Composing Differential Privacy and Secure Computation: A case study on scaling private record linkage

Xi He
Duke University
hexi88@cs.duke.edu

Cheryl Flynn
AT&T Labs-Research
cflynn@research.att.com

Ashwin Machanavajjhala
Duke University
ashwin@cs.duke.edu

Divesh Srivastava
AT&T Labs-Research
divesh@research.att.com

ABSTRACT

Private record linkage (PRL) is the problem of identifying pairs of records that are similar as per an input matching rule from databases held by two parties that do not trust one another. We identify three key desiderata that a PRL solution must ensure: (1) perfect precision and high recall of matching pairs, (2) a proof of end-to-end privacy, and (3) communication and computational costs that scale subquadratically in the number of input records. We show that all of the existing solutions for PRL – including secure 2-party computation (S2PC), and their variants that use non-private or differentially private (DP) blocking to ensure subquadratic cost – violate at least one of the three desiderata. In particular, S2PC techniques guarantee end-to-end privacy but have either low recall or quadratic cost. In contrast, no end-to-end privacy guarantee has been formalized for solutions that achieve subquadratic cost. This is true even for solutions that compose DP and S2PC: DP does not permit the release of any exact information about the databases, while S2PC algorithms for PRL allow the release of matching records.

In light of this deficiency, we propose a novel privacy model, called *output constrained differential privacy*, that shares the strong privacy protection of DP, but allows for the truthful release of the output of a certain function applied to the data. We apply this to PRL, and show that protocols satisfying this privacy model permit the disclosure of the true matching records, but their execution is insensitive to the presence or absence of a single non-matching record. We find that prior work that combine DP and S2PC techniques even fail to satisfy this end-to-end privacy model. Hence, we develop novel protocols that provably achieve this end-to-end privacy guarantee, together with the other two desiderata of PRL. Our empirical evaluation also shows that our protocols obtain high recall, scale near linearly in the size of the input databases and the output set of matching pairs, and have communication and computational costs that are at least 2 orders of magnitude smaller than S2PC baselines.

1 INTRODUCTION

Organizations are increasingly collecting vast amounts of data from individuals to advance science, public health, and resource management and governance. In a number of scenarios, different organizations would like to collaboratively analyze their data in order to mine patterns that they cannot learn from their individual datasets. For instance, hospitals or health workers in neighboring cities might want to identify HIV positive patients who have sought care in multiple cities to quantify the mobility patterns of patients, and hence the spread of the virus. This requires finding patients who occur in multiple databases even though the patient records might not have the same primary key across databases. This problem is called *record linkage*, and has been well studied for the last several decades [6, 9, 13]. In a collaborative analysis across organizations, privacy is always a concern. In particular, one of the collaborating parties, say Hospital A, should not be able to tell whether or not a record is in the database of the other party, say Hospital B, if that record does not appear in the match output. Privacy constraints arise due to concerns from individuals who provide their data, such as hospital patients, or due to contractual or legal obligations that organizations have to the individuals in their data. This has led to a field of research called *private record linkage* (PRL).

Traditional PRL techniques aim to solve the linkage problem with a strong privacy goal – no information should be leaked beyond (a) the sizes of the datasets, and (b) the set of matching records. However, this strong privacy goal (which we call S2PC) [15] comes with a high cost. Existing techniques that achieve this goal either require cryptographically secure comparisons of all pairs of records (and hence are inefficient), or are restricted to equi-joins (and thus have very low recall). Hence, we formalize our problem as follows: *given private databases D_A and D_B held by two semi-honest parties, and a matching rule m , design a protocol Π that outputs pairs of matching records to both parties and satisfies three desiderata: (1) correctness in terms of perfect precision and high recall of matches, (2) provable end-to-end privacy guarantee, and (3) efficiency in terms of sub-quadratic communication and computational cost in n , where $n = \max(|D_A|, |D_B|)$. There are two sources of the cost incurred by PRL: (1) the number of cryptographic operations, and (2) the time taken for each cryptographic operation. Our protocols aim to reduce the number of cryptographic operations (i.e., the number of secure pairwise comparisons), the first source of cost, while using existing techniques to securely compare pairs of records.*

Techniques that securely compare all pairs of records (APC) have a quadratic cost and hence fail to meet the efficiency requirement

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

CCS '17, October 30–November 3, 2017, Dallas, TX, USA

© 2017 Association for Computing Machinery.

ACM ISBN 978-1-4503-4946-8/17/10...\$15.00

<https://doi.org/10.1145/3133956.3134030>

of our problem. On the other hand, techniques for efficient private set intersection (PSI) [12, 32] satisfy all three desiderata for equality-like matching rules, but result in poor recall for general fuzzy matching rules. When records in D_A and D_B come from the same discrete domain, one could expand D_A by adding all records that could potentially match with a record in D_B , and then find matches by running PSI on the expanded D_A and D_B . However, this technique can be very inefficient: the expanded databases could be much larger than the input databases for complex matching functions or when data are high dimensional. A long line of work [5, 18, 19, 21, 24, 34] has considered scaling APC by using *blocking*, which is a standard technique for scaling non-private record linkage with a small loss in recall of matching pairs. However, blocking can reveal sensitive properties of input records. We show that such hybrid protocols do not ensure an end-to-end privacy guarantee even in solutions where the blocking step satisfies a strong privacy notion, called differential privacy (DP) [10]. This negative result is in contrast to other success stories [2, 16, 26, 27, 30, 31, 39] on composing DP and secure computation. These settings either consider a client-server model where all data sits on the server or consider aggregated functions across partitioned data where the privacy goals of DP and secure computation do not conflict. In the case of scaling PRL, neither blocking nor DP blocking naturally composes with the strong privacy guarantee of S2PC. To our knowledge, this work presents the first solution to the above open problem, and makes the following contributions:

- We propose and formalize three desiderata for the PRL problem: (1) correctness, or perfect precision and high recall of matches, (2) provable end-to-end privacy, or insensitivity to the presence or absence of an individual record that is not a matching record, and (3) efficiency, or communication and computational costs that scale subquadratically in the input size. We show that all of the existing solutions for PRL violate at least one of these three desiderata. (§ 2)
- This motivates us to develop a novel privacy definition, which we call *Output Constrained DP*. Protocols satisfying this notion are allowed to truthfully return the output of a specific function, but must be insensitive to the presence or absence of individual records that do not affect the function output. (§ 3.1)
- We adapt the notion of Output Constrained DP to the context of PRL. Under this privacy notion, computationally bounded adversaries cannot distinguish two different protocol executions when a single *non-matching* record is replaced by another non-matching record in one of the databases. This privacy notion, named *DPRL*, allows protocols to truthfully release the set of matching records. (§ 3.2)
- We show that prior attempts [5, 19, 24] to scale PRL using blocking do not satisfy our privacy definition DPRL (Theorem 4.7), and hence fail to achieve stronger privacy guarantees including differential privacy or S2PC. (§ 4)
- We develop novel protocols for private record linkage that leverage blocking strategies. Our protocols ensure end-to-end privacy (Theorems 4.5 and 4.10), provide at least as much recall as the non-private blocking strategy (Theorems 4.4 and 4.11), and achieve subquadratic scaling (Theorems 4.8 and 4.11).
- Using experiments on real and synthetic data, we investigate the 3-way trade-off between recall, privacy, and efficiency. Our key findings are: our protocols (1) are at least 2 orders of magnitude more efficient than S2PC baselines, (2) achieve a high recall and end-to-end privacy, and (3) achieve near linear scaling in the size of the input databases and the output set of matching pairs on real and synthetic datasets. (§ 5)

2 PROBLEM SETTING & STATEMENT

In this section, we formulate our problem: finding pairs of records that are similar as per an input matching rule while ensuring three desiderata: *correctness*, *privacy*, and *efficiency*. We then discuss prior attempts to solve this problem and how they do not satisfy one or more of the three aforementioned desiderata, thus motivating the need for a novel solution.

2.1 The Private Record Linkage Problem

Consider two parties Alice and Bob who have databases D_A and D_B . Let records in D_A come from some domain Σ_A and let the records in D_B come from domain Σ_B . Let $m : \Sigma_A \times \Sigma_B \rightarrow \{0, 1\}$ denote a *matching rule*, and let $D_A \bowtie_m D_B$ denote the set of matching pairs $\{(a, b) | a \in D_A, b \in D_B, m(a, b) = 1\}$. A matching rule can be distance-metric based: two records match if their distance is less than a threshold. For example, Euclidean distance is typically used for numeric attributes, whereas for string attributes, the distance metric is typically based on q-grams [7, 8, 35], phonetic encoding [20], or edit distance over strings [3, 29, 33]. A matching rule can also be conjunctions of predicates over different types of attributes. For instance, two records match if their names differ by at most 2 characters and their phone numbers differ by at most 1 digit. Alice and Bob would like to jointly compute $D_A \bowtie_m D_B$.¹

Our goal is to design a protocol Π that Alice and Bob can follow to compute $D_A \bowtie_m D_B$, while satisfying the following three desiderata – correctness, privacy and efficiency.

- *Correctness*: Let $O_\Pi \subseteq D_A \times D_B$ denote the set of pairs output by the protocol Π as the set of matching pairs. The protocol is *correct* if (a) the protocol returns to both Alice and Bob the same output O_Π , and (b) $O_\Pi = D_A \bowtie_m D_B$, and incorrect otherwise. Note that if Alice and Bob indeed receive the same output, O_Π can only be incorrect in one way – some matching pairs $(a, b) \in D_A \bowtie_m D_B$ are not present in O_Π . This ensures perfect precision – no false positives. Hence, we quantify the correctness of a protocol Π using a measure called *recall*, which is computed as:

$$r_\Pi(D_A, D_B) = \frac{|O_\Pi \cap (D_A \bowtie_m D_B)|}{|D_A \bowtie_m D_B|}. \quad (1)$$

We require Π to have a high recall (close to 1). This precludes trivial protocols that output an empty set.

- *Privacy*: We assume that the data in D_A and D_B are sensitive. As part of the protocol Π , Alice would like no one else (including Bob) to learn whether a specific non-matching record a is in or out of D_A ; and analogously for Bob. This precludes the

¹The standard record linkage problem involves learning a matching function in addition to computing the matches. Although the problem considered in this paper and in the private record linkage literature ignores this crucial aspect of record linkage, we have chosen to also use this term for continuity with existing literature on the topic.

Methods	Correctness	Privacy	Efficiency
APC	✓	✓	✗
PSI	✗	✓	✓
PSI+X	✓	✓	✗
PRL+ $\mathcal{B}$	✓	✗	✓
PRL+ $\mathcal{B}_{DP}$	✓	✗	✓

Table 1: Summary of Prior Work

trivial solution wherein Bob sends D_B to D_A in the clear so that Alice can compute $D_A \bowtie_m D_B$ using standard techniques in the record linkage literature [6]. It also precludes the trivial solution wherein Alice and Bob send their records to a trusted third party in the clear who can then compute $D_A \bowtie_m D_B$. Formally stating a privacy definition is challenging (as we will see later in the paper) and is a key contribution of this paper. We will assume throughout the paper that Alice and Bob are semi-honest, i.e., they follow the protocol honestly, but are curious about each others' databases. We also assume that Alice and Bob are computationally bounded, i.e., they are probabilistic polynomially bounded turing machines.

- *Efficiency*: Jointly computing matching records would involve communication and computational cost. We assume that each record in the database has $O(1)$ length; i.e., it does not grow with $n = \max(|D_A|, |D_B|)$. The communication and computational costs are bounded below by the output size, i.e. $\Omega(M)$, where $M = |D_A \bowtie_m D_B|$. If M is quadratic in n , then the costs have to be quadratic in n to ensure high recall. Hence, we consider problems with sub-quadratic output size, and we say that the protocol is efficient if both the communication and computational costs are sub-quadratic in n , i.e., $o(n^2)$.

We formalize our problem statement as follows.

PROBLEM 1 (PRL). *Let D_A and D_B be private databases held by two semi-honest parties, and let m be a matching rule. Design a protocol Π that outputs pairs of matching records to both parties such that (1) Π ensures high recall close to 1, (2) Π provably guarantees privacy, and (3) Π has sub-quadratic communication and computational cost.*

2.2 Prior Work

Before describing our solution, we outline five approaches for the PRL problem from prior work – APC, PSI, PSI+X, PRL+ $\mathcal{B}$ and PRL+ $\mathcal{B}_{DP}$. Table 1 summarizes their (in)ability to satisfy our three desiderata stated in Problem 1. Other related work on composing S2PC and DP is discussed in § 6.

2.2.1 All-Pairwise Comparisons (APC). One approach to solve the PRL problem, which we call APC, works as follows: (1) design a secure 2-party algorithm that takes as input a record $a \in D_A$ and a record $b \in D_B$ and outputs to both parties the pair (a, b) if the value of $m(a, b) = 1$ without leaking any additional information, and (2) run the secure comparison algorithm for every pair of records in $D_A \times D_B$. The secure comparison primitive can be implemented either using garbled circuits [40] or (partially) homomorphic encryption [28], depending on the matching rule. APC

achieves a recall of 1, but requires a quadratic communication and computational cost for $|D_A| \times |D_B|$ secure pairwise comparisons.

APC provides a strong end-to-end privacy guarantee – it leaks no information other than the sizes of the databases and the set of matching records. This guarantee is formalized as follows.

Definition 2.1 (IND-S2PC [15]). A 2-party protocol Π that computes function f satisfies IND-S2PC if for any D_A , and for every pair of D_B and D'_B where $f(D_A, D_B) = f(D_A, D'_B)$, the view of Alice during the execution of Π over (D_A, D_B) is computationally indistinguishable from the view over (D_A, D'_B) , i.e. for any probabilistic polynomial adversary T ,

$$\begin{aligned} & \Pr[T(\text{VIEW}_A^\Pi(D_A, D_B)) = 1] \\ & \leq \Pr[T(\text{VIEW}_A^\Pi(D_A, D'_B)) = 1] + \text{negl}(\kappa); \end{aligned} \quad (2)$$

and the same holds for the view of Bob over (D_A, D_B) and (D'_A, D_B) for $f(D_A, D_B) = f(D'_A, D_B)$. $\text{negl}(\kappa)$ refers to any function that is $o(\kappa^{-c})$, for all constants c , and $\text{VIEW}_A^\Pi(D_A, \cdot)$ ($\text{VIEW}_B^\Pi(\cdot, D_B)$ resp.) denotes the view of Alice (Bob resp.) during an execution of Π .

The IND-S2PC definition uses κ as a “security” parameter to control various quantities. The size of the adversary is polynomial in κ , and the output of the protocol is at most polynomial in κ . The views of the protocol execution are also parameterized by κ .

