: Improved primitives for mpc over mixed arithmetic-binary circuits. In: Micciancio, D., Ristenpart, T. (eds.) Advances in Cryptology – CRYPTO 2020. pp. 823–852. Springer International Publishing, Cham (2020)
36. Frederiksen, T.K., Keller, M., Orsini, E., Scholl, P.: A unified approach to mpc with preprocessing using ot. In: Proceedings, Part I, of the 21st International Conference on Advances in Cryptology – ASIACRYPT 2015 - Volume 9452. p. 711–735. Springer-Verlag, Berlin, Heidelberg (2015). https://doi.org/10.1007/978-3-662-48797-6_29, https://doi.org/10.1007/978-3-662-48797-6_29
37. Ghazi, B., Ilango, R., Kamath, P., Kumar, R., Manurangsi, P.: Towards separating computational and statistical differential privacy. In: 2023 IEEE 64th Annual Symposium on Foundations of Computer Science (FOCS). pp. 580–599 (2023). https://doi.org/10.1109/FOCS57990.2023.00042

38. Ghosh, A., Roughgarden, T., Sundararajan, M.: Universally utility-maximizing privacy mechanisms. SIAM Journal on Computing **41**(6), 1673–1693 (2012). https://doi.org/10.1137/09076828X, https://doi.org/10.1137/09076828X
39. Goldreich, O.: Foundations of Cryptography: Volume 2, Basic Applications. Cambridge University Press, USA (2004)
40. Goldreich, O., Goldwasser, S., Nussboim, A.: On the implementation of huge random objects. SIAM Journal on Computing **39**(7), 2761–2822 (2010). https://doi.org/10.1137/080722771, https://doi.org/10.1137/080722771
41. Goyal, V., Khurana, D., Mironov, I., Pandey, O., Sahai, A.: Do Distributed Differentially-Private Protocols Require Oblivious Transferl. In: Chatzigiannakis, I., Mitzenmacher, M., Rabani, Y., Sangiorgi, D. (eds.) 43rd International Colloquium on Automata, Languages, and Programming (ICALP 2016). Leibniz International Proceedings in Informatics (LIPIcs), vol. 55, pp. 29:1–29:15. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, Dagstuhl, Germany (2016). https://doi.org/10.4230/LIPIcs.ICALP.2016.29, http://drops.dagstuhl.de/opus/volltexte/2016/6308
42. Goyal, V., Mironov, I., Pandey, O., Sahai, A.: Accuracy-privacy tradeoffs for two-party differentially private protocols. In: Canetti, R., Garay, J.A. (eds.) Advances in Cryptology – CRYPTO 2013. pp. 298–315. Springer Berlin Heidelberg, Berlin, Heidelberg (2013)
43. Groce, A., Katz, J., Yerukhimovich, A.: Limits of computational differential privacy in the client/server setting. In: Ishai, Y. (ed.) Theory of Cryptography. pp. 417–431. Springer Berlin Heidelberg, Berlin, Heidelberg (2011)
44. Groce, A., Rindal, P., Rosulek, M.: Cheaper private set intersection via differentially private leakage. Proceedings on Privacy Enhancing Technologies **2019**, 6–25 (07 2019). https://doi.org/10.2478/popets-2019-0034
45. Haitner, I., Mazor, N., Shaltiel, R., Silbak, J.: Channels of small log-ratio leakage and characterization of two-party differentially private computation. In: Hofheinz, D., Rosen, A. (eds.) Theory of Cryptography. pp. 531–560. Springer International Publishing, Cham (2019)
46. Haitner, I., Mazor, N., Silbak, J., Tsfadia, E.: On the complexity of two-party differential privacy. In: Proceedings of the 54th Annual ACM SIGACT Symposium on Theory of Computing. p. 1392–1405. STOC 2022, Association for Computing Machinery, New York, NY, USA (2022). https://doi.org/10.1145/3519935.3519982, https://doi.org/10.1145/3519935.3519982
47. He, X., Machanavajjhala, A., Flynn, C., Srivastava, D.: Composing differential privacy and secure computation: A case study on scaling private record linkage. In: Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security. p. 1389–1406. CCS '17, Association for Computing Machinery, New York, NY, USA (2017). https://doi.org/10.1145/3133956.3134030, https://doi.org/10.1145/3133956.3134030
48. Hirt, M., Maurer, U., Zikas, V.: Mpc vs. sfe: Unconditional and computational security. In: Proceedings of the 14th International Conference on the Theory and Application of Cryptology and Information Security: Advances in Cryptology. p. 1–18. ASIACRYPT '08, Springer-Verlag, Berlin, Heidelberg (2008). https://doi.org/10.1007/978-3-540-89255-7_1, https://doi.org/10.1007/978-3-540-89255-7_1
49. Ishai, Y., Kushilevitz, E., Lindell, Y., Petrank, E.: On combining privacy with guaranteed output delivery in secure multiparty computation. In: Proceedings of the 26th Annual International Conference on Advances in Cryptology. p. 483–500.

CRYPTO'06, Springer-Verlag, Berlin, Heidelberg (2006). https://doi.org/10.1007/11818175_29, https://doi.org/10.1007/11818175_29
50. Kasiviswanathan, S.P., Lee, H.K., Nissim, K., Raskhodnikova, S., Smith, A.: What can we learn privately? SIAM Journal on Computing **40**(3), 793–826 (2011). https://doi.org/10.1137/090756090, https://doi.org/10.1137/090756090
51. Keeler, D., Komlo, C., Lepert, E., Veitch, S., He, X.: Dprio: Efficient differential privacy with high utility for prio. Proceedings on Privacy Enhancing Technologies (2023)
52. Keller, H., Möllering, H., Schneider, T., Tkachenko, O., Zhao, L.: Secure noise sampling for dp in mpc with finite precision. Cryptology ePrint Archive, Paper 2023/1594 (2023), https://eprint.iacr.org/2023/1594, https://eprint.iacr.org/2023/1594
53. Keller, M.: Mp-spdz: A versatile framework for multi-party computation. In: Proceedings of the 2020 ACM SIGSAC Conference on Computer and Communications Security. p. 1575–1590. CCS '20, Association for Computing Machinery, New York, NY, USA (2020). https://doi.org/10.1145/3372297.3417872, https://doi.org/10.1145/3372297.3417872
54. Keller, M., Orsini, E., Scholl, P.: Mascot: Faster malicious arithmetic secure computation with oblivious transfer. In: Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security. p. 830–842. CCS '16, Association for Computing Machinery, New York, NY, USA (2016). https://doi.org/10.1145/2976749.2978357, https://doi.org/10.1145/2976749.2978357
55. Khurana, D., Maji, H.K., Sahai, A.: Black-box separations for differentially private protocols. In: Sarkar, P., Iwata, T. (eds.) Advances in Cryptology – ASIACRYPT 2014. pp. 386–405. Springer Berlin Heidelberg, Berlin, Heidelberg (2014)
56. Kifer, D., Abowd, J.M., Ashmead, R., Cumings-Menon, R., Leclerc, P., Machanavajjhala, A., Sexton, W., Zhuravlev, P.: Bayesian and frequentist semantics for common variations of differential privacy: Applications to the 2020 census (2022)
57. Krehbiel, S.: Choosing epsilon for privacy as a service. Proceedings on Privacy Enhancing Technologies **2019**, 192 – 205 (2019)
58. Lindell, Y.: How to Simulate It – A Tutorial on the Simulation Proof Technique, pp. 277–346. Springer International Publishing, Cham (2017)
59. Lindell, Y.: Secure multiparty computation. Commun. ACM **64**(1), 86–96 (dec 2021). https://doi.org/10.1145/3387108, https://doi.org/10.1145/3387108
60. Lipmaa, H., Toft, T.: Secure equality and greater-than tests with sublinear online complexity. In: Fomin, F.V., Freivalds, R., Kwiatkowska, M., Peleg, D. (eds.) Automata, Languages, and Programming. pp. 645–656. Springer Berlin Heidelberg, Berlin, Heidelberg (2013)
61. Mazloom, S., Gordon, S.D.: Secure computation with differentially private access patterns. In: Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security. p. 490–507. CCS '18, Association for Computing Machinery, New York, NY, USA (2018). https://doi.org/10.1145/3243734.3243851, https://doi.org/10.1145/3243734.3243851
62. McGregor, A., Mironov, I., Pitassi, T., Reingold, O., Talwar, K., Vadhan, S.: The limits of two-party differential privacy. In: 2010 IEEE 51st Annual Symposium on Foundations of Computer Science. pp. 81–90 (2010). https://doi.org/10.1109/FOCS.2010.14
63. Mcgregor, A., Mironov, I., Pitassi, T., Reingold, O., Talwar, K., Vadhan, S.P.: The limits of two-party differential privacy. Electronic Colloquium on Computational Complexity **18**, 106 (2011)

64. Mehner, L., von Voigt, S.N., Tschorsch, F.: Towards explaining epsilon: A worst-case study of differential privacy risks. 2021 IEEE European Symposium on Security and Privacy Workshops (EuroS&PW) pp. 328–331 (2021)
65. Mironov, I., Pandey, O., Reingold, O., Vadhan, S.: Computational differential privacy. In: Halevi, S. (ed.) Advances in Cryptology - CRYPTO 2009. pp. 126–142. Springer Berlin Heidelberg, Berlin, Heidelberg (2009)
66. Nishide, T., Ohta, K.: Multiparty computation for interval, equality, and comparison without bit-decomposition protocol. In: Okamoto, T., Wang, X. (eds.) Public Key Cryptography - PKC 2007, 10th International Conference on Practice and Theory in Public-Key Cryptography, Beijing, China, April 16-20, 2007, Proceedings. Lecture Notes in Computer Science, vol. 4450, pp. 343–360. Springer (2007). https://doi.org/10.1007/978-3-540-71677-8_23, https://doi.org/10.1007/978-3-540-71677-8_23
67. Rotaru, D., Wood, T.: Marbled circuits: Mixing arithmetic and boolean circuits with active security. In: Progress in Cryptology – INDOCRYPT 2019: 20th International Conference on Cryptology in India, Hyderabad, India, December 15–18, 2019, Proceedings. p. 227–249. Springer-Verlag, Berlin, Heidelberg (2019). https://doi.org/10.1007/978-3-030-35423-7_12, https://doi.org/10.1007/978-3-030-35423-7_12
68. Roth, E., Noble, D., Falk, B.H., Haeberlen, A.: Honeycrisp: large-scale differentially private aggregation without a trusted core. In: Proceedings of the 27th ACM Symposium on Operating Systems Principles. p. 196–210. SOSP '19, Association for Computing Machinery, New York, NY, USA (2019). https://doi.org/10.1145/3341301.3359660, https://doi.org/10.1145/3341301.3359660
69. Schoppmann, P., Vogelsang, L., Gascón, A., Balle, B.: Secure and scalable document similarity on distributed databases: Differential privacy to the rescue. Proceedings on Privacy Enhancing Technologies **2020**, 209 – 229 (2020)
70. Vadhan, S.: The Complexity of Differential Privacy, pp. 347–450. Springer International Publishing, Cham (2017). https://doi.org/10.1007/978-3-319-57048-8_7, https://doi.org/10.1007/978-3-319-57048-8_7
71. Viola, E.: The complexity of distributions. SIAM Journal on Computing **41**(1), 191–218 (2012). https://doi.org/10.1137/100814998, https://doi.org/10.1137/100814998
72. Watson, T.W.: The Computational Complexity of Randomness. Ph.D. thesis (2013), https://www.proquest.com/dissertations-theses/computational-complexity-randomness/docview/1441711479/se-2, copyright - Database copyright ProQuest LLC; ProQuest does not claim copyright in the individual underlying works; Zuletzt aktualisiert - 2023-03-03
73. Wei, C., Yu, R., Fan, Y., Chen, W., Wang, T.: Securely sampling discrete gaussian noise for multi-party differential privacy. In: Proceedings of the 2023 ACM SIGSAC Conference on Computer and Communications Security. p. 2262–2276. CCS '23, Association for Computing Machinery, New York, NY, USA (2023). https://doi.org/10.1145/3576915.3616641, https://doi.org/10.1145/3576915.3616641