In PRL, let $f_{\bowtie_m}$ be the function that takes as inputs D_A and D_B , and outputs a triple $(|D_A|, |D_B|, D_A \bowtie_m D_B)$. The view of Alice, $\text{VIEW}_A^\Pi(D_A, \cdot)$, includes $(D_A, r, m_1, \dots, m_t)$, where r represents the outcome of Alice's internal coin tosses, and m_i represents the i -th message it has received. The output received by Alice after an execution of Π on (D_A, D_B) , denoted $\text{O}_A^\Pi(D_A, D_B)$ is implicit in the party's own view of the execution. The view of Bob can be similarly defined. In addition, the output size of VIEW will be (at most) polynomial in κ . Intuitively, IND-S2PC ensures that the adversary Alice cannot distinguish any two databases D_B and D'_B from her view given the constraint $f(D_A, D_B) = f(D_A, D'_B)$, and the same applies to Bob. This IND-S2PC definition is a necessary condition for the standard simulation-based definition (Theorem A.2 in Appendix A.1).

To summarize, APC guarantees end-to-end privacy and provides a recall of 1, but violates the efficiency requirement.

2.2.2 Private Set Intersection (PSI). We call the next class of approaches PSI, since they were originally designed for efficient private set intersection. Like APC, PSI also ensures IND-S2PC and the parties only learn the sizes of the databases and the set of matching records. The algorithms are efficient, but only ensure high recall for equality predicate like matching rules [12, 32].

The basic protocol works as follows: Alice defines a polynomial $p(x)$ whose roots are her set of elements $a \in D_A$. She sends the homomorphic encryptions of the coefficients to Bob. For each element $b \in D_B$, Bob computes the encrypted values $\tilde{b} = r \cdot p(b) + b$, where r is a random value, and sends them back to Alice. These values are decrypted by Alice and then matched with D_A . If $b \notin D_A$, then the decrypted value of $\tilde{b}$ will be a random value not matching any records in D_A ; otherwise, it will find a match from D_A . The basic protocol described thus far required $O(|D_A| + |D_B|)$ communications and $O(|D_A \times D_B|)$ operations on encrypted values. [12] further optimizes the computational cost with Horner's rule and

cryptographic hashing to replace a single high-degree polynomial with several low-degree polynomials. This reduces the computational cost to $O(|D_B| \cdot \ln \ln |D_A|)$, and hence is sub-quadratic in n , for $n = \max(|D_A|, |D_B|)$. State of the art PSI techniques [32] further improve efficiency.

PSI techniques are limited to equality like matching functions, and extensions [12, 41] allow for matching rules that require exact match on at least t out of T features. However these techniques achieve poor recall for general matching rules. For example, they do not extend to matching rules that involve conjunctions and disjunctions of similarity functions evaluated on multiple attributes. They also do not extend to complex distance metrics, such as $\text{CosineSimilarity}(\text{First Name}) > 0.9$ OR $\text{CosineSimilarity}(\text{Last Name}) > 0.9$, which are typical in record linkage tasks [13].

2.2.3 PSI with Expansion (PSI+X). The PSI technique can be used to achieve high recall for general matching rules by using the idea of *expansion*. Suppose D_A and D_B have the same domains, i.e., $\Sigma_A = \Sigma_B = \Sigma$. For every record $a \in D_A$, one could add all records $a' \in \Sigma$ such that $m(a, a') = 1$ to get an expanded database D_A^x . An equi-join between D_A^x and D_B returns the required output $D_A \bowtie_m D_B$, and satisfies IND-S2PC. However, the expanded dataset can be many orders of magnitude larger than the original dataset making this protocol, PSI+X, inefficient (in the size of the original datasets). Moreover, enumerating all matches per record is hard for a complex matching function. For instance, if the matching function m can encode Boolean 3-CNF formulas, then finding values for a such that $m(a, a') = 1$ could be an intractable problem. In such a case, any efficient expansion algorithm may need to enumerate a superset of matches, further increasing the computational cost. Lastly, even for relatively simple matching functions, we empirically illustrate low recall of PSI and inefficiency of PSI+X protocols respectively in § 5.

2.2.4 PRL with Blocking (PRL+B). Blocking is commonly used to scale up non-private record linkage. Formally,

Definition 2.2 (Blocking ($\mathcal{B}$)). Given k bins $\{\mathcal{B}_0, \dots, \mathcal{B}_{k-1}\}$, records in D_A and D_B are hashed by $\mathcal{B}$ to a subset of the k bins. The set of records in D_A (respectively D_B) falling into the i^{th} bin are represented by $\mathcal{B}_i(D_A)$ (respectively $\mathcal{B}_i(D_B)$). A blocking strategy $\mathcal{B}^S \subseteq [0, k) \times [0, k)$ specifies pairs of bins of D_A and D_B that are compared, i.e. records in $\mathcal{B}_i(D_A)$ are compared with records in $\mathcal{B}_j(D_B)$ if $(i, j) \in \mathcal{B}^S$.

We sometimes use $\mathcal{B}$ to refer to the entire blocking algorithm as well as the blocking functions used in the algorithm. We refer to the set of pairs of records that are compared by a blocking strategy as *candidate matches*. A blocking strategy $\mathcal{B}^S$ is *sub-quadratic* if the number of candidate matches

$$\text{cost}_{\mathcal{B}^S}(D_A, D_B) = \sum_{(i,j) \in \mathcal{B}^S} |\mathcal{B}_i(D_A)| |\mathcal{B}_j(D_B)|$$

is $o(n^2)$, for $n = \max(|D_A|, |D_B|)$. Blocking techniques are useful as a pre-processing step [18, 21, 34] to achieve sub-quadratic efficiency and high recall. We can use blocking as a pre-processing step for APC – secure comparison is performed only for the candidate matches – resulting in an efficient protocol with high recall. However, the blocking strategy itself can leak information about

the presence or absence of a record in the database. This was illustrated using an attack by Cao et al. [5]. This is because the number of candidate matches can vary significantly even if D_B and D'_B differ in only one record. We formally prove this negative result for a large class of blocking techniques which use locality sensitive hashing (LSH). A majority of the hash functions used by blocking algorithms like q-gram based hash signatures [1] or SparseMap [34] are instances of LSH.

Definition 2.3 (Locality Sensitive Hashing (LSH)[14]). A family of functions H is said to be (d_1, d_2, p_1, p_2) -sensitive, where $d_2 > d_1$ and $p_1 > p_2$, if for all $h \in H$, (1) if $\text{dist}(a, b) \leq d_1$, then $\Pr[h(a) = h(b)] \geq p_1$, and (2) if $\text{dist}(a, b) > d_2$, then $\Pr[h(a) = h(b)] \leq p_2$.

An LSH-based blocking considers a set of bins where each bin consists of records with the same hash values for all $h \in H$. A popular blocking strategy is to compare all the corresponding bins, and results in a set of candidate matches $\{(a, b) | h(a) = h(b) \forall h \in H, a \in D_A, b \in D_B\}$. In general, we can show that any LSH based blocking cannot satisfy IND-S2PC.

THEOREM 2.4. *An LSH based blocking with a family of (d_1, d_2, p_1, p_2) -sensitive hashing functions H cannot satisfy IND-S2PC.*

The proof can be found in Appendix B.1.1.

2.2.5 PRL with DP Blocking (PRL+ $\mathcal{B}_{DP}$). Differential privacy has arisen as a gold standard for privacy in situations where it is ok to reveal statistical properties of datasets but not reveal properties of individuals. An algorithm satisfies differential privacy if its output does not significantly change when adding/removing or changing a single record in its input. More formally,

Definition 2.5 ((ϵ, δ) -Differential Privacy[10]). A randomized mechanism $M : \mathcal{D} \rightarrow \mathcal{O}$ satisfies (ϵ, δ) -differential privacy (DP) if

$$\Pr[M(D) \in O] \leq e^\epsilon \Pr[M(D') \in O] + \delta \quad (3)$$

for any set $O \subseteq \mathcal{O}$ and any pair of neighboring databases $D, D' \in \mathcal{D}$ such that D and D' differ by adding/removing a record.

A recent line of work has designed differentially private blocking algorithms as a preprocessing step to APC. DP hides the presence or absence of a single record, and hence the number of candidate matches stays roughly the same on D_B and D'_B that differ in a single record. While this approach seems like it should satisfy all three of our desiderata, we have found that none of the protocols presented in prior work (on DP Blocking) [5, 19, 24] provide an end-to-end privacy guarantee. In fact, each paper in this line of work finds privacy breaches in the prior work. We also show in the proof of Theorem 4.7 (Appendix B.1.2) that even the most recent of these protocols in [5] does not satisfy an end-to-end privacy guarantee. This is because of a fundamental disconnect between the privacy guarantees in the two steps of these algorithms. DP does not allow learning any fact about the input datasets with certainty, while IND-S2PC (and PRL protocols that satisfy this definition) can reveal the output of the function f truthfully. On the other hand, while DP can reveal aggregate properties of the input datasets with low error, protocols that satisfy IND-S2PC are not allowed to leak any information beyond the output of f . Hence, DP and IND-S2PC do not naturally compose.

To summarize, none of the prior approaches that attempt to solve Problem 1 satisfy all three of our desiderata. Approaches that satisfy a strong privacy guarantee (IND-S2PC) are either inefficient or have poor recall. Efficient PRL with blocking or DP blocking fail to provide true end-to-end privacy guarantees. A correct conceptualization of an end-to-end privacy guarantee is critical for achieving correctness, privacy and efficiency. Hence, in the following sections, we first define an end-to-end privacy guarantee for PRL to address this challenge (§ 3), and then present algorithms in this privacy framework to achieve sub-quadratic efficiency and high recall (§ 4).

3 OUTPUT CONSTRAINED DP

Designing efficient and correct algorithms for PRL is challenging and non-trivial because there is no existing formal privacy framework that enables the trade-off between correctness, privacy, and efficiency. In this section, we propose a novel privacy model to achieve this goal.

3.1 Output Constrained Differential Privacy

Both IND-S2PC (Def. 2.1) and DP (Def. 2.5) ensure the privacy goal of not revealing information about individual records in the dataset. However, there is a fundamental incompatibility between the two definitions. IND-S2PC reveals the output of a function truthfully; whereas, nothing truthful can be revealed under differential privacy. On the other hand, DP reveals noisy yet accurate (to within an approximation factor) aggregate statistics about all the records in the dataset; but, nothing other than the output of a pre-specified function can be revealed under IND-S2PC.

The difference between these privacy definitions can be illustrated by rephrasing the privacy notions in terms of a distance metric imposed on the space of databases. Without loss of generality, assume Alice is the adversary. Let $\mathcal{G} = (V, E)$ denote a graph, where V is the set of all possible databases that Bob could have and E is a set of edges that connect neighboring databases. The distance between any pair of databases is the shortest path distance in $\mathcal{G}$. Intuitively, the adversary Alice's ability to distinguish protocol executions on a pair of databases D_B and D'_B is larger if the shortest path between the databases is larger.

DP can be represented by the set of edges that connect neighboring databases that differ in the presence or absence of one record, $|D_B \setminus D'_B \cup D'_B \setminus D_B| = 1$. This means, any pair of databases D_B and D'_B are connected in this graph by a path of *finite* length that is equal to the size of their symmetric difference. While an adversary can distinguish protocol executions between some pair of “far away” databases, the adversary can never tell with certainty whether the input was a specific database. On the other hand, under IND-S2PC, *every pair* of databases that result in the same output for $f(D_A, \cdot)$ for a given D_A are neighbors. However, there is neither an edge nor a path between databases that result in different outputs. Thus the output constraint divides the set of databases into disjoint complete subgraphs (in fact equivalence classes).

Example 3.1. Consider databases with domain $\{1, 2, 3, 4, 5, 6\}$. Given $D_A = \{1, 2\}$, the graph $\mathcal{G}$ for the database instances for D_B are shown in Figure 1. For the graph of differential privacy in Figure 1(a), every pair of database instances that differ in one record is connected by an edge and form a neighboring pair. For

Figure 1: Neighboring databases for (a) DP, and (b) IND-S2PC for Example 3.1.

instance, $D_B = \{1\}$ and $D'_B = \{1, 2\}$ are neighbors under DP. Figure 1(b) considers an output which consists of the size of D_B and the intersection between D_B and D_A . Hence, all the instances in $\mathcal{G}_{IND-S2PC}$ have the same dataset size and have the same intersection with $D_A = \{1, 2\}$. For example, the fully connected 6 database instances all have 2 records, but have no intersection with D_A . The instance $\{1, 2\}$ has no neighboring databases, as it is same as the output, and hence none of the records in this database instance requires privacy protection.

Comparing these two graphs, we can see that all instances in $\mathcal{G}_{DP}$ are connected, and hence an adversary can not distinguish protocol executions on any pair of databases with certainty, but is allowed to learn statistical properties (with some error). This is not true under $\mathcal{G}_{IND-S2PC}$, where some instances are disconnected. For instance, an adversary can distinguish between protocol executions on $\{1, 2\}$ and $\{1, 5\}$ since they give different outputs when matched with D_A .

From Example 3.1, it is clear that the privacy guarantees given by DP and IND-S2PC are different. To ensure scalable record linkage with formal privacy guarantees, we need the best of both worlds: *the ability to reveal records that appear in the match truthfully, the ability to reveal statistics about non-matching records, and yet not reveal the presence or absence of individual non-matching records in the dataset.* Hence, we propose a weaker, but end-to-end, privacy definition for the two party setting.

Definition 3.2 (f -Neighbors). Given function $f : \mathcal{D}_A \times \mathcal{D}_B \rightarrow \mathcal{O}$ and $D_A \in \mathcal{D}_A$. For any pairs of datasets D_B, D'_B , let $\Delta(D_B, D'_B) = D_B \setminus D'_B \cup D'_B \setminus D_B$. This is the symmetric difference between D_B and D'_B , and is the set of records that must be deleted and added to D_B to get D'_B . D_B and D'_B are neighbors w.r.t to $f(D_A, \cdot)$, denoted by $\mathcal{N}(f(D_A, \cdot))$ if

- (1) $f(D_A, D_B) = f(D_A, D'_B)$,
- (2) $\Delta(D_B, D'_B) \neq \emptyset$, and
- (3) there is no database $D''_B \in \mathcal{D}_B$, where $f(D_A, D_B) = f(D_A, D''_B)$, such that $\Delta(D_B, D'_B) \subset \Delta(D_B, D''_B)$.

$\mathcal{N}(f(\cdot, D_B))$ is similarly defined.

The third condition ensures that D_B and D'_B are minimally different in terms of record changes.

Definition 3.3 (Output Constrained DP). A 2-party PRL protocol Π for computing function $f : \mathcal{D}_A \times \mathcal{D}_B \rightarrow \mathcal{O}$ is $(\epsilon_A, \epsilon_B, \delta_A, \delta_B, f)$ -constrained differential privacy (DP) if for any $(D_B, D'_B) \in \mathcal{N}(f(D_A, \cdot))$, the views of Alice during the execution of Π to any probabilistic polynomial-time adversary T satisfies

$$\begin{aligned} & \Pr[T(\text{VIEW}_A^\Pi(D_A, D_B)) = 1] \\ & \leq e^{\epsilon_B} \Pr[T(\text{VIEW}_A^\Pi(D_A, D'_B)) = 1] + \delta_B \end{aligned} \quad (4)$$

and the same holds for the views of Bob with ϵ_A and δ_A .