A The Efficient Samplability of Distributions

We consider a probability distribution Dist efficiently samplable if there exists a PPT Turing Machine (TM) that maps 1^κ to a sample from Dist. That is, if there exists a probabilistic TM that runs in strict polynomial time and, using

only its unbiased randomness tape, samples from Dist exactly. We do not delve into the broader literature on what distributions can be sampled (exactly or approximately) under certain constraints (see, for instance [71, 40, 72]), but rather settle for some basic remarks. Most critically, we note that since the sampler runs in polynomial time, it can read at most $2^{poly(\kappa)}$ coins from its randomness tape, with $poly(\cdot)$ being some polynomial.[19] Since the randomness tape is the only source of randomness in the sampler, the sampler is deterministic if one considers the coins as input. This means that there are only $2^{poly(\kappa)}$ possible executions of the sampler, each giving a fixed output. This has two direct consequences:

- **All probability densities in Dist must be multiples of $2^{-poly(\kappa)}$.**
- **The support of Dist can contain at most $2^{poly(\kappa)}$ distinct elements.**

The first of these, of course, implies the second and the second can similarly be directly realised by that the sampler can write at most $poly(\kappa)$ elements on its output tape, since it is strict PPT. The restriction on the support is anyhow useful to include explicitly, since it implies that only discrete distributions on a sufficiently small support can be efficiently sampled. This rules out, most directly, sampling from the reals (as in the usual Laplace or Gaussian distributions) but also, a bit more subtly, sampling distributions whose support is a finite of size $q > 2^{poly(\kappa)}$. The first restriction however is the more important one, and in particular, it rules out distributions such as

- **Bernoulli trials of general parameters:** There are parameters of α such that the Bernoulli distribution $Ber(\alpha)$ can not be efficiently sampled. One example of this is $\alpha = 1/3$. In particular, $Ber(\alpha)$ is efficiently samplable iff α is a multiple of $2^{-poly(\kappa)}$ for some polymial $poly(\cdot)$.
- **Truncated geometric and discrete gaussian distributions of general parameters:** There are parameters of α such that the geometric distribution (discrete Laplace) on a finite (small) support $Geo(\alpha)$ can not be efficiently sampled. This is due to the need to generate probability densities of the form $\frac{1}{2\alpha}e^{\frac{-|z|}{\alpha}}$, which generally cannot be expressed as an inverse power of two. The same reasoning holds for the discrete gaussian.

The takeaway from these examples, since we still would like to use these kinds of distributions when constructing DP mechanisms, is that one must either relax the demand for *strict* polynomial time or the demand that the samples are from *exactly* Dist rather than from a good approximation of Dist. Indeed, is is shown in [14] that the discrete gaussian (with support on the integers) can be sampled in *expected* polynomial time. In practice, settling for expected polynomial time is arguably not at all problematic, at least in the central model. The problems arise on the theory side when trying to prove security in a protocol, since the literature on secure computation heavily favours strict polynomial time, meaning that directly slotting in a mechanism that potentially runs, say, exponentially long might prove a large obstacle to proving security of the protocol as a whole.

[19] This simple fact was made aware to us by [2, 14].

The other alternative is to settle for approximating the distribution in question, a strategy arguably more readily usable in conjunction with formally proving secure computation, as we do in this work. The challenge to this approach however is to include this sampling error into the DP guarantee in question, for instance letting it be a part of the δ parameter or including it in a computational error term in the CDP notion in question.

B The Ideal/real Paradigm, Standalone and UC security

We now give a brief introduction to the real/ideal-world paradigm of security and it's two standard versions, the *standalone security model* and the *universal composability (UC) security model*. Due to their complicated nature, we will not be able to describe them in full formal detail and we refer to [11, 22] for details on the UC model and to [10, 39, 58] for details on the standalone model. In our summary here, we lean upon those in [44, 59, 34].

The core idea of the ideal/real paradigm of security is to define an *ideal world* that is secure by definition, i.e. which formulates what computations are supposed to be done and what it means formally to have that done securely (for instance, specifying what types of information leakage are not to be seen as a violation of security). Then the security of the actual protocol in question, defining the *real world*, is asserted by a simulation proof that the adversary cannot know if it is interacting with the ideal world or the real world. That is, the intuition is that if the adversary cannot tell if it is interacting with the real protocol or a version of the protocol that is secure by definition, then the protocol should be seen as essentially as secure as that in the ideal world.

The ideal world works as follows: There is an incorruptible third party called the *ideal functionality* which is given the inputs of all of the parties. The functionality then performs the computation in question, perhaps incorporating some well-defined allowed influence of the adversary, and then forwards the result to the players who output it. The functionality thereby defines what it means to be secure, and what computations should be possible, by merit of being incorruptible. However, when observing the protocol execution, it is potentially very easy to distinguish such an ideal world from the real world, for instance by observing the number of messages sent. Therefore, the ideal world also must include a *simulator*, or *ideal-world adversary*, whose task it is to generate a view that is indistinguishable from that of the real-world adversary and to do this by using only the information available to it in the ideal world (essentially, the information given to it by the ideal functionality). An important quantifier of the strength of the simulation argument is the efficiency of the simulator, since it describes how much work is needed to turn the allowed information leakage into the real one, with an efficient simulator giving a stronger guarantee of security. Therefore, in the literature on secure computation, one typically requires the simulator to be efficient, although this is not always the case for CDP using the ideal/real-world

paradigm.[20]

So the core idea is that the ideal world (with parties, ideal functionality and simulator) in some sense looks like the real world (with parties and adversary). This begs the question of who they should look the same to – who is the distinguisher? Here is where the standalone and UC security models start diverging. In the standalone model, the distinguisher is the adversary, meaning that the distinguisher itself takes part in the protocol. More precisely, the task of the simulator is to use only information available in the ideal world and generate an output distribution that is indistinguishable from the view[21] of the real-world adversary.