If $\epsilon_A = \epsilon_B = \epsilon$, $\delta_A = \delta_B = \delta$, we simply denote it as (ϵ, δ, f) -constrained DP. Similar to DP, Output Constrained DP satisfies composition properties that are useful for protocol design.

THEOREM 3.4 (SEQUENTIAL COMPOSITION). Given Π_1 is $(\epsilon_1, \delta_1, f)$ -constrained DP, and Π_2 is $(\epsilon_2, \delta_2, f)$ -constrained DP, then applying these two protocols sequentially, i.e. $\Pi_2(D_A, D_B, \Pi_1(D_A, D_B))$ satisfies $(\epsilon_1 + \epsilon_2, \delta_1 + \delta_2, f)$ -constrained DP.

THEOREM 3.5 (POST-PROCESSING). Given Π is (ϵ, δ, f) -constrained DP, and let $O^\Pi(D_A, D_B)$ be the output after the execution of Π , then any probabilistic polynomial (in κ) function $g(O^\Pi(D_A, D_B))$ satisfies (ϵ, δ, f) -constrained DP.

See Appendix B.2.1 and B.2.2 for the proofs of Theorem 3.4 and Theorem 3.5 respectively. Output constrained DP inherits other desirable properties from DP, for instance, its robustness to attacks [17, 36]. We omit details due to space constraints.

3.2 Differential Privacy for Record Linkage

PRL can be a direct application of Output Constrained Differential Privacy by considering $f_{\bowtie_m}$. We have the following theorem to define the neighboring databases for PRL.

THEOREM 3.6 (NEIGHBORS FOR PRL). Given the function $f_{\bowtie_m}$ in PRL, if $(D_B, D'_B) \in \mathcal{N}(f_{\bowtie_m}(D_A, \cdot))$ for a given $D_A \in \mathcal{D}$, then $|D_B| = |D'_B|$, D_B and D'_B must differ in only one pair of non-matching records with respect to the given D_A , i.e. $D'_B = D_B - b + b'$ and $b \neq b'$, where $m(b, a) = 0$ and $m(b', a) = 0$ for all $a \in D_A$.

PROOF. The output constraint $f_{\bowtie_m}(D_A, D_B) = f_{\bowtie_m}(D_A, D'_B)$ implies that $|D_B| = |D'_B|$ and $D_A \bowtie_m D_B = D_A \bowtie_m D'_B$. If D_B and D'_B differ in a matching record, then their matching outputs with a given D_A are different. Hence D_B and D'_B must differ in one or more non-matching records. In addition, to ensure $|D_B| = |D'_B|$, the number of non-matching records added to D_B to get D'_B must be the same as the number of non-matching records removed from D_B . If $\Delta(D_B, D'_B)$ contains more than one pair of record additions and deletions, a subset of $\Delta(D_B, D'_B)$ can give a valid D''_B such that $f_{\bowtie_m}(D_A, D_B) = f_{\bowtie_m}(D_A, D''_B)$. Hence, a neighboring pair D_B, D'_B differ by exactly one pair of non-matching records. $\square$

Next we define the privacy guarantee that allows us to design efficient PRL protocols with provable privacy guarantees.

Definition 3.7 (DPRL). A 2-party PRL protocol Π for computing function $f_{\bowtie_m} : \mathcal{D}_A \times \mathcal{D}_B \rightarrow \mathcal{O}$ is $(\epsilon_A, \epsilon_B, \delta_A, \delta_B)$ -DPRL if Π satisfies $(\epsilon_A, \epsilon_B, \delta_A, \delta_B, f_{\bowtie_m})$ -constrained DP.

3.3 Related Privacy Definitions

In this section we discuss related privacy definitions and their connections with DPRL. First, both DPRL and IND-S2PC assume a computationally bounded model. We show that DPRL is a weaker guarantee than IND-S2PC.

THEOREM 3.8. All IND-S2PC protocols for record linkage satisfy $(0, \text{negl}(\kappa))$ -DPRL.

PROOF. IND-S2PC for record linkage is equivalent to DPRL with $\epsilon = 0$ and $\delta = \text{negl}(\kappa)$. The δ in DPRL is always greater than $\text{negl}(\kappa)$ but smaller than $o(1/n)$. $\square$

Hence, APC, PSI, and PSI+X techniques that satisfy IND-S2PC, guarantee $(0, \text{negl}(\kappa))$ -DPRL as well.

Indistinguishable computationally differential privacy (IND-CDP-2PC) [25] is another privacy notion under a computationally bounded model, and is a direct extension of DP to the two party setting where both parties are computationally bounded. DPRL is weaker than IND-CDP-2PC. Formally

THEOREM 3.9. If a protocol for record linkage satisfies $\epsilon/2$ -IND-CDP-2PC, then it satisfies (ϵ, δ) -DPRL.

The factor 2 arises since neighboring databases protected by DPRL have a symmetric difference of 2, while neighboring databases under IND-CDP-2PC have a symmetric difference of 1. The detailed proof can be found in Appendix A.2.

Blowfish Privacy [17] generalizes differential privacy to problems where constraints on the input database must hold (e.g., when certain query answers have been released by the database exactly). Output Constrained DP, including DPRL, is an extension of Blowfish in two ways: (1) from a computationally unbounded model to a computationally bounded model; (2) from a single-party setting to a two-party setting. Note that with the output constraint $f_{\bowtie_m}(D_A, D_B) = f_{\bowtie_m}(D_A, D'_B)$ for record linkage, the number of different records between neighboring databases D_B and D'_B is only two. This is not necessarily true for other applications of Output Constrained DP, or Blowfish Privacy. This property is desirable for DP based algorithms since larger distances between neighboring databases typically require larger perturbation to hide the difference between neighbors resulting in poorer utility.

Another instantiation of Blowfish privacy, called Protected DP [22], aims to ensure the privacy of a protected subpopulation. In contrast, an unprotected “targeted” subpopulation receives no privacy guarantees. In DPRL, one could think of the non-matching records as the protected subpopulation, and the matching records as the targeted subpopulation. However, unlike in Protected DP, in DPRL the set of protected records are learned as an output of the DPRL protocol, and hence are not available as an input to the protocol like the targeted subpopulation in the Protected DP algorithms.

4 PROTOCOLS FOR DPRL

In this section, we introduce protocols that satisfy DPRL and permit a 3-way trade-off between correctness, privacy and efficiency. We first present a class of protocols that achieves $(\epsilon, \text{negl}(\kappa))$ -DPRL by using a blocking strategy that satisfies local differential privacy (DP). Though these protocols achieve high recall with a sufficiently small privacy parameter, they only achieve a constant factor speedup in

efficiency. Next, we present the Laplace Protocol (LP) that achieves all three desiderata of high recall, privacy and subquadratic efficiency. This protocol hides non-matching records by adding Laplace noise to the blocking strategy. We also show that attempts from prior work to use Laplace noise in blocking fail to satisfy DPRL (Theorem 4.7). Moreover, we design a Sort & Prune (SP) heuristic that is used in conjunction with LP (as well as the local DP based protocols) and helps additionally tradeoff efficiency and recall. Finally, we present the Greedy Match & Clean heuristic optimization (GMC), that can further improve efficiency. All the protocols presented in this section are proven to satisfy DPRL.

4.1 Local DP Protocol

Let $\mathcal{B}$ be a blocking that randomly hashes records into a pre-specified set of k bins, such that for all $i \in [1 \dots k]$,

$$\Pr[\mathcal{B}(b) = i] \leq e^\epsilon \Pr[\mathcal{B}(b') = i]. \quad (5)$$

Such a blocking $\mathcal{B}$ satisfies ϵ -local DP (as defined in Appendix A.3), since each record is perturbed *locally* independent of the other records. We show that protocols that combine a local differentially private blocking with IND-S2PC protocols for record linkage can achieve $(\epsilon, \text{negl}(\kappa))$ -DPRL.

THEOREM 4.1. *All IND-S2PC protocols for record linkage with ϵ -local differentially private blocking satisfies $(\epsilon, \text{negl}(\kappa))$ -DPRL.*

The proof can be found in Appendix C.2. Such local differentially private protocols can be constructed from well known local differentially private algorithms based on randomized response (RR) [11] or the Johnson-Lindenstrauss (JL) transformation [4], where each record is hashed independent of others. We refer the reader to Appendix C.2 for a concrete blocking algorithm based on RR. We show that while this algorithm permits high recall and privacy, it does not improve efficiency by more than a constant factor (a function of ϵ) (Theorem C.1). Whether any local DP based blocking algorithms can achieve subquadratic efficiency is an interesting open question.

4.2 Laplace Protocol (LP)

4.2.1 Algorithm Description. In this protocol, Alice and Bob agree on a blocking function $\mathcal{B}$ with k bins and strategy $\mathcal{B}^S$, which we take as input to the protocol. The Laplace Protocol (LP, as shown in Algorithm 1) works by inserting a carefully chosen number of *dummy* records into each bin of the blocking strategy such that the bin sizes are differentially private. While candidate matches may contain dummy records, they do not contribute to the output set of matches, because the dummy records do not match *any* record. These candidate matches are then securely matched using an IND-S2PC algorithm.

In the first step (Lines 1-4) of the protocol shown in Algorithm 1, Alice and Bob take their inputs D_A and D_B , the agreed blocking protocol $\mathcal{B}$, and privacy parameters $\epsilon_A, \epsilon_B, \delta_A$, and δ_B as input, and compute *noisy* bins $\tilde{\mathcal{B}}(D_A)$ and $\tilde{\mathcal{B}}(D_B)$ respectively. The noisy bins are constructed as follows (Algorithm 2). Records in D are first hashed into bins according to the blocking protocol $\mathcal{B}$, and $\mathcal{B}(D)$ denotes the set of bins of records from D . Then the counts of the bins are perturbed using noise drawn from a truncated and discretized Laplace distribution, such that the noisy counts satisfy

Algorithm 1: Laplace Protocol (LP)

Input : $D_A, D_B, \epsilon_A, \epsilon_B, \delta_A, \delta_B, \mathcal{B}(\text{including } \mathcal{B}^S)$
Output : O

```

1 // Alice performs the following:
2  $\tilde{\mathcal{B}}(D_A) \leftarrow \text{LapNoise}(D_A, \mathcal{B}, \epsilon_A, \delta_A)$ ;
3 // Bob performs the following:
4  $\tilde{\mathcal{B}}(D_B) \leftarrow \text{LapNoise}(D_B, \mathcal{B}, \epsilon_B, \delta_B)$ ;
5 // Alice and Bob perform the following:
6  $O = \emptyset$ ;
7 // Sort & prune  $\mathcal{B}^S$  (§ 4.3)
8 for  $(i, j) \in \mathcal{B}^S$  do
9   for  $a \in \tilde{\mathcal{B}}_i(D_A)$  and  $b \in \tilde{\mathcal{B}}_j(D_B)$  do
10    | Add  $\text{SMC}(a, b)$  to  $O$ ;
11   end
12 // Greedy match & clean (§ 4.4)
13 end
14 return  $O$ ;
```

Algorithm 2: Add Laplace Noise

```

1 function  $\text{LapNoise}(D, \mathcal{B}, \epsilon, \delta)$ ;
2 for  $\mathcal{B}_i \in \mathcal{B}$  do
3    $\eta_i \sim \text{Lap}(\epsilon, \delta, \Delta \mathcal{B})$ ;
4    $\tilde{\mathcal{B}}_i(D) \leftarrow \text{add } \eta_i^+ = \max(\eta_i, 0)$  dummy records to  $\mathcal{B}_i(D)$ ;
5 end
6 return  $\tilde{\mathcal{B}}(D)$ ;
```

(ϵ, δ) -DPRL. The Laplace noise depends on not only the privacy parameters ϵ and δ , but also the sensitivity of the given blocking protocol $\mathcal{B}$.

Definition 4.2 (Sensitivity of $\mathcal{B}$). The sensitivity of the blocking strategy $\mathcal{B}$ for Bob, denoted by $\Delta \mathcal{B}_B$ is

$$\max_{D_A \in \mathcal{D}} \max_{(D_B, D'_B) \in \mathcal{N}(f_{\bowtie_m}(D_A, \cdot))} \sum_{i=0}^k ||\mathcal{B}_i(D_B) - \mathcal{B}_i(D'_B)||,$$

the maximum bin count difference between D_B and D'_B for any $(D_B, D'_B) \in \mathcal{N}(f_{\bowtie_m}(D_A, \cdot))$ for all $D_A \in \mathcal{D}$. $\Delta \mathcal{B}_A$ for Alice is similarly defined.

If the hashing of $\mathcal{B}$ is the same for Alice and Bob, then $\Delta \mathcal{B}_A = \Delta \mathcal{B}_B = \Delta \mathcal{B}$. We assume this in our paper. If $\mathcal{B}$ hashes each record to at most k' bins, then $\Delta \mathcal{B} = 2k'$.

Definition 4.3 ($\text{Lap}(\epsilon, \delta, \Delta \mathcal{B})$). A random variable follows the $\text{Lap}(\epsilon, \delta, \Delta \mathcal{B})$ distribution if it has a probability density function

$$\Pr[\eta = x] = p \cdot e^{-(\epsilon/\Delta \mathcal{B})|x - \eta^0|}, \quad \forall x \in \mathbb{Z}, \quad (6)$$

where $p = \frac{e^{\epsilon/\Delta \mathcal{B}} - 1}{e^{\epsilon/\Delta \mathcal{B}} + 1}$, and $\eta^0 = -\frac{\Delta \mathcal{B} \ln((e^{\epsilon/\Delta \mathcal{B}} + 1)(1 - (1 - \delta)^{1/\Delta \mathcal{B}}))}{\epsilon}$.

This distribution has a mean of η_0 and takes both positive and negative values. LP draws a noise value η from this distribution, and truncates it to 0 if η is negative. Then, η dummy records are added to the bin. These dummy records lie in an expanded domain, such that they do not match with any records in the true domain.

After Alice and Bob perturb their binned records, they will initiate secure matching steps to compare candidate matches, i.e. records

in $\tilde{\mathcal{B}}_i(D_A) \times \tilde{\mathcal{B}}_j(D_B)$ if $(i, j) \in \mathcal{B}^S$. For each candidate match (a, b) , Alice and Bob participate in a two party secure matching protocol $SMC(a, b)$ that outputs the pair (a, b) to both Alice and Bob if $m(a, b) = 1$ (true matching pair) and null otherwise. Secure matching can be implemented either using garbled circuits [40] or (partially) homomorphic encryption [28], depending on the matching rule (see Appendix C.1 for an example).