Definition 16 (Standalone security, reformulation of Def. 4 in [10]). *We say that a protocol π is a secure protocol for the functionality $\mathcal{F}$ if for all efficient adversaries $\mathcal{A}$, there exists an efficient simulator $\mathcal{S}$ (corrupting the same parties as $\mathcal{A}$) such that the joint output of the honest parties and $\mathcal{A}$ in the real world is computationally indistinguishable from the joint output of the honest parties and $\mathcal{S}$ in the ideal world, i.e. when the output distributions in the ideal and real worlds are computationally indistinguishable.*

There are many different flavors of the security definition, for instance in the type of indistinguishability (sometimes one requires the distributions to be identical or have negligible statistical distance) or in the inclusions of specific requirements on the *correctness* (such as requiring that the outputs in the real and ideal worlds are identical or statistically close if there are no corruptions, as done in [39, 58]). The version of the notion which is used in MPRV [65] within the definition of SIM$^+$-CDP (Definition 8) has such an extra correctness requirement, as well as demanding efficient protocols and removing the efficiency requirement of the simulator.

Definition 17 (Standalone security as in MPRV [65], Reformulated). *We say that a protocol π is a secure protocol for the functionality $\mathcal{F}$ if it fulfills Definition 16 with the following changes:*

1. *π must be efficiently computable (PPT);*
2. *π must have perfect correctness, that is, in an honest execution of π, its output distribution is identical to that of $\mathcal{F}$;*
3. *the simulator is allowed to be inefficient.*

In the standalone model, as the name implies, the security of the protocol is considered in isolation, meaning that the distinguisher is constrained to what other protocols it can run in order to try and distinguish the worlds. Making

[20] Most notably, in MPRV [65], the simulators are allowed to be inefficient (computationally unbounded) in the definition of SIM$^+$-CDP.

[21] In Definition 16 it is the output rather than the view of the adversary that is considered. These two formulations are equivalent since the adversary is allowed to simply output its entire view as output.

such a restriction makes proving security technically easier, for example by allowing so-called rewinding techniques. The drawback of the model is precisely that it considers protocol security in isolation, opening up the possibility that a protocol thought to be secure loses all of its security properties when it is run in parallel to some other processes. Since it can be argued that such composition of protocols and processes is the rule rather than the exception in modern computer systems, it is highly desirable to be able to prove that a protocol remains secure also when other protocols are run in composition to it.

There are many ways to compose protocols and some of them are easier to deal with than others. For instance, the usual formulations of the standalone model guarantee that security is preserved under *sequential composition*, meaning as long as all the surrounding protocols are run sequentially (one after another). The most powerful type of composition results are those when the security of the protocol is preserved regardless of how the surrounding protocols are run (in particular when they run concurrently to the protocol in question). This is called *universal composition* and the entire point of the UC (Universal Composability) security framework is that protocols proven within it remain secure under universal composition. This means, in particular, that if a protocol π realises the ideal functionality $\mathcal{F}$, then any other protocol that uses $\mathcal{F}$ as a subprocedure does not lose its security properties if $\mathcal{F}$ is replaced by a copy of π. In the UC framework, the distinguisher goes from taking part in the protocol (as in the standalone model) to being an external entity that observes and interacts with the system from the outside. The distinguisher is captured in an entity called *the environment*, which is an entity in both worlds that selects the initial inputs to all parties, interacts arbitrarily with the adversary and then, based on the outputs of each party at the end, tries to distinguish between the two worlds. In other words, the environment gets to play with one of the worlds and depending only on the input-output behaviour of this world it tries to determine which world it is playing with.

Definition 18 (UC security [44, 11]). *We say that an efficient protocol π UC-securely realises the ideal functionality $\mathcal{F}$ if for all efficient real-world adversaries $\mathcal{A}$ there exists an efficient simulator*[22] *$\mathcal{S}$ (corrupting the same parties as $\mathcal{A}$) such that for all efficient environments E, the statistical distance between E's output when interacting with the ideal world and that when interacting with the real world is negligible in the security parameter κ.*

B.1 The Arithmetic Black-box

In Figure 5 we present the ideal functionality $\mathcal{F}_{ABB}$ of the arithmetic black-box. The ABB is at times formulated slightly differently, such as only operating within the arithmetic domain, not including conversions between the domains or including conversions in both directions in between the binary and arithmetic

[22] Also called *ideal-world adversary.*

representations. We choose the flavor of ABB that is used in [35], simply because it includes the operations we need but nothing more. For more details on the ABB see, for instance, [23, 60].

Functionality $\mathcal{F}_{ABB}$

Parameters: A modulus q that is either a prime or a power of 2.

Input:

– Upon input $(\mathsf{Input}, P_i, \mathsf{type}, \mathsf{id}, x)$ from P_i and input $(\mathsf{Input}, P_i, \mathsf{type}, \mathsf{id})$ from all other parties, if id is a fresh identifier and $(\mathsf{type}, x) \in \{\{\mathsf{binary}\} \times \mathbb{Z}_2, \{\mathsf{arithmetic}\} \times \mathbb{Z}_q)\}$ then store $(\mathsf{type}, \mathsf{id}, x)$.

Linear combination:

– Upon input $(\mathsf{LinComb}, \mathsf{type}, \mathsf{id}, (\mathsf{id}_j)_{j=1}^m, c, (c_j)_{j=1}^m)$ from all parties, if
  - each id_j is stored in memory, and,
  - $c, c_j \in \mathbb{Z}_2$ if type is binary and $c, c_j \in \mathbb{Z}_q$ if type is $\mathsf{arithmetic}$,

  then
  - retrieve $((\mathsf{type}, \mathsf{id}_1, x_1), \ldots, (\mathsf{type}, \mathsf{id}_m, x_m))$,
  - compute $y \leftarrow \sum c_j \cdot x_j \mod 2$ if type is binary and $y \leftarrow \sum c_j \cdot x_j \mod q$ if type is $\mathsf{arithmetic}$,
  - store $(\mathsf{type}, \mathsf{id}, y)$.

Multiplication:

– Upon input $(\mathsf{Mult}, \mathsf{type}, \mathsf{id}, \mathsf{id}_1, \mathsf{id}_2)$ from all parties, if $\mathsf{id}_1, \mathsf{id}_1$ are stored in memory then
  - retrieve $(\mathsf{type}, \mathsf{id}_1, x_1), (\mathsf{type}, \mathsf{id}_2, x_2))$,
  - compute $y \leftarrow x_1 \cdot x_2 \mod 2$ if type is binary and $y \leftarrow x_1 \cdot x_2 \mod q$ if type is $\mathsf{arithmetic}$,
  - store $(\mathsf{type}, \mathsf{id}, y)$.