4.2.2 Correctness Analysis. Compared to the original non-private blocking protocol $\mathcal{B}$, no records are deleted, and dummy records do not match any real record. Hence,

THEOREM 4.4. *Algorithm 1 gives the same recall as the non-private blocking protocol $\mathcal{B}$ it takes as input.*

4.2.3 Privacy Analysis. Next, we show that LP satisfies DPRL.

THEOREM 4.5. *Algorithm 1 satisfies $(\epsilon_A, \epsilon_B, \delta_A, \delta_B)$ -DPRL.*

PROOF. We prove privacy for Bob (the proof for Alice is analogous). In this protocol, Alice with input data D_A has a view consisting of (1) the number of candidate matching pairs arising in each $(i, j) \in \mathcal{B}^S$, (2) the output for each candidate matching pair. Algorithm 1 is the composition of two steps: (a) add dummy records to bins, and (b) secure comparison of records within bins.

Consider a neighboring pair $(D_B, D'_B) \in N(f_{\bowtie_m}(D_A, \cdot))$ for a given D_A . By Theorem 3.6, D_B and D'_B differ in only one non-matching record with respect to D_A , i.e. $D'_B = D_B - b_* + b'_*$ and $b_* \neq b'_*$, where $m(b_*, a) = 0$ and $m(b'_*, a) = 0$ for all $a \in D_A$. D_B and D'_B can differ by at most $\Delta\mathcal{B}$ in their bin counts. We show in Lemma B.2 (Appendix) that Algorithm 2 adds a sufficient number of dummy records to hide this difference: with probability $1 - \delta_B$, the probabilities of generating the same noisy bin counts for Bob, and hence the same number of candidate matching pairs consisting in each $(i, j) \in \mathcal{B}^S$ from D_B and D'_B are bounded by e^{ϵ_B} . Thus, Step (a) ensures (ϵ_B, δ_B) -DPRL for Bob. Given a fixed view from Step (a) which consists of the noisy bin counts and encrypted records from $\tilde{\mathcal{B}}(D_B)$, Alice's view regarding the output for each candidate matching pair (a, b) is the same. The encrypted records for a given noisy bin counts can only differ in b_* and b'_* , but both of them lead to the same output for each candidate matching, because they do not match any records in D_A . Each secure pairwise comparison satisfies $(0, \text{negl}(\kappa))$ -DPRL, and since there are at most n^2 comparisons (recall $\kappa > n = \max(|D_A|, |D_B|)$). Thus Step (b) satisfies $(0, \text{negl}(\kappa))$ -DPRL.

Therefore, using similar arguments for Alice and sequential composition, we get that Algorithm 1 satisfies DPRL. $\square$

THEOREM 4.6. *If Algorithm 1 (LP) takes $\eta_0 = \ln^2 n \cdot \Delta\mathcal{B}/\epsilon$ for Eqn. (6), then LP satisfies $(\epsilon_A, \epsilon_B, o(1/n^k), o(1/n^k))$ -DPRL, for any $k > 0$, where $n = \max(|D_A|, |D_B|)$.*

PROOF. (sketch) Taking $\eta_0 = \ln^2 n \cdot \Delta\mathcal{B}/\epsilon$, the failing probability $\delta = 1 - (1 - \frac{1}{n \ln n (e^{\epsilon/\Delta\mathcal{B}} + 1)})^{\Delta\mathcal{B}} \leq \frac{c}{n \ln n}$ for some constant c (in terms of $\epsilon, \Delta\mathcal{B}$). Hence $\delta = o(1/n^k)$ for all $k > 0$. $\square$

LP only adds non-negative noise to the bin counts. One could instead add noise that could take positive and negative values, and suppress records if the noise is negative. We call this protocol LP-2. This is indeed the protocol proposed by prior work [5, 19, 24]

that combined APC with DP blocking. However, we show that this minor change in LP results in the protocol violating DPRL (even though the noise addition seems to satisfy DP)! Hence, LP-2 also does not satisfy IND-CDP-2PC (by Theorem 3.9).

THEOREM 4.7. *For every non-negative $\epsilon, \delta < \frac{p^{\Delta\mathcal{B}}}{2e^{\epsilon}}$, there exists a pair of neighboring databases for which LP-2 does not ensure (ϵ, δ) -DPRL, where $p = \frac{e^{\epsilon/\Delta\mathcal{B}} - 1}{e^{\epsilon/\Delta\mathcal{B}} + 1}$.*

PROOF. (sketch) The output of the record suppression step is dependent on the ratio between the matching and non-matching records in the bin. This introduces a correlation between the matching and non-matching records. Consider a neighboring pair D_B and D'_B that differ by a non-matching pair (b_*, b'_*) for a given D_A . If b_* is in a bin full of non-matching records with D_A , and b'_* is in a bin full of matching records with D_A (except b'_*). D_B is more likely to output all matching pairs than D'_B if some record is suppressed. The detailed proof can be found in Appendix B.1.2. $\square$

4.2.4 Efficiency Analysis. Last, we present our result on the efficiency of LP. Note that the communication and computational costs for LP are the same as $O(\text{cost}_{\mathcal{B}^S})$, where $\text{cost}_{\mathcal{B}^S}$ is the number of candidate matches, if you consider the communication and computational costs associated with a single secure comparison as a constant. Hence, we analyze efficiency in terms of the number of candidate pairs $\text{cost}_{\mathcal{B}^S}$ in LP.

THEOREM 4.8. *Given a blocking protocol $\mathcal{B}$ with k bins and blocking strategy $\mathcal{B}^S$, such that the number of candidate matches for D_A and D_B , $\text{cost}_{\mathcal{B}^S}(D_A, D_B)$, is sub-quadratic in n , i.e. $o(n^2)$, where $n = \max(|D_A|, |D_B|)$. If (1) the number of bins k is $o(n^c)$ for $c < 2$, and (2) each bin of a party is compared with $O(1)$ number of bins from the opposite party, then the expected number of candidate matches in Algorithm 1 is sub-quadratic in n .*

PROOF. Given ϵ and δ , the expected number of dummy records added per bin $\mathbb{E}(\eta^+)$ is a constant denoted by c_η (Def. 4.3). Each bin of a party is compared with at most c_b bins from the opposite party, where c_b is a constant. The number of candidate matches in LP is a random variable, denoted by COST , with expected value

$$\begin{aligned} \mathbb{E}(\text{COST}) &= \sum_{(i,j) \in \mathcal{B}^S} \mathbb{E}(|\tilde{\mathcal{B}}_i(D_A)| |\tilde{\mathcal{B}}_j(D_B)|) \\ &= \sum_{(i,j) \in \mathcal{B}^S} |\mathcal{B}_i(D_A)| |\mathcal{B}_j(D_B)| + \sum_{(i,j) \in \mathcal{B}^S} \mathbb{E}(\eta_i^+) \mathbb{E}(\eta_j^+) \\ &\quad + \sum_{(i,j) \in \mathcal{B}^S} (\mathbb{E}(\eta_i^+) |\mathcal{B}_j(D_B)| + \mathbb{E}(\eta_j^+) |\mathcal{B}_i(D_A)|) \\ &< \text{cost}_{\mathcal{B}^S}(D_A, D_B) + c_\eta^2 c_b k + 2c_\eta c_b n. \end{aligned}$$

Since $\text{cost}_{\mathcal{B}^S}(D_A, D_B)$ and k are sub-quadratic in n , $\mathbb{E}(\text{COST})$ is also sub-quadratic in n . When δ is a negligible term as defined in Theorem 4.6, the noise per bin is $O(\ln^2 n)$. As k is $o(n^c)$ for $c < 2$, the expected value of COST is still sub-quadratic in n . $\square$

Conditions (1) and (2) in the above theorem are satisfied by, for instance, sorted neighborhood, and distance based blocking [6] (we use the latter in our experiments). While the asymptotic complexity of LP is sub-quadratic, it performs at least a constant number of secure comparisons for each pair $(i, j) \in \mathcal{B}^S$ even if there are no real

records in $\mathcal{B}_i(D_A)$ and $\mathcal{B}_j(D_B)$. We can reduce this computational overhead with a slight loss in recall (with no loss in privacy) using a heuristic we describe in the next section.

4.3 Sort & Prune $\mathcal{B}^S$ (SP)

Algorithm 1 draws noise from the same distribution for each bin, and hence the expected number of dummy records is the same for every bin. The bins with higher noisy counts will then have a higher ratio of true to dummy records. This motivates us to match candidate pairs in bins with high noisy counts first. Instead of comparing bin pairs in $\mathcal{B}^S$ in a random or index order, we would like to sort them based on the noisy counts of $\tilde{\mathcal{B}}(D_A)$ and $\tilde{\mathcal{B}}(D_B)$. Given a list of descending thresholds $\bar{t} = [t_1, t_2, t_3 \dots]$, the pairs of bins from the matching strategy $\mathcal{B}^S$ can be sorted into groups denoted by $\mathcal{B}^{S, t_l}$ for $l = 1, 2, \dots$, where

$$\mathcal{B}^{S, t_l} = \{(\tilde{\mathcal{B}}_i(D_A) > t_l \wedge |\tilde{\mathcal{B}}_j(D_B)| > t_l) | (i, j) \in \mathcal{B}^S\}.$$

Each group consists of bin pairs from $\mathcal{B}^S$ with both noisy counts greater than the threshold.

We let the thresholds $\bar{t}$ be the deciles of the sorted noisy bin sizes of $\tilde{\mathcal{B}}(D_A)$ and $\tilde{\mathcal{B}}(D_B)$. As the threshold decreases, the likelihood of matching true records instead of dummy records drops for bins. Alice and Bob can stop this matching process before reaching the smallest threshold in $\bar{t}$. If the protocol stops at a larger threshold, the recall is smaller. In the evaluation, if the protocol stops at 10% percentile of the noisy bin counts, the recall can reach more than 0.95. This allows a trade-off between recall and efficiency for a given privacy guarantee. We show that this step also ensures DPRL.

COROLLARY 4.9. *Algorithm 1 with sort & prune step (SP) satisfies $(\epsilon_A, \epsilon_B, \delta_A, \delta_B)$ -DPRL.*

PROOF. Similar to the proof in Theorem 4.5, Alice with input data D_A has a view consisting of (1) the number of candidate matching pairs arising in each $(i, j) \in \mathcal{B}^S$, and (2) the output for each candidate matching pair. As SP is a post-processing step based on the noisy bin counts, which is part of Alice's original view, the overall protocol still satisfies the same DPRL guarantee by Theorem 3.5 (post-processing). $\square$

We next present an optimization that also uses a form of post-processing to significantly reduce the number of secure pairwise comparisons in practice, but whose privacy analysis is more involved than that of SP.

4.4 Greedy Match & Clean (GMC)

LP executes a sequence of secure comparison protocols, one per candidate pair. After every comparison (or a block of comparisons), Alice and Bob learn a subset of the matches O . Based on the current output O , Alice and Bob can greedily search matching pairs in the clear from their respective databases (Lines 5,10 in Algorithm 3), and add the new matching pairs to the output set O until no new matching pairs can be found. In addition, Alice and Bob can remove records in the output from the bins $\tilde{\mathcal{B}}(D_A)$ and $\tilde{\mathcal{B}}(D_B)$ to further reduce the number of secure pairwise comparisons (Lines 4,9). We can see that this optimization step is not simply post-processing, because it makes use of the true record in plain text for matching. In traditional differential privacy, when the true data is used for

Algorithm 3: Greedy match and clean

Input: $O, \tilde{\mathcal{B}}(D_A), \tilde{\mathcal{B}}(D_B)$

```

1 repeat
2   // Alice performs the following:
3    $O_A \leftarrow \pi_A O, O_B \leftarrow \pi_B O$ ;
4    $\tilde{\mathcal{B}}(D_A) \leftarrow \tilde{\mathcal{B}}(D_A) - O_A$ ;
5    $O' \leftarrow \text{PlainMatch}(O_B, \tilde{\mathcal{B}}(D_A))$ ;
6   Add  $O'$  to  $O$  and send  $O$  to Bob;
7   // Bob performs the following:
8    $O_A \leftarrow \pi_A O, O_B \leftarrow \pi_B O$ ;
9    $\tilde{\mathcal{B}}(D_B) \leftarrow \tilde{\mathcal{B}}(D_B) - O_B$ ;
10   $O' \leftarrow \text{PlainMatch}(O_A, \tilde{\mathcal{B}}(D_B))$ ;
11  Add  $O'$  to  $O$  and send  $O$  to Alice;
12 until  $O$  received by Alice has no updates;
```

computation, the privacy guarantee decays. However, we show that this is not true for the GMC step in the setting of DPRL.

THEOREM 4.10. *Algorithm 1 with the greedy match & clean step (GMC) in Algorithm 3 satisfies $(\epsilon_A, \epsilon_B, \delta_A, \delta_B)$ -DPRL.*

PROOF. First consider the privacy for Bob. Alice with input data D_A , has a view consisting of (1) the number of candidate matching pairs arising in each $(i, j) \in \mathcal{B}^S$, (2) the output for each candidate matching pair, (3) the output from plaintext comparisons with output records.

Consider a neighboring pair $(D_B, D'_B) \in N(f_{\times_m}(D_A, \cdot))$ for a given D_A . By Theorem 3.6, D_B and D'_B differ in only one non-matching record with respect to D_A , i.e. $D'_B = D_B - b_* + b'_*$ and $b_* \neq b'_*$, where $m(b_*, a) = 0$ and $m(b'_*, a) = 0$ for all $a \in D_A$. D_B and D'_B can differ by at most $\Delta \mathcal{B}$ in their bin counts. Similar to the proof for Theorem 4.5, the first step of the protocol adds dummy records to bins, and satisfies (ϵ_B, δ_B) -DPRL.

In the second step, given a fixed view VIEW^* from the first step which consists of the noisy bin counts and encrypted records from $\tilde{\mathcal{B}}(D_B)$, Alice's view regarding the output for each candidate matching pair (a, b) is the same regardless (a, b) are compared securely or in plaintext. Alice's view regarding the output from plaintext comparisons with the records in the output set is also the same for a fixed VIEW^* from the first step. The encrypted records for a given noisy bin counts can only differ in b_* and b'_* , and they will never be pruned away. Both of them also lead to the same output for secure pairwise comparisons or plaintext comparisons, because they do not match any records in D_A . Thus Step (b) satisfies $(0, \text{negl}(\kappa))$ -DPRL.

Therefore, using similar arguments for Alice and sequential composition, we get that Algorithm 1 satisfies DPRL. $\square$

With the same privacy guarantee, LP with the GMC step can even improve the efficiency of LP without sacrificing recall.

THEOREM 4.11. *LP with the greedy match & clean step (GMC) performs no more secure pairwise comparisons than LP, and outputs at least as many matching pairs as LP.*

We refer the reader to Appendix B.3.2 for the proof. Both SP and GMC are also applicable on the local DP based protocols for the

similar reasoning. Hence, we will only show how each optimization helps improve the efficiency of the basic LP in the evaluation.

5 EVALUATION

We empirically evaluate the correctness, privacy, and efficiency of the protocols proposed in § 4. Our experiments demonstrate the following results:

- The Laplace Protocol (LP, which includes all the optimizations) proposed in § 4 is over 2 orders of magnitude more efficient than the baseline approaches while still achieving a high recall and end-to-end privacy. (§ 5.2.1)
- At any given level of privacy, LP incurs a computational cost that is near linear in the input database size. (§ 5.2.1)
- Greedy match & clean and Sort & prune optimization help reduce communication and computation costs. The former results in 50% lower cost than unoptimized LP in some cases. (§ 5.2.2)
- We explore the 3-way trade-offs between correctness, privacy, and efficiency of LP. (§ 5.2.3)

5.1 Evaluation Setup

5.1.1 Datasets and Matching Rules. *Taxi dataset (Taxi):* To simulate linkage in the location domain, we extract location distribution information from the TLC Trip Record Data [37]. Each record includes a pickup location in latitude-longitude coordinates (truncated to 6 decimal places) and the date and hour of the pickup time. Taking the original dataset as D_A , we create D_B by perturbing the latitude-longitude coordinates of each record in D_A with random values uniformly drawn from $[-\theta, +\theta]^2$, where $\theta = 0.001$. Each day has approximately 300,000 pickups. The data size can be scaled up by increasing the number of days, T . We experiment with $T = 1, 2, 4, 8, 16$, with $T = 1$ being the default. Any pair of records $a, b \in \Sigma$ are called a match if they have the same day and hour, and their Euclidean distance in location is no larger than θ . The location domain is within the bounding box (40.711720N, 73.929670W) and (40.786770N, 74.006600W). We project the locations into a uniform grid of 16×16 cells with size 0.005×0.005 . A blocking strategy $\mathcal{B}^S$ based on the pickup time and grid is applied to both datasets, resulting in $(16 \times 16 \times 24T)$ bins. $\mathcal{B}^S$ compares pairs of bins that are associated with the same hour, and corresponding/neighborhood grid cells. Thus, each bin in $\mathcal{B}(D_A)$ is compared with 9 bins in $\mathcal{B}(D_B)$.

Abt and Buy product dataset (AB): These datasets are synthesized from the online retailers Abt.com and Buy.com [23] who would like to collaboratively study the common products they sell as a function of time. Each record in either dataset consists of a product name, brand and the day the product was sold. The product names are tokenized into trigrams, and hashed into a bit vector with a bloom filter having domain $\Sigma = \{0, 1\}^{50}$. We consider 16 brands, and sample 5,000 records per day from the original datasets for Abt and Buy each. The data size can be scaled up with T for $T = 1, 2, 4, 8, 16$, with 1 being the default for T . Any pair of records $a, b \in \Sigma$ are called a match if (a) they are sold on the same day, (b) they are of the same brand, and (c) the hamming distance between their vectorized names is no more than $\theta = 5$. A blocking strategy hashes records having the same value for day and brand into the same bin, resulting in $16T$ bins, and compares records falling in the corresponding bins.

Figure 2: The average log(cost) of LP, APC, PSI+X and non-private matching (np) for the Taxi and AB datasets vs log(data size). LP give lower costs than the baselines PSI+X and APC for all values of $\epsilon = 0.1, 0.4, 1.6$ and $\delta = 10^{-5}$, and scales near linearly.

5.1.2 Protocols: We evaluate four DPRL protocols: (1) Laplace protocol (LP), (2) all-pairwise comparisons (APC), (3) private set intersection (PSI), and (4) PSI with expansion (PSI+X). The default LP consists of the basic protocol described in Algorithm 1 along with optimization steps (SP and GMC) in § 4.3 and 4.4.

5.1.3 Metrics: There are three dimensions in the trade-off space: correctness, privacy and efficiency. The correctness of a protocol is measured by the *recall*, which is the fraction of the matching pairs output by the algorithm, as defined in Eqn. (1), with larger values close to 1 being better. The *privacy* metric is specified in advance for each algorithm using parameters ϵ, δ . For AP, PSI, and PSI+X, $\epsilon = 0$ and $\delta = \text{negl}(\kappa)$ by Theorem 3.8. We consider $\epsilon_A = \epsilon_B = \epsilon$ and $\delta_A = \delta_B$ for $\epsilon \in \{0.1, 0.4, 1.6\}$ and $\delta \in \{10^{-9}, 10^{-7}, 10^{-5}\}$ for LP. The default value for ϵ and δ is 1.6 and 10^{-5} , respectively. Finally, we define *efficiency* of APC and LP protocols for a given dataset as the number of secure pairwise comparisons, and denote this by *cost*. The cost of PSI and PSI+X can be estimated as $\gamma n \ln \ln(n)$, where γ is the expansion factor, or the ratio of sizes of the expanded and true databases. This represents the number of operations on encrypted values. For PSI, γ is 1. We use the number of secure comparison/operations on encrypted values rather than the wallclock times as a measure of efficiency, since these operations dominate the total time. We discuss wallclock times in more detail in § 5.2.4.

5.2 Results and Discussions

5.2.1 Efficiency and scalability. In this section, we empirically investigate how LP scales as the data size increases ($T \in \{1, 2, 4, 8, 16\}$) in comparison to baselines APC and PSI+X, when all the algorithms achieve 100% recall. We do not include PSI as its recall is close to 10%. LP is evaluated at privacy parameter $\epsilon \in \{1.6, 0.4, 0.1\}$ and fixed $\delta = 10^{-5}$. At each ϵ , we report the average number of candidate pairs for LP over 10 runs for each value of T . To achieve 100% recall, PSI+X expands each record b in D_B to every other record b' within a θ -ball around b . We add 2,369,936 records per record in the AB dataset, and $1000^2\pi$ records per record in the Taxi dataset.

In Figure 2, we report the log(base 10) value of the average cost, $\log(\text{cost})$, with respect to the log value of data size $\log(n)$ for PSI+X,

Figure 3: The average cost with standard deviation of LP protocols with five settings: (1) ‘basic’ – the basic LP protocol in Algorithm 1, (2) ‘basic+SP’ – the basic LP with sort & prune step in § 4.3, (3) ‘basic+GMC’ – the basic protocol with greedy match & clean step in Algorithm 3, (4) ‘basic+GMC+S’ – the previous protocol with addition sorting step described in § 4.3, and (5) ‘basic+GMC+SP’ – the protocol stops at recall less than full recall.

APC, and LP with varying ϵ and the non-private setting (np) when they achieve a recall of 1.0. Results for Taxi are shown on the left, and AB are shown on the right. For both datasets, the baseline methods, PSI+X and APC, have data points and line segments above LP for the plotted data size range. When the Taxi dataset has a size of $10^{5.5}$, LP at $\epsilon = 0.1$ costs an order of magnitude less than APC, as shown by the leftmost brown point (APC) and blue point (LP, $\epsilon=0.1$) in Figure 2(left). As the data size increases, the gap between APC and LP gets larger. When data size increases by 16 times (the right most points in the plots), LP at $\epsilon = 0.1$ costs over 2 orders of magnitude less than APC. When ϵ increases, the cost of LP shifts downward towards the non-private setting (np). When $\epsilon = 1.6$, LP has 3 orders of magnitude lower cost than APC for the given range of data sizes. The line for np is the lower bound for LP, where no dummy records are added to the bins. Similar observations are found in Figure 2(right) for the AB dataset, where LP improves APC by up to 2 orders for the plotted data size range.

PSI+X has a much larger cost than both APC and LP, mainly due to the fact that the expansion factor is far larger than the data size. We also observe that the lines that pass through the points of APC for both Taxi and AB datasets have a slope of 2, which corresponds to the quadratic communication and computational cost of APC. LP and PSI+X have slopes of values slightly larger than 1, and thus are linear time. Thus, for sufficiently large data sizes, PSI+X can beat APC. However, we do not expect PSI+X to beat LP due to the large expansion factor. Similar results are observed when the protocol stops before achieving full recall (Figure 5 in Appendix D).

5.2.2 Optimization steps. We next study the effectiveness of the optimization steps for LP. We study 5 protocols as shown below:

- ‘basic’: the basic LP Algorithm 1 with no heuristic optimizations;
- ‘basic+SP’: the basic LP with the sort & prune step (SP). SP stops the protocol when the threshold reaches the 10% percentile of the noisy bin counts of $\hat{\mathcal{B}}(D_A)$ and $\hat{\mathcal{B}}(D_B)$. Together with the sorting step, bins pairs with insufficient counts can be pruned away, resulting in a recall slightly smaller than the highest possible recall;

- ‘basic+GMC’: the greedy match & clean step (GMC) in Algorithm 3 is applied to the basic LP;
- ‘basic+GMC+S’: in addition to the previous protocol, bins are sorted in order of size. Pruning is omitted so that the highest possible recall is achieved;
- ‘basic+GMC+SP’: the same protocol as ‘basic+GMC+S’, except it prunes the bins with counts in the bottom 10% percentile.

Hence, the default LP can be also denoted by ‘basic+GMC+S’ if recall is 1.0 and ‘basic+GMC+SP’ if recall is less than 1.0.

In Figure 3, we report the average cost with the standard deviation across 10 runs of the above mentioned protocols at $\epsilon = 0.1, 0.4, 1.6$ and $\delta = 10^{-5}$ for the Taxi and AB datasets when $T = 1$ and $T = 16$. Several interesting observations arise from this plot.

First, the most significant drop in cost is due to GMC. The protocols with the greedy step have smaller cost than other protocols for all ϵ and datasets. For the Taxi datasets at $T = 1$ or $T = 16$, ‘basic+GMC’ saves the cost of ‘basic’ by over 50% when $\epsilon = 1.6$. As ϵ decreases, these relative savings reduce because more dummy records are added and cannot be matched or removed by this greedy step. For the AB datasets, ‘basic+GMC’ reduces the cost of ‘basic’ by up to 16% at $\epsilon = 1.6$ and 11% at $\epsilon = 0.1$.

Next, adding the sorting step to GMC (GMC+S) improves upon GMC when the data sizes are large ($T=16$). For instance, when $\epsilon = 0.1$ and $T = 16$, ‘basic+GMC+S’ can further bring the cost down by approximately 8.0×10^6 candidate pairs for the AB datasets, and by 2.0×10^9 for the Taxi datasets.

Third, the cost of ‘basic+GMC+SP’ is reported at a recall reaching above 0.95. The reduction with respect to ‘basic+GMC+S’ is relatively small, but the absolute reduction in cost is significant in some setting. For instance, the number of candidate pairs is reduced by 5.0×10^6 for the AB datasets when $\epsilon = 0.1$ and $T = 16$.

Last, for the AB dataset at $T = 16$, ‘basic+SP’ has a smaller variance in cost than ‘basic’ at $\epsilon = 0.1$. Similarly, ‘basic+GMC+SP’ has a smaller variance in cost than ‘basic+GMC’. This implies the sort & prune step can help prune away bins, and hence reduce the variance introduced by dummy records.

Figure 4: LP with varying privacy settings plotted over the default Taxi datasets and AB datasets. Each trade-off line between recall and the efficiency ($\text{cost(LP)}/\text{cost(APC)}$) corresponds to the default LP at a privacy setting (ϵ, δ) . Figure 4(a) varies ϵ and Figure 4(b) varies δ .

5.2.3 Three-way trade-offs. All the DPRL baseline methods including APC, PSI and PSI+X, have a fixed and strong privacy guarantee where $\epsilon = 0$ and $\delta = \text{negl}(\kappa)$. Hence, each baseline has a single point in a plot between recall and efficiency for a given data size, where APC and PSI+X have a point with full recall and high cost, and PSI has a point with low recall and low cost. Here, we will show that LP allows a trade-off between recall and efficiency for a given privacy guarantee. The efficiency metric used here is the ratio of the cost(LP) to the cost(APC) .

Figure 4(a) illustrates the case when both Alice and Bob require (ϵ, δ) -DPRL protection where $\epsilon = \{0.1, 0.4, 1.6\}$ and fixed $\delta = 10^{-5}$. In Figure 4(b), we vary the values of δ for $\delta \in \{10^{-9}, 10^{-7}, 10^{-5}\}$ with fixed $\epsilon = 1.6$. Each data point in the plot corresponds to the average $\text{cost(LP)}/\text{cost(APC)}$ and average recall of the default LP for a given (ϵ, δ) and the default data size with $T = 1$. The default LP allows the sort & prune step as described in § 4.3 with a list of thresholds that are the 90%, 80%, ..., 0% percentiles of the sorted bin sizes of $\mathcal{B}(D_A)$ and $\mathcal{B}(D_B)$. We report the average recall and $\text{cost(LP)}/\text{cost(APC)}$ for each percentile. This gives a trade-off line for each ϵ and δ value.

We observe that all the trade-off lines obtain a high recall at very small values of $\text{cost(LP)}/\text{cost(APC)}$. Even at $\epsilon = 0.1$, LP incurs 100 times smaller cost than APC. LP has a slightly larger cost for AB dataset. In Figure 4(a), the trade-off lines between recall and efficiency shift rightwards as the privacy parameter ϵ gets smaller. In other words, the cost is higher for a stronger privacy guarantee in order to output the same recall. Similar observations are found in Figure 4(b). However, the trade-off lines are more sensitive to ϵ than δ . The red lines in Figure 4(a) and the red lines in Figure 4(b) correspond to the same privacy setting. As δ reduces by 10000 times from 10^{-5} to 10^{-9} , the trade-off line of LP for the Taxi datasets shifts the ratio of costs by at most 0.001 as shown in Figure 4(b) (left) while the trade-off line increases the ratio of costs to 0.07 as ϵ reduces from 1.6 to 0.1 (Figure 4(a)).

As the Taxi and AB dataset have different data distributions over bins, the shapes of the trade-off lines are different. AB datasets are more skewed and have some bins with large counts. These bins also have many matching pairs, and hence we see a steep rise for

the first part of the trade-off lines for the AB datasets. When the data size increases, if the distribution of matching pairs remains similar, the trade-off lines between the efficiency and recall tends to stay the same. These trade-off lines can be useful when choosing the recall, privacy and efficiency for larger datasets.

5.2.4 Wall clock times. We implemented APC and LP in python, and implemented operations on encrypted records using the Paillier homomorphic cryptosystem using the python-paillier library [38]. As all algorithms require a one-time encryption of records we exclude this cost and only measure the cost of operations on the encrypted records. On a 3.1 GHz Intel Core i7 machine with 16 GB RAM, we found that computing the Hamming distance of two encrypted records with dimension $d = 50$ takes an average of $t_s = 77$ ms. That is, for datasets of size $n = 5000$, APC would take over 22 days to complete! Additionally, for the same dataset with $\epsilon = 1.6$, LP would only take 80 hours to achieve a recall of 1. In comparison, the wall clock time of LP ignoring the time spent in comparisons of encrypted records was only 120 seconds. We believe that this order of magnitude difference in time for secure operations and normal operations is true independent of the library or protocol used for secure comparisons. Thus, the computational cost of LP is dominated by the cost of secure comparison. How to improve the unit cost of each secure pairwise comparison is an important research topic, and is orthogonal to our research. Hence, in this evaluation, we focused only on the number of secure comparisons/operations on encrypted values to measure efficiency.