Converting from binary to arithmetic:

– Upon input $(\mathsf{ConvertB2A}, \mathsf{id}, \mathsf{id}')$ from all parties, if id' is present in memory then retrieve $(\mathsf{binary}, \mathsf{id}', x)$ and store $(\mathsf{arithmetic}, \mathsf{id}, x)$.

Output:

– Upon input $(\mathsf{Output}, \mathsf{type}, \mathsf{id})$ from all honest parties, if id' is present in memory then retrieve $(\mathsf{type}, \mathsf{id}, x)$ and output it to the adversary. Wait for input from the adversary of the form $(\mathsf{Deliver}, b)$, where $b \in \mathbb{Z}_2$. if $b = 1$ then output x to all parties, otherwise output $\perp$.

Fig. 5: The ideal functionality for the arithmetic black-box.

C On SFE with DP Leakage

As noted shortly in Section 4, one cryptographic task that SIM*-CDP can handle but SIM+-CDP cannot is that of computing a differentially private mechanism whilst allowing the adversary to receive leakage throughout the protocol, as long as that leakage is DP, in particular when some leakage occurs before the corrupted party chooses their input. Joint computation of functions whilst allowing DP leakage has been studied in a few different settings with regards to output functions and adversarial models [61, 44, 69, 5]. Of particular interest to us is the work of [44] where it is proposed an ideal functionality in UC for this setting which is then realised with respect to private set intersection (PSI) in the presence of active corruptions. One reason that the ideal functionality of [44] cannot be expressed as an instantiation of SFE (the functionality used in SIM+-CDP) is that the functionality in [44] relaxes the guarantee of *input independence*, meaning that the corrupted party can choose their input based on the input of the honest party.

The PSI protocol of [44] outputs the exact set intersection (to only one of the parties, the other gets no output) and therefore their protocol as a whole intuitively cannot be SIM*-CDP. If one would instead realise their ideal functionality for computing a function with leakage, and enforce that all possible combinations of leakage functions and the output function to the corrupted party is DP (when seen as a composition), then SIM*-CDP can be achieved. Below in Figure 6 we re-iterate the ideal functionality from [44] but augmented to have two potentially different classes of leakage functions for each party. The need for this is that since f_1 and f_2 need not be the same, as in the case when only one of them gets an output, then one can allow the party whose output function is DP with better parameters to have leakage functions that use up more of the privacy budget. In Definition 19 we reiterate [44]'s definition of SFE with DP leakage, reformulated for consistency with our notation.

Definition 19 (SFE with DP leakage [44]). *A protocol π securely realises f with leakage $(\mathcal{L}_1, \mathcal{L}_2)$ if π is a UC-secure protocol for $\mathcal{F}_{SFE\ with\ leakage}^{f,\mathcal{L}_1,\mathcal{L}_2}$ (see Figure 6). We say that the protocol realises f with $(\varepsilon_\kappa, \varepsilon_\kappa)$-SDP leakage if it realises f with $(\mathcal{L}_1, \mathcal{L}_2)$ if for every $(L_{ji}^{pre}, L_{ji}^{post}) \in \mathcal{L}_1 \cup \mathcal{L}_2$, the probabilistic function $D := D_1 || D_2 \to (L_{ji}^{pre}(D), L_{ji}^{post}(D))$ is $(\varepsilon_\kappa, \varepsilon_\kappa)$-SDP.*

D Proofs

D.1 Proof of Proposition 1

Proof overview. We now prove Proposition 1, which in short says that in the plain model (without setup assumptions) there exists tasks and parameter regimes for which SIM+-CDP can be satisfied but not SIM*-CDP with the same parameters. This follows from the results that some ideal functionalities cannot

Functionality $\mathcal{F}_{\text{SFE with leakage}}^{f,\mathcal{L}_1,\mathcal{L}_2}$

Parameters:

- A function $f = (f_1, f_2) : (\{0,1\}^*)^2 \to (\{0,1\}^*)^2$.
- Two classes of functions $\mathcal{L}_j = \{(L^{pre}, L^{post})_{j1}, \ldots\}$, $j \in \{1,2\}$, with $L_{ji}^{pre}, L_{ji}^{post} : \{0,1\}^* \to \{0,1\}^*$.

No corruptions:

- Upon D_1 from P_1 and D_2 from P_2, deliver $f_1(D_1||D_2)$ to P_1 and $f_2(D_1||D_2)$ to P_2.

Party P_c corrupted (P_h is honest):

- Upon D_h from P_h and (Leak, L^{pre}) from P_c, if there exists an element $(L^{pre}, \cdot)$ in $\mathcal{L}_c$ then send $L^{pre}(D_h)$ to P_c, otherwise send $\perp$.
- Upon D_c and $(\mathsf{Leak}, L^{post})$ from P_c, if there exists an element (L^{pre}, L^{post}) in $\mathcal{L}_c$ then send $L^{post}(D_h)$ to P_c, otherwise send $\perp$. Regardless, also send $f_c(D_1||D_2)$ to P_c.
- Upon $(\mathsf{Deliver}, b)$ from P_c, if $b = 1$ then send $f_h(D_1||D_2)$ to P_h, otherwise send $\perp$.

Fig. 6: The ideal functionality for reactive two-party SFE with abort and leakage.

be realised with UC security in the plain model, and this particularly holds for a large class of same-output probabilistic two-party functionalities, as proven in [12]. Such results yield the desired separation after noting that some optimal SDP mechanisms fall within that class of functionalities, and that they are directly computable with SIM$^+$-CDP by use of general-purpose *standalone secure* two-party computation.

Background. The UC part of the proof is in the *plain model*, meaning that no setup assumptions (such as having a common reference string) are made, and that one by default only has access to authenticated (not necessarily secure) channels [12, 13]. In the plain model, it has been shown that there are many functionalities that cannot be realised with UC-security. At the same time, it is known that in the standalone model (also here without any assumptions apart from that there are authenticated channels), all two-party PPT functionalities can be realised. Further, any protocol that securely realises an ideal functionality computing an SDP mechanism in the standalone model also satisfies SIM$^+$-CDP with unchanged parameters. Therefore, if there exists a task for which there is an optimal mechanism and this mechanism can be realised with standalone security but not UC-security (without setup assumptions), then our desired protocol will follow. We now state some definitions and results needed for our proof.