6 RELATED WORK

In addition to the prior work [5, 19, 24] that attempted to combine DP and secure computation techniques in order to scale-up the PRL problem, there are other efforts that take similar approaches, but focus on solving different problems. Wagh et al. [39] formalized the notion of *differentially private oblivious RAM (DP ORAM)* and their corresponding protocols significantly improved the bandwidth overheads with a relaxed privacy guarantee. This privacy notion considers a client-server model where all data sit on a single server,

while DPRL considers two party computation. Moreover, the protocols for DP ORAM only consider the trade-off between privacy and efficiency while DPRL considers an additional trade-off dimension: correctness. Several efforts [2, 16, 26, 27, 30, 31] also integrated DP with SMC in a distributed setting where data is vertically or horizontally partitioned between parties. The difference is that these papers focus on aggregate functions over the partitioned data, such as join size, marginal counts and sum, while PRL requires matching individual record pairs. This matching of individual record pairs does not naturally compose with DP, and hence motivated DPRL, a new privacy model for efficient PRL.

7 CONCLUSION

In this work, we propose a novel privacy model, called *output constrained differential privacy*, that shares the strong privacy protection of differential privacy, but allows for the truthful release of the output of a certain function on the data. We showed that this new privacy model can be applied to record linkage to define differential privacy for record linkage (DPRL). Under this framework, we proposed novel protocols for efficient PRL that satisfy three desiderata: correctness, privacy and efficiency. This is an important advance, since none of the prior techniques achieves all three desiderata. Despite this advance, further investigation into the practicality of DPRL protocols is a direction for future research. This includes investigation into their wall clock times in a specific operational environment and over datasets with more complex matching functions. Additional directions for future research include identifying DPRL protocols that further reduce the computational complexity of record linkage, such as applying a data-dependent blocking strategy, extending two-party DPRL to a multi-party setting, and generalizing the notion of output constrained differential privacy to other applications beyond private record linkage.

Acknowledgements: This work was supported by NSF grant 1253327, 1408982, 1443014, and DARPA & SPAWAR under contract N66001-15-C-4067. The U.S. Government is authorized to reproduce and distribute reprints for Governmental purposes not withstanding any copyright notation thereon. The views, opinions, and/or findings expressed are those of the author(s) and should not be interpreted as representing the official views or policies of the Department of Defense or the U.S. Government.

REFERENCES

- [1] Ali Al-Lawati, Dongwon Lee, and Patrick McDaniel. 2005. Blocking-aware Private Record Linkage. In *IQIS*.
- [2] Dima Alhadidi, Noman Mohammed, Benjamin C. M. Fung, and Mourad Debbabi. 2012. Secure Distributed Framework for Achieving ϵ -differential Privacy. In *PETS*.
- [3] Mikhail J. Atallah, Florian Kerschbaum, and Wenliang Du. 2003. Secure and Private Sequence Comparisons. In *WPES*.
- [4] Jeremiah Blocki, Avrim Blum, Anupam Datta, and Or Sheffet. 2012. The Johnson-Lindenstrauss Transform Itself Preserves Differential Privacy. In *FOCS*.
- [5] Jianneng Cao, Fang-Yu Rao, Elisa Bertino, and Murat Kantarcioglu. 2015. A hybrid private record linkage scheme: Separating differentially private synopses from matching records. In *ICDE*.
- [6] Peter Christen. 2012. *Data Matching: Concepts and Techniques for Record Linkage, Entity Resolution, and Duplicate Detection*. Springer Publishing Company, Incorporated.
- [7] Tim Churches and Peter Christen. 2004. Blind data linkage using n-gram similarity comparisons. In *KDD*.
- [8] Tim Churches and Peter Christen. 2004. Some methods for blindfolded record linkage. *BMC Medical Informatics and Decision Making* 4, 1 (2004), 1.
- [9] Xin Luna Dong and Divesh Srivastava. 2013. Big Data Integration. *VLDB* (2013).
- [10] Cynthia Dwork. 2006. Differential Privacy. In *ICALP 2006*.
- [11] Cynthia Dwork and Aaron Roth. 2014. The Algorithmic Foundations of Differential Privacy. *Found. Trends Theor. Comput. Sci.* (2014).
- [12] Michael J. Freedman, Kobbi Nissim, and Benny Pinkas. 2004. Efficient Private Matching and Set Intersection. In *EUROCRYPT*.
- [13] Lise Getoor and Ashwin Machanavajjhala. 2013. Entity Resolution for Big Data. In *KDD*.
- [14] Aristides Gionis, Piotr Indyk, and Rajeev Motwani. 1999. Similarity Search in High Dimensions via Hashing. In *VLDB*.
- [15] Oded Goldreich. 2004. *Foundations of Cryptography: Volume 2, Basic Applications*. Cambridge University Press, New York, NY, USA.
- [16] Slawomir Goryczka, Li Xiong, and Vaidy Sunderam. 2013. Secure Multiparty Aggregation with Differential Privacy: A Comparative Study. In *EDBT*.
- [17] Xi He, Ashwin Machanavajjhala, and Bolin Ding. 2014. Blowfish Privacy: Tuning Privacy-utility Trade-offs Using Policies. In *SIGMOD*.
- [18] Ali Inan, Murat Kantarcioglu, Elisa Bertino, and Monica Scannapieco. 2008. A Hybrid Approach to Private Record Linkage. In *ICDE*.
- [19] Ali Inan, Murat Kantarcioglu, Gabriel Ghinita, and Elisa Bertino. 2010. Private Record Matching Using Differential Privacy. In *EDBT*.
- [20] Alexandros Karakasidis and Vassilios S. Verykios. 2009. Privacy Preserving Record Linkage Using Phonetic Codes. In *BCI*.
- [21] Dimitrios Karapiperis and Vassilios S. Verykios. 2015. An LSH-Based Blocking Approach with a Homomorphic Matching Technique for Privacy-Preserving Record Linkage. *TKDE* (2015).
- [22] Michael Kearns, Aaron Roth, Zhiwei Steven Wu, and Grigory Yaroslavtsev. 2016. Private algorithms for the protected in social network search. *Proceedings of the National Academy of Sciences* (2016).
- [23] Hanna Köpcke, Andreas Thor, and Erhard Rahm. 2010. Benchmark datasets for entity resolution. http://dbs.uni-leipzig.de/en/research/projects/object_matching/fever/benchmark_datasets_for_entity_resolution. (2010).
- [24] Mehmet Kuzu, Murat Kantarcioglu, Ali Inan, Elisa Bertino, Elizabeth Durham, and Bradley Malin. 2013. Efficient Privacy-aware Record Integration. In *EDBT*.
- [25] Ilya Mironov, Omkant Pandey, Omer Reingold, and Salil Vadhan. 2009. Computational Differential Privacy. In *CRYPTO*.
- [26] Noman Mohammed, Dima Alhadidi, Benjamin C. M. Fung, and Mourad Debbabi. 2014. Secure Two-Party Differentially Private Data Release for Vertically Partitioned Data. *IEEE Transactions on Dependable and Secure Computing* (2014).
- [27] Arjun Narayan and Andreas Haeberlen. 2012. DJoin: Differentially Private Join Queries over Distributed Databases. In *OSDI*.
- [28] Pascal Paillier. 1999. *Advances in Cryptology*. Springer Berlin Heidelberg, Chapter Public-Key Cryptosystems Based on Composite Degree Residuosity Classes.
- [29] Chaoyi Pang, Lifang Gu, David Hansen, and Anthony Maeder. 2009. *Privacy-Preserving Fuzzy Matching Using a Public Reference Table*.
- [30] Manas Pathak, Shantanu Rane, and Bhiksha Raj. 2010. Multiparty Differential Privacy via Aggregation of Locally Trained Classifiers. In *Advances in Neural Information Processing Systems 23*, J. Lafferty, C. Williams, J. Shawe-taylor, R.s. Zemel, and A. Culotta (Eds.). http://books.nips.cc/papers/files/nips23/NIPS2010_0408.pdf
- [31] Martin Pettai and Peeter Laud. 2015. Combining Differential Privacy and Secure Multiparty Computation. In *ACSAC*.
- [32] Benny Pinkas, Thomas Schneider, and Michael Zohner. 2016. Scalable Private Set Intersection Based on OT Extension. *Cryptology ePrint Archive*, Report 2016/930. (2016). <http://eprint.iacr.org/2016/930>.
- [33] Pradeep Ravikumar and Stephen E. Fienberg. 2004. A secure protocol for computing string distance metrics. In *In PSDM held at ICDM*. 40–46.
- [34] Monica Scannapieco, Ilya Figotin, Elisa Bertino, and Ahmed K. Elmagarmid. 2007. Privacy Preserving Schema and Data Matching. In *SIGMOD*.
- [35] Rainer Schnell, Tobias Bachteler, and Jörg Reiher. 2009. Privacy-preserving record linkage using Bloom filters. *BMC Medical Informatics and Decision Making* 9, 1 (2009), 41.
- [36] Adam Smith. 2015. The Privacy of Secured Computations. In *Crypto & Big Data Workshop*.
- [37] NYC Taxi and Limousine Commission. 2013. TLC Trip Record Data. http://www.nyc.gov/html/tlc/html/about/trip_record_data.shtml. (2013).
- [38] Brian Thorne. 2016. Python-paillie. <https://readthedocs.org/projects/python-paillier/>. (2016).
- [39] Sameer Wagh, Paul Cuff, and Prateek Mittal. 2016. Root ORAM: A Tunable Differentially Private Oblivious RAM. *CoRR abs/1601.03378* (2016).
- [40] Andrew Chi-Chih Yao. 1986. How to Generate and Exchange Secrets. In *SFCS*.
- [41] Qingsong Ye, Ron Steinfeld, Josef Pieprzyk, and Huaxiong Wang. 2009. Efficient Fuzzy Matching and Intersection on Private Datasets. In *ICISC*.

A RELATED PRIVACY DEFINITIONS

A.1 Simulation-Based S2PC

The standard simulation-based definition for SMC is defined below.

Definition A.1 (SIM-S2PC). [15] For a functionality f , a 2-party protocol Π which computes f provides simulation-based secure 2-party computation (SIM-S2PC) if for all data sets D_A, D_B of polynomial sizes (in κ), there exist probabilistic polynomial-time algorithms (simulators), denoted by S_A and S_B such that the distribution of S_A (resp., S_B) is computationally indistinguishable from VIEW_A^Π (resp., VIEW_B^Π), i.e. for any probabilistic polynomial-time (in κ) adversary T ,

$$\begin{aligned} & \Pr[T(S_A(D_A, f_A(D_A, D_B)), f(D_A, D_B))] = 1] \\ & \leq \Pr[T(\text{VIEW}_A^\Pi(D_A, D_B), O^\Pi(D_A, D_B)) = 1] + \text{negl}(\kappa) \\ & \Pr[T(S_B(D_A, f_B(D_A, D_B)), f(D_A, D_B))] = 1] \\ & \leq \Pr[T(\text{VIEW}_B^\Pi(D_A, D_B), O^\Pi(D_A, D_B)) = 1] + \text{negl}(\kappa). \end{aligned} \quad (7)$$

If f is deterministic, Alice gains no additional knowledge other than its respective input (D_A) and output ($f_A(D_A, D_B)$); similarly for Bob. When randomized functionalities are concerned, augmenting the view of the semi-honest party by the output of the other party is essential. In this case, for any protocol Π that computes the randomized functionality f , it does not necessarily hold that $O^\Pi(D_A, D_B) = f(D_A, D_B)$. Rather, these two random variables must be identically distributed. In order to study the possibility of composing DP and S2PC, we choose the indistinguishability-based definition for PRL, which is implied from SIM-S2PC.

THEOREM A.2. *SIM-S2PC implies IND-S2PC.*

PROOF. Given the protocol Π , for all possible inputs (D_A, D_B), there exists a global simulator S_A such that the distribution of S_A is computationally indistinguishable from the view of Alice. As $f(D_A, D_B) = f(D_A, D'_B)$, S_A takes the same input and hence will have the same distribution for D_B and D'_B . Hence, the views over (D_A, D_B) or (D_A, D'_B) are indistinguishable. $\square$

Any algorithm that satisfies SIM-CDP also satisfies IND-CDP [25], but it is unknown if the converse holds.

A.2 Computationally Differential Privacy

Mironov et al. [25] defines a privacy notion, known as indistinguishable computationally differential privacy (IND-CDP-2PC). This notion is a direct extension of DP in two party setting where both parties are computationally bounded. Formally, we have

Definition A.3 (IND-CDP-2PC). A 2-party protocol Π for computing function f satisfies $(\epsilon_A(\kappa), \epsilon_B(\kappa))$ -indistinguishable computationally differential privacy (IND-CDP-2PC) if $\text{VIEW}_A^\Pi(D_A, \cdot)$ satisfies $\epsilon_B(\kappa)$ -IND-CDP, i.e. for any probabilistic polynomial-time (in κ) adversary T , for any neighboring databases (D_B, D'_B) differing in a single row,

$$\begin{aligned} & \Pr[T(\text{VIEW}_A^\Pi(D_A, D_B)) = 1] \\ & \leq e^{\epsilon_B} \Pr[T(\text{VIEW}_A^\Pi(D_A, D'_B)) = 1] + \text{negl}(\kappa). \end{aligned} \quad (9)$$

The same holds for Bob's view for any neighbors (D_A, D'_A) and ϵ_A .

A.3 Local Differential Privacy

The local model is usually considered in the model where individuals do not trust the curator with their data. The local version of differential privacy is defined as follows.

Definition A.4 (ϵ -Local Differential Privacy). [11] A randomized mechanism $M : \Sigma \rightarrow \mathcal{O}$ satisfies ϵ -local differential privacy if

$$\Pr[M(r) = O] \leq e^\epsilon \Pr[M(r') = O] \quad (10)$$

for any set $O \subseteq \mathcal{O}$, and any records $r, r' \in \Sigma$ and $\epsilon > 0$.

B THEOREMS & PROOFS

B.1 Privacy Leakage in Prior Work

B.1.1 Theorem 2.4 (Limitations of PRL with Blocking). Given (d_1, d_2, p_1, p_2) -sensitive $H = \{h_0, \dots, h_{|H|-1}\}$, we use $H(\cdot)$ for a record to denote the list of hashing values $[h_0(\cdot), \dots, h_{|H|-1}(\cdot)]$. An LSH-based blocking considers a set of bins where records associated with the same value for $H(\cdot)$ are hashed to the same bin. A popular blocking strategy is to compare all the corresponding bins, and results in a set of candidate matches $\{(a, b) | h(a) = h(b) \forall h \in H, a \in D_A, b \in D_B\}$, i.e. $\{(a, b) | H(a) = H(b), \forall a \in D_A, b \in D_B\}$. We can show that any LSH based blocking cannot satisfy IND-S2PC as stated in Theorem 2.4. Here is the proof.

PROOF. Take a pair of databases (D_B, D'_B) where $f_{\bowtie_m}(D_A, D_B) = f_{\bowtie_m}(D_A, D'_B)$. Let the symmetric difference between D_B and D'_B be (b, b') and $\text{dist}(b, b') > d_2$. Hence, with high probability $1 - p_2^{|H|}$, we have $H(b) \neq H(b')$, and $|\mathcal{B}_{H(b)}(D_B)| - |\mathcal{B}_{H(b)}(D'_B)| = 1$ and $|\mathcal{B}_{H(b')}(D'_B)| - |\mathcal{B}_{H(b')}(D_B)| = 1$ as the rest of records are the same in D_B and D'_B . Alice as a semi-honest adversary can set her dataset such that $|\mathcal{B}_{H(b)}(D_A)| \neq |\mathcal{B}_{H(b')}(D_A)|$. Then, with high probability, the following inequality holds

$$\begin{aligned} & \text{cost}_{\mathcal{B}^S}(D_A, D_B) - \text{cost}_{\mathcal{B}^S}(D_A, D'_B) \\ & = (|\mathcal{B}_{H(b)}(D_B)| - |\mathcal{B}_{H(b)}(D'_B)|) |\mathcal{B}_{H(b)}(D_A)| \\ & \quad + (|\mathcal{B}_{H(b')}(D_B)| - |\mathcal{B}_{H(b')}(D'_B)|) |\mathcal{B}_{H(b')}(D_A)| \\ & = |\mathcal{B}_{H(b)}(D_A)| - |\mathcal{B}_{H(b')}(D_A)| \neq 0. \end{aligned} \quad (11)$$

Hence, Alice can distinguish D_B and D'_B by $\text{cost}_{\mathcal{B}^S}(D_A, D_B) \neq \text{cost}_{\mathcal{B}^S}(D_A, D'_B)$ with high probability $1 - p_2^{|H|}$. Other blocking strategies can be similarly shown. Therefore, this LSH-based PRL does not satisfy IND-S2PC. $\square$

B.1.2 Theorem 4.7 (Limitations of PRL with DP Blocking of Prior Approaches/LP-2). Several prior works [5, 19, 24] combine PRL techniques with differentially private blocking (PRL + $\mathcal{B}_{DP}$). These approaches can be summarized in three steps: (1) DP blocking, (2) records addition and suppression, (3) secure pair-wise comparisons based on blocking strategy $\mathcal{B}^S$. In the first step, Alice and Bob process their data independently. Each party generates an ϵ -differentially private partition of the data, where each partition is associated with a noisy count $\tilde{o}_i = |\mathcal{B}_i(D_B)| + \eta_i$, where $\Pr[\eta_i = x] = p e^{-\epsilon/\Delta \mathcal{B} \cdot |x|}$, for $x \in \mathbb{Z}$ and $p = \frac{e^{\epsilon/\Delta \mathcal{B} - 1}}{e^{\epsilon/\Delta \mathcal{B}} + 1}$ is the normalized factor². $\Delta \mathcal{B}$ is the sensitivity of the blocking strategy (Def 4.2).

²We use discrete version of Laplace distribution to avoid rounding.

Next, for each partition, if the noise η_i is positive, dummy records are added; otherwise, records in that partition are suppressed randomly to obtain the published count. This results in new bins, denoted by $\{\tilde{\mathcal{B}}_i(D_A)\}$ and $\{\tilde{\mathcal{B}}_j(D_B)\}$. In the last step, Alice and Bob jointly compare record pairs (a, b) , where $a \in \tilde{\mathcal{B}}_i(D_A)$ and $b \in \tilde{\mathcal{B}}_j(D_B)$ for all $(i, j) \in \mathcal{B}^S$ as in APC. They only exchange the true records (a, b) if they match. [5] considers a third party for identifying candidate pairs for Alice and Bob, so that Alice and Bob has no direct access to the noisy bins of the opposite party, but has access to the number of secure comparisons. However, this hybrid protocol above does not satisfy (ϵ, δ) -DPRL as stated in Theorem 4.7. The failure to satisfy DPRL is mainly caused by the record suppression step for the negative noise drawn from a zero-mean Laplace distribution, as shown in the following proof.

PROOF. Without loss of generality, we consider Alice as the adversary. For any arbitrary ϵ and small $\delta < \frac{p^{\Delta\mathcal{B}}}{2e^\epsilon}$, there exists a counter example fails (ϵ, δ) -DPRL. For simplicity, we illustrate how to construct counterexamples using a blocking strategy $\mathcal{B}$ with sensitivity $\Delta\mathcal{B} = 2$, where Alice and Bob use the same hashing and each record is hashed to at most 1 bin. For other blocking strategies, counterexamples can be similarly constructed.

Fix a D_A , consider D_B such that $\mathcal{B}_0(D_B) = \{b_*\}$ and $\mathcal{B}_1(D_B) = \{b_1, \dots, b_{n_1}\}$, where $1 \leq n_1 < \frac{p^2}{e^\epsilon \delta - 1}$. (Note that $\frac{p^2}{e^\epsilon \delta} > 2$ because $\delta < \frac{p^2}{2e^\epsilon}$.) In addition, all records in $\mathcal{B}_1(D_B)$ can find some matching ones from D_A , but b_* does not match any record in D_A . A neighboring database D'_B can be constructed from D_B by removing b_* from $\mathcal{B}_0$, and adding another b'_* that can be hashed to $\mathcal{B}_1$. It is easy to see that $(D_B, D'_B) \in \mathcal{N}(f_{\bowtie_m}(D_A, \cdot))$.

Without a third party [19, 24], Alice and Bob has access to the number of secure comparisons and the noisy bin counts (in addition to the input data sizes and the matching output). Consider a set of views of Alice VIEW^* with output that contains all matching pairs from $\mathcal{B}_1(D_A) \bowtie_m \mathcal{B}_1(D_B)$ and noisy counts for bin $\mathcal{B}_0$ and $\mathcal{B}_1$ for Bob being 0 and n_1 respectively. Let $\tilde{\mathcal{B}}(D_A)$ be the noisy bins that Alice uses for the final secure pairwise comparisons. The probabilities to generate these views from D_B and D'_B are respectively:

$$\Pr[\text{VIEW}^*|\tilde{\mathcal{B}}(D_A), D_B] = \Pr[\eta_0 = 0] \Pr[\eta_1 = 0] = p^2, \quad (12)$$

$$\begin{aligned} \Pr[\text{VIEW}^*|\tilde{\mathcal{B}}(D_A), D'_B] &= \Pr[\eta_0 = 1] \Pr[\eta_1 = -1 \& \text{suppress } b'_*] \\ &= p^2 / (e^\epsilon (n_1 + 1)) > \delta \end{aligned} \quad (13)$$

The inequality above is due to $n_1 < \frac{p^2}{e^\epsilon \delta} - 1$. Hence, we have

$$\begin{aligned} \Pr[\text{VIEW}^*|\tilde{\mathcal{B}}(D_A), D_B] &= (e^\epsilon + n_1 e^\epsilon) \Pr[\text{VIEW}^*|\tilde{\mathcal{B}}(D_A), D'_B] \\ &> e^\epsilon \Pr[\text{VIEW}^*|\tilde{\mathcal{B}}(D_A), D'_B] + \delta. \end{aligned} \quad (14)$$

Hence, (ϵ, δ) -DPRL is violated.

With a third party [5], Alice and Bob has access to the final output, and the total number of secure pairwise comparisons, but not the noisy bin counts. We can construct examples where knowing the number of secure comparisons leaks the noisy bin counts. After which the previous arguments (for the case with no third party) can show that this protocol does not ensure DPRL for all epsilon and delta. For instance, consider Alice has only 1 record in $\tilde{\mathcal{B}}_0(D_A)$, and more than 1 records in other bins, if the output $O = D_A \bowtie_m D_B$, and the total number of secure pairwise comparisons is $|O| + 1$.

This secure pairwise matching that returns false can only happen between a record of Bob from $\tilde{\mathcal{B}}_0(D_B)$ with the record from $\tilde{\mathcal{B}}_0(D_A)$. Hence, Alice can infer the noisy counts of $\tilde{\mathcal{B}}(D_B)$. Then the argument for the case with no third party can be used. $\square$

In addition, by Theorem 3.8 and Theorem 3.9, DPRL is weaker than IND-S2PC and IND-CDP-2PC, we have the following result.

COROLLARY B.1. *LP-2 satisfies neither IND-CDP-2PC nor IND-S2PC.*

B.2 Properties of Output Constrained DP

B.2.1 Theorem 3.4 (Sequential Composition).

PROOF. Consider Alice as a probabilistic polynomial-time (in κ) adversary T , with input D_A . (D_B, D'_B) are neighbors w.r.t. $f(D_A, \cdot)$. We have the probabilities of distinguishing D_B and D'_B bounded by

$$\begin{aligned} &\Pr[T(\text{VIEW}_A^{\Pi_2, \Pi_1}(D_A, D_B)) = 1] \\ &\leq \int_x \Pr[T(\text{VIEW}_A^{\Pi_2}(D_A, D_B, x)) = 1] \cdot \\ &\quad \Pr[x = \text{VIEW}_A^{\Pi_1}(D_A, D_B)] dx \\ &\leq \int_x (e^{\epsilon_2} \Pr[T(\text{VIEW}_A^{\Pi_2}(D_A, D'_B, x)) = 1] + \delta_2) \cdot \\ &\quad \Pr[x = \text{VIEW}_A^{\Pi_1}(D_A, D_B)] dx \\ &\leq \int_x (e^{\epsilon_2} \Pr[T(\text{VIEW}_A^{\Pi_2}(D_A, D'_B, x)) = 1]) \cdot \\ &\quad (e^{\epsilon_1} \Pr[x = \text{VIEW}_A^{\Pi_1}(D_A, D'_B)] + \delta_1) dx + \delta_2 \\ &\leq e^{\epsilon_1 + \epsilon_2} \Pr[T(\text{VIEW}_A^{\Pi_2, \Pi_1}(D_A, D'_B)) = 1] + \delta_1 + \delta_2 \end{aligned}$$

$\square$

B.2.2 Theorem 3.5 (Post-processing).

PROOF. Since g is efficient and in composition with T can be used as adversary itself. If $g(\text{O}^\Pi(D_A, D_B))$ does not satisfy (ϵ, δ) -IND-DPRL, then Π does not satisfy (ϵ, δ) -IND-DPRL. $\square$

B.2.3 Theorem 3.9 (Relation with IND-CDP-2PC). We show that DPRL is weaker than IND-CDP-2PC.

PROOF. $\epsilon/2$ -IND-CDP-2PC is equivalent to ϵ -IND-DP-2PC, where neighboring databases have a symmetric difference of 2. The set of neighboring databases for DPRL is a subset of that for ϵ -IND-DP-2PC, and hence (ϵ, δ) -DPRL is weaker than $\epsilon/2$ -IND-DP-2PC. $\square$

B.3 Properties for DPRL Protocols

B.3.1 Theorem 4.5 (Privacy of Laplace Protocol).

LEMMA B.2. *With probability $1 - \delta$, the probability for Alice having the same view from neighboring databases $(D_B, D'_B) \in \mathcal{N}(f_{\bowtie_m}(D_A, \cdot))$ is bounded by e^ϵ .*

PROOF. Given $(D_B, D'_B) \in \mathcal{N}(f_{\bowtie_m}(D_A, \cdot))$ and $\mathcal{B}$, the maximum difference in the bin counts of D_B and D'_B is $\Delta\mathcal{B}$. Let $\mathcal{B}_\Delta$ be the set of bins that D_B and D'_B have different counts, and $\sum_{i \in \mathcal{B}_\Delta} |\mathcal{B}_i(D_B) - \mathcal{B}_i(D'_B)| \leq \Delta\mathcal{B}$. If all the noise for these bins are non-negative, then

the probability to output the same noisy counts $(c_0, \dots, c_{k-1})$ from D_B and D'_B is bounded by

$$\begin{aligned} & \ln \left(\frac{\Pr[(c_0, \dots, c_{k-1} | D_B)]}{\Pr[(c_0, \dots, c_{k-1} | D'_B)]} \right) \\ &= \ln \left(\frac{\prod_{i=0}^{k-1} \Pr[\eta_i = c_i - |\mathcal{B}_i(D_B)|]}{\prod_{i=0}^{k-1} \Pr[\eta_i = c_i - |\mathcal{B}_i(D'_B)|]} \right) \\ &= \sum_{i \in \mathcal{B}_\Delta} \ln(\Pr[\eta_i = c_i - |\mathcal{B}_i(D_B)|]) - \ln(\Pr[\eta_i = c_i - |\mathcal{B}_i(D'_B)|]) \\ &\leq \epsilon / \Delta \mathcal{B} \cdot \left(\sum_{i \in \mathcal{B}_\Delta} |\mathcal{B}_i(D_B) - \mathcal{B}_i(D'_B)| \right) \leq \epsilon \end{aligned}$$

The probability to draw a negative noise η from $Lap(\epsilon, \delta, \Delta \mathcal{B})$ is

$$\Pr[\eta < 0] = \sum_{i=-1}^{-\infty} p \cdot e^{-(\epsilon / \Delta \mathcal{B})(x - \eta^0)} = \frac{e^{-\eta^0 \epsilon / \Delta \mathcal{B}}}{e^{\epsilon / \Delta \mathcal{B}} + 1} \quad (15)$$

Given $\eta_0 = -\frac{\Delta \mathcal{B} \ln((e^{\epsilon / \Delta \mathcal{B}} + 1)(1 - (1 - \delta)^{1 / \Delta \mathcal{B}}))}{\epsilon}$, we have $\Pr[\eta < 0] = 1 - (1 - \delta)^{1 / \Delta \mathcal{B}}$. For each neighboring pair, at most $\Delta \mathcal{B}$ bins differ and fail to have $\Pr[\eta \geq 0]$. Hence, the overall failing probability is $1 - (1 - \Pr[\eta < 0])^{\Delta \mathcal{B}} = \delta$. With $1 - \delta$, the probability of having the same view from PRL neighboring databases is bounded by e^ϵ . $\square$

B.3.2 Theorem 4.11 (Correctness & Efficiency of GMC).

PROOF. First, we will show that the efficiency of LP with the greedy match & clean step (GMC) is better than LP alone. The first part of the protocol that adds dummy records is the same. The second part of the protocol without GMC compares all the candidate matches using the secure matching protocol $SMC(a, b)$. On the other hand, with GMC, if a record pair (a, b) is compared securely, then (a, b) must be one of the candidate matches. Hence, the number of the secure pairwise comparisons with GMC will be no more than the protocol without GMC.