Definition 20 (Unpredictable function family [12]). *A probabilistic function family $f = \{f_\kappa\}_{\kappa \in \mathbb{N}}$ with $f\kappa : \mathcal{D}^2 \to \mathcal{R}$ is said to be unpredictable if there exists a polynomial $p(\cdot)$ such that for infinitely many κ: $\exists D_1, D_2 \in \mathcal{D}$ such that:*

1. $\forall D_2' \in \mathcal{D}, z \in \mathcal{R} : \mathbb{P}(f_\kappa(D_1, D_2') = z) \leq \frac{1}{p(\kappa)}$.
2. $\forall D_1' \in \mathcal{D}, z \in \mathcal{R} : \mathbb{P}(f_\kappa(D_1', D_2) = z) \leq \frac{1}{p(\kappa)}$.

Intuitively, f is unpredictable if (asymptotically) there are no choices of inputs any one of the parties can make such that the function output is almost always the same, regardless of the inputs of the other party. In other words, each party can choose its input such that the function output will not have almost all its probability mass at one output event. This is a pretty weak requirement on a probabilistic function. In particular, we have that the randomised response mechanism for binary functionalities (i.e. where for a binary function $f : \{0,1\}^2 \to \{0,1\}$, $f(D_1, D_2)$ is output with probability $\frac{e^\varepsilon}{e^\varepsilon+1}$ and otherwise its negation is output) is unpredictable, with p in the definition of unpredictability being chosen as a suitable constant. Further, it is easy to verify that for all binary functions, no pure SDP mechanism can have a higher probability of returning the true value than randomised response. Thus we can set u to be 1 iff the mechanism outputs the correct evaluation of a fixed arbitrary binary function, and α to be $\frac{e^\varepsilon}{e^\varepsilon+1}$, i.e. the probability of a correct answer when using randomised response with parameter ε.

Lemma 4 (Reformulation of Theorem 6.1 in [12]). *Let $\mathcal{M} = \{\mathcal{M}_\kappa\}$ be a family of unpredictable PPT two-input same-output functions and let $\mathcal{F}$ be the ideal functionality that returns (to both players) a sample from $\mathcal{M}(D_1, D_2)$ when given D_1 from party 1 and D_2 from party 2. Then $\mathcal{F}$ cannot be UC-realised in the plain model by any non-trivial protocol.*

The notion of a non-trivial protocol is an extremely broad one, essentially only requiring that all parties get outputs in the case that all parties are honest and the adversary does not prevent any messages from being delivered. It is immediately clear that the protocol which realises the ideal functionality of randomised response in the standalone model with perfect correctness by use of general-purpose two-party computation is indeed non-trivial.

Proof (of Proposition 1). What needs to be presented is a choice of ε and a task (α, u) together with proofs that it can be solved with $(\varepsilon, 0)$-SIM$^+$-CDP but not $(\varepsilon, 0)$-SIM*-CDP in the plain model. As explained in the preceding paragraphs, we define the utility function by choosing the boolean XOR function $f : \{0,1\}^2 \to \{0,1\}$ and set $u(D, z) = 1 \iff z = f(D_1, D_2) := D_1 \oplus D_2$. We set $\alpha := \frac{e^\varepsilon}{e^\varepsilon+1}$. Further, set ε such that there exists a polynomial p in κ such that α is a multiple of $2^{-p(\kappa)}$ for large enough κ. The randomised response mechanism with parameter ε, by construction, has α-usefulness for u, since α is exactly the probability with which the mechanism outputs the true answer. We may now make the following observations:

- There is no $(\varepsilon, 0)$-SDP mechanism with higher utility than α for u. This follows directly from a simple contradiction argument, namely that if a boolean mechanism has utility above α, then it cannot be $(\varepsilon, 0)$-SDP since there exists a choice of database such that the privacy loss random variable is above e^ε.

– If there exists a polynomial p in κ such that α is a multiple of $2^{-p(\kappa)}$ for large enough κ, then randomised response is computable in strict polynomial time, which implies (by the possibility of computational perfect-correctness general-purpose two-party computation in the standalone model, see e.g. Theorem 7.1.2 of [39]) that the ideal functionality that computes randomised response with respect to f can be realised in the standalone model.

The second observation above implies that there is a protocol which is $(\varepsilon, 0)$-SIM$^+$-CDP for the task (α, u), namely the one that uses general-purpose two-party computation to realise (with perfect correctness) the ideal functionality which performs randomised response with respect to f. The first observation above gives that randomised response is optimal in the sense that no other mechanism has higher utility and that any mechanism with utility indistinguishable from that of randomised response, also must have an output distribution which is indistinguishable from that of randomised response, due to the boolean output range.

Thus the only thing that remains to be shown is that there is no protocol that has a utility indistinguishable from that of the protocol above whilst satisfying $(\varepsilon, 0)$-SIM*-CDP. This follows directly from the fact that the randomised response ideal functionality is unpredictable (in the sense of Definition 20) together with Lemma 4. In particular, any $(\varepsilon, 0)$-SIM*-CDP protocol which has utility indistinguishable from α has an output distribution that is indistinguishable from that the randomised response ideal functionality, which implies that the protocol UC-realises said functionality, which is a contradiction to Lemma 4 since randomised response is unpredictable. Finally, we note that since the impossibility holds not only for protocols with exactly the utility α but also those with utility computationally indistinguishable from α, the impossibility result holds for SIM*-CDP protocols both with respect to real-world accuracy and ideal-world accuracy.

D.2 Proof of Proposition 2

Proof overview. We now prove Proposition 2, which says that there are tasks such that in some parameter regimes, they can be solved with SIM*-CDP but not SIM$^+$-CDP. The whole idea of the proof is that, under sufficient complexity assumptions, both UC-security and standalone security allow general-purpose two-party computation, meaning that any PPT functionality can be realised. Now the requirement of perfect correctness in SIM$^+$-CDP means on the other side that no non-PPT ideal functionalities can be used to achieve SIM$^+$-CDP, whereas this is not the case for SIM*-CDP, since there the correctness requirement is computational rather than perfect. Again, we use the concrete proof strategy applied in Section D.1, namely showing that there is a parameter regime for which the randomised response mechanism can be realised with SIM*-CDP but not with SIM$^+$-CDP and this gives the result by the fact that randomised response is optimal for pure SDP boolean mechanisms.