Next, we will show the correctness of LP with GMC. Let O_{LP}, O_{LP+GMC} be the final output of LP protocol without GMC and with GMC. We would like to show that if $(a, b) \in O_{LP}$, then $(a, b) \in O_{LP+GMC}$. Suppose this is not true, then there exists a matching pair $(a, b) \in O_{LP}$, but $(a, b) \notin O_{LP+GMC}$. If so, then one of the records in (a, b) must be removed from the bins before its turn of secure pairwise comparison $SMC(a, b)$. Without loss of generality, let's say a is cleaned from Alice's bins before $SMC(a, b)$. The condition to remove a is that a has already been in the current output. Hence, Bob is able to compare a with all his records in plain text and identify this matching pair (a, b) . This leads to a contradiction. Hence, $O_{LP} \subseteq O_{LP+GMC}$. Moreover, if a matching pair (a, b) is not a candidate match based on the blocking strategy $\mathcal{B}^S$, and if a has been already found matching with another record of Bob, then GMC can add (a, b) into O_{LP+GMC} . Hence, it is possible that LP with GMC gains even more matching pairs than LP alone. $\square$

C ADDITIONAL PROTOCOLS

C.1 Example for secure pairwise match

Here we give an example for the function $SMC(a, b)$ that outputs (a, b) if they match; null otherwise. The matching rule is that Euclidean distance of a and b is less than θ . First, Party Alice creates a homomorphic public/private key pair (pk, pr) , and sends the public

Algorithm 4: Secure Match a and b

```

1 function SecureMatch  $(a, b, \theta)_{pk, pr}$ ;
   Input :  $a, b \in \{0, 1\}^d$ , hamming distance threshold  $\theta$ ,
           public/private key pair  $(pk, pr)$ 
   Output:  $(a, b)$  or  $\emptyset$ 
2 Alice: randomly generates an id  $a_{id}$  and sends to Bob;
3 Bob: randomly generates an id  $b_{id}$  and an integer  $r$ ;
4 Bob: initiates  $s$  with  $E_{pk}(r)$ ;
5 for  $i \in [0, \dots, d - 1]$  do
6   Alice: sends to Bob  $E_{pk}(a[i])$ ;
7   Bob: updates
        $s = s +_h E_{pk}(a[i]) +_h (E_{pk}(a[i]) \times_h (-2b[i])) +_h E_{pk}(b[i])$ ;
8 end
9 Bob: sends  $(b_{id}, s)$  to Alice;
10 Alice: decrypts  $s = D_{pr}(s)$ ;
11 if  $s \leq \theta + r$  (secure integer comparison) then
12   return  $(a, b)$ ;
13 else
14   return  $\emptyset$ ;
15 end
```

key pk to party Bob. Let $E_{pk}(\cdot)$ denote the encryption function with public key pk and $D_{pr}(\cdot)$ the decryption function with private key pr . Paillier's cryptosystem supports the following operations on the encrypted plain texts m_1 and m_2 without the knowledge of the private key:

- Addition: $E_{pk}(m_1 + m_2) = E_{pk}(m_1) +_h E_{pk}(m_2)$;
- Multiplication with constant c : $E_{pk}(cm_1) = c \times_h E_{pk}(m_1)$

These two operations allow secure computation of Euclidean distances, i.e. $dist(a, b) = \sum_i (a[i] - b[i])^2 = \sum_i (a[i]^2 - 2a[i]b[i] + (b[i])^2)$, and also hamming distances for bit vectors, i.e. $dist(a, b) = \sum_i |a[i] - b[i]| = \sum_i a[i] + b[i] - 2a[i]b[i]$.

As summarized in Algorithm 4, given bit vectors a and b , Alice will send to Bob the encrypted values $(a_{id}, \{E_{pk}(a[i])\})$ where a_{id} is a randomly generated record identifier for record a . Next, party Bob computes for each of its records b_{id} the value $E_{pk}(a[i]) +_h E_{pk}(a[i]) \times_h (-2b[i]) +_h E_{pk}(b[i])$ which is equal to $E_{pk}(|a[i] - b[i]|)$ for all i , and computes the encrypted $E_{pk}(\sum_i |a[i] - b[i]|)$. A random number r is generated and added to the encrypted distance, such that the true distance is hidden from Alice if (a, b) is not a matching pair. Party Bob creates the message $(b_{id}, E_{pk}(\sum_i |a[i] - b[i]| + r))$ for each record pair comparison. Alice can then decrypt the message with her private key and obtain the relative distance $d = \sum_i |a[i] - b[i]| + r$. Since Bob knows $\theta + r$, a secure comparison protocol, such as Yao's garbled circuit [40], can be used to evaluate if $d \leq \theta + r$. If this algorithm outputs "True", Alice and Bob will exchange their true record values.

C.2 Local DP Protocol

C.2.1 Theorem 4.1. Let $\mathcal{B}$ be a blocking that randomly hashes records into a pre-specified set of k bin, such that $\frac{\Pr[\mathcal{B}(b)=i]}{\Pr[\mathcal{B}(b')=i]} \leq e^\epsilon$. Such a blocking $\mathcal{B}$ satisfies ϵ -local DP (Appendix A.3). Protocols

that combine a local differentially private blocking with IND-S2PC protocols for record linkage to achieve $(\epsilon, \text{negl}(\kappa))$ -DPRL.

PROOF. (sketch) We prove privacy for Bob (the proof for Alice is analogous). In this protocol, Alice with input Data D_A has a view consisting of (1) the number of candidate matching pairs arising in each $(i, j) \in \mathcal{B}^S$, (2) the output for each candidate matching pair. Consider a neighboring pair $(D_B, D'_B) \in N(f_{\bowtie_m}(D_A, \cdot))$ for a given D_A . By Theorem 3.6, D_B and D'_B differ in only one non-matching record with respect to D_A , i.e. $D'_B = D_B - b_* + b'_*$ and $b_* \neq b'_*$, where $m(b_*, a) = 0$ and $m(b'_*, a) = 0$ for all $a \in D_A$. Given both b_*, b'_* can be hashed into the same bin with probability ratio bounded by e^ϵ , the probabilities of generating the same number of candidate matching pairs from D_B and D'_B are also bounded by the same ratio. The encrypted records only differ in b_* and b'_* , and both of them lead to the same output for each candidate matching, because they do not match any records in D_A . $\square$

In this work, we use randomized response (RR) [11] as an example to achieve DPRL. Other local DP algorithms, such as Johnson-Lindenstrauss (JL) transform [4] can be similarly applied.

C.2.2 RR based Blocking. Given a fixed hash function $h: \Sigma \rightarrow [0, k-1]$, records in D_B are hashed into k bins, $\mathcal{B}_0, \dots, \mathcal{B}_{k-1}$ respectively. Let us define a RR based on this fixed hashing function with privacy budget ϵ_B for Bob. Each record $b \in D_B$ is randomly hashed into $\mathcal{B}_{h(b)}$ with probability $p_B = \frac{e^{\epsilon_B}}{k-1+e^{\epsilon_B}}$ and the other $(k-1)$ bins with probability $q_B = \frac{1}{k-1+e^{\epsilon_B}}$. We denote the resulted bins by $\tilde{\mathcal{B}}(D_B)$, and the resulted bin for each record $b \in D_B$ by $\tilde{\mathcal{B}}(b)$. Similarly, using the same fixed hashing function and randomized response, Alice's records D_A are randomly hashed into the k bins, $\tilde{\mathcal{B}}(D_A)$ with corresponding p_A, q_A based on ϵ_A . This randomized response with probabilities (p_A, q_A) and (p_B, q_B) ensures ϵ_A -local DP and ϵ_B -local DP respectively.

Consider a basic blocking strategy $\mathcal{B}^S = \{(i, i) | i \in [0, k)\}$, all corresponding bins are compared. If the hash function h is a LSH, then matching records are likely fall into the same bin as Alice and Bob use the same hash function. The probability that such records (a, b) appear in the same bin after randomization $\Pr[\tilde{\mathcal{B}}(a) = \tilde{\mathcal{B}}(b) | h(a) = h(b)]$ is $p_A p_B + (k-1)q_A q_B$. This probability increases with the privacy budget ϵ_A, ϵ_B , and hence recall will improves.

A further trade-off between correctness, privacy and efficiency is allowed by considering a general blocking strategy $\mathcal{B}^S = \{(i, (i+j)\%k) | i \in [0, k), j \in [0, k')\}$, each record in $\tilde{\mathcal{B}}_i(D_A)$ is securely compared with k' neighboring bins of Bob. The basic blocking strategy corresponds to the case where $k' = 1$. As k' increases, more candidate matching pairs are securely compared, resulting potentially higher recall and more communication and computation cost. When $k' = k$, the resulted protocol is equivalent to the all pairwise comparisons baseline.

C.2.3 Optimal RR Probability. If the window size k' is given in advance as a parameter for the efficiency, the expected recall of randomized response can be further optimized. Let p_i^B be the probability for a record $b \in D_B$ to be randomly hashed into $\mathcal{B}_{(h(b)+i)\%k}$. To ensure ϵ_B -DPRL, the probabilities to hash any pairs of bins should be bounded by e^{ϵ_B} . The hashing probability for Alice records, p_i^A , is similarly defined and constrained. As each record of Alice's can

fall into any bin, and each bin of Alice's is compared with k' neighboring bins of Bob's, the probability that a matching pair (a, b) is compared after randomization is

$$p(a, b) = \sum_{i=0}^{k-1} \sum_{j=0}^{k'-1} p_i^A p_{i+j\%k}^B. \quad (16)$$

The expected recall can be improved by maximizing $p(a, b)$ (Eqn. (16)) with the constraints on

- (a) Ratio: $\frac{p_i^B}{p_{i'}^B} \leq e^{\epsilon_B}, \forall i \neq i'; \frac{p_i^A}{p_{i'}^A} \leq e^{\epsilon_A}, \forall i \neq i';$
- (b) Sum: $\sum_{i=0}^{k-1} p_i^A = 1; \sum_{i=0}^{k-1} p_i^B = 1.$

The solution in Section C.2.2 where we assign a high probability to a single bin, and a low probability to the rest of the bins is a valid solution to the maximization problem, but it is not always the optimal solution. In general, this optimization can be solved by existing tools for quadratic objectives with linear constraints, such as quadratic programming.

Here, we present an explicit solution for the special case when $\epsilon_B = \epsilon \geq 0$ and $\epsilon_A \rightarrow \infty$. This case corresponds to the situation where Alice's data is public and only Bob's data requires ϵ -DPRL protection. Then the expected recall can be maximized by solving the following linear optimization:

$$\max_{p_0^B, \dots, p_{k-1}^B} \sum_{i=0}^{k'-1} p_i^B \text{ s.t. } \sum_{i=0}^{k-1} p_i^B = 1, \text{ and } \frac{p_i^B}{p_{i'}^B} \leq e^\epsilon \forall i \neq i'$$

The expected recall is maximized with value $p(a, b) = \frac{k'e^\epsilon}{k-k'+k'e^\epsilon}$ at

$$p_i^B = \begin{cases} \frac{e^\epsilon}{k-k'+k'e^\epsilon}, & \text{for } i = 0, \dots, k' - 1 \\ \frac{1}{k-k'+k'e^\epsilon}, & \text{for } i = k', \dots, k - 1 \end{cases}$$

If Bob's records are uniformly distributed over the bins, then the compression ratio w.r.t all pairwise comparisons (APC) is $\rho = \frac{k'}{k}$. The maximized expected recall can be written as

$$p(a, b) = \frac{\rho e^\epsilon}{1 - \rho + \rho e^\epsilon}. \quad (17)$$

This equation explicitly form the relationship between correctness $(p(a, b))$, privacy (ϵ) and efficiency (ρ) of this protocol.

THEOREM C.1. *The basic RR mechanism achieves a constant factor speedup in efficiency given $\epsilon_B = \epsilon \geq 0, \epsilon_A \rightarrow \infty, \delta_A = \delta_B = \text{negl}(\kappa)$ and recall r .*

PROOF. Given a recall $r = p(a, b)$, we have $\rho = 1 - \frac{e^\epsilon(1-r)}{r+e^\epsilon(1-r)}$ based on Eqn. (17). The improvement in efficiency ρ is a constant factor in terms of r and ϵ , independent of n . $\square$

Next, we compute the optimal amongst a restricted class of strategies for the more general case where $\epsilon_A = \epsilon_B = \epsilon$, though the explicit form for this case is unknown yet. The strategies we consider are those where Alice and Bob (a) use symmetric probabilities to assign a bin to each record, and (b) they both assign a high probability $p_\top$ to place a record from bin i to bins i through $(i+x-1)\%k$ (for some $1 \leq x \leq k'$), and a low probability $p_\perp$ to assign a record from bin i to the rest of the bins. Note that, when only one of Alice or Bob is randomizing their records, $x = k'$ results in the RR probabilities that optimize the expected recall.

In order to satisfy the constraints in the above maximization problem, we need $p_{\top} = \frac{e^{\epsilon}}{k-x+x \cdot e^{\epsilon}}$ and $p_{\perp} = \frac{1}{k-x+x \cdot e^{\epsilon}}$.

The expected recall can be derived as follows. Without loss of generality consider a matching pair (a, b) that fall into bin 0. There are 3 ways (a, b) are matched after randomization:

(i): Both Alice and Bob randomize their records to a $\top$ bin (i.e., some bin $0 \leq j \leq x-1$). Since $x \leq k'$, these records are definitely compared. This occurs with probability $\frac{x(x+1)}{2} p_{\top}^2$.

(ii): Only one of Alice and Bob randomize their records to a $\top$ bin (i.e., some bin $0 \leq j \leq x-1$). There $2k'x - x(x+1)$ ways in which exactly one of a or b is randomized to a $\top$ bin, but still end up getting compared by the algorithm. This occurs with probability $(2k'x - x(x+1)) p_{\top} \cdot p_{\perp}$.

(iii): Both Alice and Bob randomize their records to a $\perp$ bin. This occurs with probability $(kk' - (2k'x - \frac{x(x+1)}{2})) p_{\perp}^2$. In total, the probability that (a, b) are compared is expressed in terms of x as

$$p_{(a,b)}(x) = \frac{x(x+1)}{2} p_{\top}^2 + (2k'x - x(x+1)) p_{\top} \cdot p_{\perp} + (kk' - (2k'x - \frac{x(x+1)}{2})) p_{\perp}^2$$

The derivative of $p_{(a,b)}(x)$ w.r.t x is

$$p'_{(a,b)}(x) = C_1 \cdot [(e^{\epsilon} - 1)C_2x + k(2k' + e^{\epsilon} - 1)],$$

where $C_1 = \frac{e^{\epsilon}-1}{2(k-x+x e^{\epsilon})^3}$ and $C_2 = (e^{\epsilon} - 3 + 2k - 4k')$. When $C_2 > 0$, the derivative is always positive, the expected recall is maximized

when $x = k'$, as $0 < x \leq k'$. We will leave the complete analysis to the future work.

D ADDITIONAL PLOT

Figure 5 shows the $\log(\text{base } 10)$ value of the average cost with respect to the $\log$ value of data size for PSI+X, APC, and LP with $\epsilon \in \{0.1, 0.4, 1.6\}$ and $\delta = 10^{-5}$ and the non-private setting (np) when they achieve a recall > 0.95 . Similar to Figure 2, LP gives lower costs than the baselines, and scales near linearly.

Figure 5: The average $\log(\text{cost})$ vs $\log(\text{data size})$