Proof (of Proposition 2). As in the proof of Proposition 1, define $u(D, z) = 1 \iff z = f(D_1, D_2) := D_1 \oplus D_2$ i.e. with f being the XOR function. Set $\alpha(\kappa) := \frac{e^{\varepsilon_\kappa}}{e^{\varepsilon_\kappa}+1}$. Now, as opposed to the previous proof, we set ε such that the resulting randomised response mechanism can *not* be computed in strict polynomial time. In particular, set ε_κ such that $\alpha(\kappa) = 1 - 2^{-2^\kappa}$, i.e. $\varepsilon_\kappa := \frac{2^{-2^\kappa}}{1-2^{-2^\kappa}}$. That is, the mechanism we consider is that in which $f(D)$ is returned except for which probability 2^{-2^κ}.

This mechanism can be realised (with computational correctness) with UC security under the common reference string (CRS) setup assumption, since that model allows general-purpose two-party computation (see, for instance [12, 13]). That is, with utility considered in the ideal world (i.e. the ideal functionality $\mathcal{F}$ is α-useful for u), the task (α, u) can be solved with SIM*-CDP.

On the other hand, the mechanism above can not be realised in the standalone model with perfect correctness, since it requires sampling a Bernoulli trial with parameter α, which is impossible in strict polynomial time since α is not a multiple of an inverse polynomial power of 2 (see Appendix A). Further, since randomised response is optimal for boolean functionalities, there is no $(\varepsilon_\kappa, 0)$-SDP mechanism that runs in strict polynomial time and has utility above α. Thus there is no PPT mechanism which is α-useful for u and consequently there is no $(\varepsilon_\kappa, 0)$-SIM$^+$-CDP protocol which is α-useful for u either.

D.3 Proof of Proposition 3

We now prove Proposition 3, which in short says that for all protocols SIM*-CDP implies SIM-CDP with unchanged parameters. Note that this proof is analog that of the same relation between SIM$^+$-CDP and SIM-CDP, which is found in the long version of MPRV [65]. The proof follows essentially directly from the two definitions involved after noting that the mechanism (simulator) in SIM-CDP has access to all of the inputs and thus can run copies of the ideal world internally.

Proof (of Proposition 3). Let $\varepsilon_\kappa \geq 0, \delta_\kappa \in [0, 1]$ be arbitrary fixed classes of parameters. Let π be a two-party protocol that satisfies $(\varepsilon_\kappa, \delta_\kappa)$-SIM*-CDP with respect to some arbitrary fixed ideal functionality $\mathcal{F}$. This implies that for all PPT adversaries $\mathcal{A}$, the view of $\mathcal{S}_\mathcal{A}$ (the ideal-world adversary corresponding to $\mathcal{A}$) is $(\varepsilon_\kappa, \delta_\kappa)$-SDP and that the views of $\mathcal{A}$ and $\mathcal{S}_\mathcal{A}$ are computationally indistinguishable. That is,

$$\forall D \in \mathcal{D}, \mathcal{A} : \mathsf{VIEW}^{\mathcal{A}}_{\pi_{real}} \underset{c}{\approx} \mathsf{VIEW}^{\mathcal{S}_\mathcal{A}}_{\pi_{ideal}}. \tag{8}$$

This is due to the definition of UC-security, because if these two random variables are not computationally indistinguishable, then there exists an efficient environment that distinguishes the real and ideal worlds with a non-negligible advantage over guessing.

We can now turn the simulator $\mathcal{S}_{\mathcal{A}}$ into the mechanism $\mathcal{M}$ in the SIM-CDP definition by letting $\mathcal{M}$ run a copy of the ideal world protocol and then output the view of $\mathcal{S}_{\mathcal{A}}$. This is possible since in SIM-CDP, the mechanism (also at times called the simulator) $\mathcal{M}$ has access to the inputs of both the corrupted and honest parties. That is, since $\mathcal{M}$ can run a copy of the ideal world (thus making $\mathcal{M}(D)$ identically distributed to $\mathsf{VIEW}^{\mathcal{S}_{\mathcal{A}}}_{\pi_{ideal}}$), Equation 8 implies that the view of the adversary in the real world is computationally indistinguishable from the output of $\mathcal{M}(D)$ for all D, which is $(\varepsilon_\kappa, \delta_\kappa)$-SDP, and thus π is $(\varepsilon_\kappa, \delta_\kappa)$-SIM-CDP.

D.4 Proof of Lemma 1

Proof. Let $Z \sim \mathcal{M}^{p,f,\lambda}_{RTGeo}(D)$ and $Y \sim \mathcal{M}^{2B,f,\lambda}_{SRTGeo}(D)$ for arbitrary λ, D. Let p_Z and p_Y denote the probability density functions of Z and Y respectively and let F denote their cumulative distribution functions in the same manner. Since the parameter restrictions guarantee that the final sum in Y does not overflow (the result is as if the sum was done over the integers), the statistical distance between the two distributions is exactly twice the total probability mass that is affected by the truncation in Y. That is,

$$\begin{aligned}
SD(Z,Y) &= \frac{1}{2}\sum_{z\in\mathbb{Z}_p} |p_X(z) - p_Y(z)| \\
&= \sum_{z\in\mathbb{Z}_p\setminus(\bar{f}-B,\bar{f}+B)} |p_X(z) - p_Y(z)| \\
&= |F_X(\bar{f}-B) + (1 - F_X(\bar{f}+B))| \\
&= \left| \frac{e^{1/\lambda}}{e^{1/\lambda}+1} e^{-(\bar{f}-\bar{f}+B)/\lambda} \right. \\
&\quad \left. + \frac{1}{e^{1/\lambda}+1} e^{-(\bar{f}+B-\bar{f})/\lambda} \right| \\
&= e^{-B/\lambda},
\end{aligned}$$

where $\bar{f}$ is shorthand for $f(D)$. The equalities follow by inserting the formulas from Definition 11 and direct simplifications.

D.5 Proof of Lemma 2

Proof. Firstly, $\texttt{Ber}_{\hat{\alpha}}$ exactly samples a Bernoulli trial with a parameter equal to the recomposition of the first d elements of α. Call this parameter value α'. This means that the statistical distance between $\texttt{Ber}(\hat{\alpha})$ and an exact Bernoulli trial with parameter $\hat{\alpha}$ is the same as between two exact Bernoulli trials with parameter $\hat{\alpha}$ and α', respectively. This statistical distance is equal to $|\hat{\alpha} - \alpha'|$, which is at most 2^{-d} since the first 2^d bits of their decomposition are identical.

Secondly, the statistical distance between $\mathcal{M}_{\texttt{FDL}}^{\lambda,B,d,h}(D)$ and $\mathcal{M}_{SRTGeo}^{2B,h,\lambda}(D)$ is at most equal to the probability of any of the Bernoulli trials being incorrect, which due to independence is at most $B2^{-d}$.

D.6 Proof of Lemma 3

Proof. The additive usefulness follows from a standard tail bound on the geometric distribution, since the truncated geometric is at least as concentrated as the untruncated one:

$$\begin{aligned}\mathbb{P}(|Geo_{q,\lambda}(f(D)) - f(D)| \geq \nu) &= \mathbb{P}(|Geo_{q,\lambda}(0)| \geq \nu) \\ &\leq \mathbb{P}(|Geo_{\lambda}(0)| \geq \nu) \\ &= 2F_{Geo_\lambda(0)}(-\nu) \\ &= \frac{2e^{1/\lambda}}{e^{1/\lambda}+1}e^{-\nu/\lambda}.\end{aligned}$$

D.7 Proof of Theorem 1

Proof. The definition of SIM*-CDP demands two things to be shown, namely that the view of the simulator is SDP and that the protocol UC-realises the ideal functionality. The first requirement is fulfilled as the only message sent from $\mathcal{F}_{SFE}^{\hat{f}_\kappa}$ to the corrupted party is $\mathcal{M}_{RTGeo}^{q,f_\kappa,\lambda_\kappa}(D)$ and this is $(\varepsilon_\kappa, 0)$-SDP due to the fact that the range-truncated geometric mechanism is $(\varepsilon_\kappa, 0)$-SDP under the standard parametrisation specified in the theorem. The other parts of the view of $\mathcal{S}$ (like its input and randomness tape) are independent of the inputs of the honest party, thus making the view of $\mathcal{S}$ as a whole $(\varepsilon_\kappa, 0)$-SDP. Further, this holds for all types of malicious behavour of $\mathcal{S}$ since, due to the formulation of $\mathcal{F}_{SFE}$, the only way $\mathcal{S}$ can change its view is to refuse to collaborate in the protocol or change its inputs and both of those decisions would have to be made independently of the inputs of the honest party (thus making those decisions $(0, 0)$-DP as well).

The UC-realisation of the ideal functionality follows directly from the use of the arithmetic black-box and the statistical indistinguishability between $\mathcal{M}_{\texttt{FDL}}$ and $\mathcal{M}_{RTGeo}$, which follows from lemmas 1 and 2 together with the assumptions of the theorem. In particular, due to the use of $\mathcal{F}_{ABB}$, the view of the corrupted party in the hybrid world consists of only its input, random coins and the output returned from $\mathcal{F}_{ABB}$, which is exactly $\mathcal{M}_{\texttt{FDL}}$. Similarly, the view of the corrupted party in the ideal world is also only its input, random coins and output returned from $\mathcal{F}_{SFE}^{\hat{f}_\kappa}$. Therefore the simulator that simply outputs its view (after having changed its inputs and/or aborted with respect to its random coins as the hybrid-world adversary does) yields a view that is computationally indistinguishable from that of the hybrid-world adversary. Further, this simulator is strict PPT due to it performing only the same operations as the hybrid-world adversary (choosing input and abort behaviour based on its coins and then receiving one $\mathbb{Z}_q$ element), hence the theorem follows.

E Techniques for Achieving Secure MPC

In the context of MPC, we typically distinguish binary and arithmetic protocols. This classification describes the possible computations. In other words, we perform addition and multiplication in $\mathbb{F}_2$ and $\mathbb{F}_p$, respectively. In this work, we rely on secret sharing-based (SS) MPC protocols. More precisely, we use additive secret sharing (ASS). In such protocols, secret values x are shared among n parties by uniformly sampling $n-1$ random values $x_1, \ldots, x_{n-1}$ from $\mathbb{F}$, setting $x_0 \leftarrow x - \sum_{i=1}^{n} x_i$, and distributing x_i to every party p_i. We denote secret shared values as $[\![x]\!]$. We further denote $[\![x]\!] \leftarrow \mathsf{Share}(x)$, and $x \leftarrow \mathsf{Reconstruct}([\![x]\!])$ as sharing and reconstructing secrets. ASS schemes are additively homomorphic, allowing the addition of shares without interaction and hiding underlying secrets as long as there is one honest party. To allow multiplications with an ASS, one can use multiplication triples, introduced by Beaver [3]. Triples are three shared values $([\![a]\!], [\![b]\!], [\![c]\!])$, that no party knows and that fulfil $a \cdot b = c$. When multiplying two shared values $([\![x]\!], [\![y]\!])$, one reconstructs masked versions $\alpha \leftarrow \mathsf{Reconstruct}([\![x]\!] - [\![a]\!])$, $\beta \leftarrow \mathsf{Reconstruct}([\![y]\!] - [\![b]\!])$, and computes[23] $[\![z]\!] \leftarrow \alpha\beta + \beta[\![x]\!] + \alpha[\![y]\!] + [\![c]\!] = [\![x \cdot y]\!]$.

Given these ingredients, we can instantiate an active secure general-purpose MPC protocol if we have access to a secure sampling method for multiplication triples, and adversaries cannot tamper with the reconstruction procedure. In the SPDZ paper [24], the authors introduced solutions to both problems. They propose an additively homomorphic encryption scheme for sampling triples and information-theoretic message authentication codes (MACs) to secure the reconstruction procedure. Subsequent work introduced several performance improvements by instantiating the ASS over the ring $\mathbb{F}_{2^k}$ [21] or replacing the expensive homomorphic encryption with oblivious transfer [54]. Note that both improvements, to some degree, accept a higher communication for a lower computation complexity.

[23] This step requires multiplication and addition with constant terms which follows from the ASS properties.