AMPPERE: A Universal Abstract Machine for Privacy-Preserving Entity Resolution Evaluation

Yixiang Yao
yixiangy@usc.edu

USC & Information Sciences Institute
Marina del Rey, California, USA

Srivatsan Ravi
srivatsr@usc.edu

USC & Information Sciences Institute
Marina del Rey, California, USA

Tanmay Ghai
tghai@usc.edu

USC & Information Sciences Institute
Marina del Rey, California, USA

Pedro Szekely
szekely@usc.edu

USC & Information Sciences Institute
Marina del Rey, California, USA

ABSTRACT

Entity resolution is the task of identifying records in different datasets that refer to the same entity in the real world. In sensitive domains (e.g. financial accounts, hospital health records), entity resolution must meet privacy requirements to avoid revealing sensitive information such as personal identifiable information to untrusted parties. Existing solutions are either too algorithmically-specific or come with an implicit trade-off between accuracy of the computation, privacy, and run-time efficiency. We propose **AMPPERE**, an abstract computation model for performing *universal* privacy-preserving entity resolution. **AMPPERE** offers abstractions that encapsulate multiple algorithmic and platform-agnostic approaches using variants of Jaccard similarity to perform private data matching and entity resolution. Specifically, we show that two parties can perform entity resolution over their data, without leaking sensitive information. We rigorously compare and analyze the feasibility, performance overhead and privacy-preserving properties of these approaches on the Sharemind multi-party computation (MPC) platform as well as on PALISADE, a lattice-based homomorphic encryption library. The **AMPPERE** system demonstrates the efficacy of privacy-preserving entity resolution for real-world data while providing a precise characterization of the induced cost of preventing information leakage.

CCS CONCEPTS

• **Information systems** → **Entity resolution; Deduplication; Security and privacy** → **Privacy-preserving protocols.**

KEYWORDS

entity resolution, privacy-preserving technique, homomorphic encryption, secure multi-party computation, abstract machine

ACM Reference Format:

Yixiang Yao, Tanmay Ghai, Srivatsan Ravi, and Pedro Szekely. 2021. AMP-PERE: A Universal Abstract Machine for Privacy-Preserving Entity Resolution Evaluation. In *Proceedings of the 30th ACM International Conference on Information and Knowledge Management (CIKM '21), November 1–5, 2021, Virtual Event, QLD, Australia*. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/3459637.3482318>

1 INTRODUCTION

In the real world, performing entity resolution over data from varied and distributed data sources is a challenging problem. For example, consider two different patient databases belonging to different hospitals, the first one containing records for Peter and Bruce, and the second for Tony, Pet, and Brvce (Pet is an abbreviation for Peter & Brvce is a misspelling of Bruce).

The *entity resolution* (ER) task here is to identify records from the two databases that belong to the same patients. An ER algorithm will compare the information present in the records (e.g. names, home or work address, phone numbers), and determine if the records refer to the same person. The principal challenge with this is that the two hospitals may record patient information in different formats (i.e. the first may use separate fields for first and last name, whereas the second may store the full name as one field). The *privacy-preserving entity resolution* (PPER) task is an extension of the ER task, with the additional requirement to not reveal sensitive information to any party present in the computation or to an adversary. Specifically, each party learns which of its records are present on other parties, but nothing about any other records; an adversary should learn nothing about *any* records of *any* party. Illustrating this with our original example, the first hospital may learn that the second hospital has records for Peter and Bruce, but should not learn about a record for Tony, or any other information about the patients themselves unless explicitly disclosed by the second hospital. Likewise, the second hospital should *only* learn about Peter & Bruce's records in the first hospital's database.

Privacy preservation adds significant challenges and computational overhead in solving ER as the data must be obfuscated to preserve privacy, making similarity comparisons difficult and costly. This is especially the case for fuzzy or approximate matching, which is required to account for differences in formats and data structures. In our case, using encryption techniques (like secure multi-party computation [42] or homomorphic encryption [1]) for

This work is licensed under a Creative Commons Attribution International 4.0 License.

CIKM '21, November 1–5, 2021, Virtual Event, Australia.

© 2021 Copyright is held by the owner/author(s).

ACM ISBN 978-1-4503-8446-9/21/11.

<https://doi.org/10.1145/3459637.3482318>

privacy preservation disallow the use of direct or equality comparisons as there is no support for such operators over *encrypted* records. Operators, in general, including arithmetic, bit, logical, and others are limited in these tools.

Contributions This paper presents **AMPPERE**, an abstract computation model for performing privacy-preserving entity resolution offering platform-agnostic, algorithmic abstractions using variants of Jaccard similarity [19]. The provided abstractions are *universal*: they support widely used arithmetic and vector operations and first-class support for private matrix manipulation. The abstraction model is presented as a pipeline that can be instantiated via simple privacy-preserving operations. We demonstrate the feasibility of **AMPPERE**'s abstraction model by transparently implementing it atop the secure multi-party computation platform Sharemind [4] and the open-source lattice-based homomorphic encryption library PALISADE [37]. We undertake a rigorous evaluation of our system which demonstrates the efficacy of our abstraction model, and a precise characterization of the induced overheads resulting from the *strong cryptographic protections* provided by **AMPPERE**.

Roadmap The paper is structured as follows: section 2.1 algorithmically defines the ER & privacy-preserving ER problems as well as evaluation metrics; the rest of section 2 introduces the technical machinery that form the algorithmic components of **AMPPERE**; section 3 describes and analyzes **AMPPERE**, the operators supported, and how it can be adapted to concrete implementations. Section 4 reports the settings, results, and findings of our experiments; section 5 lists prior approaches to PPER, and finally section 6 presents our conclusions and directions for further work.

2 ALGORITHMIC DEFINITION AND TOOLS

2.1 Problem definition

We formalize the **ER** problem as a triple $P = (D_1, D_2, M)$ where D_1 and D_2 are two different datasets consisting of a collection of unique records. M denotes the matching record pairs between the two datasets, that is, $M = \{(r_i, r_j) \mid r_i = r_j, r_i \in D_1, r_j \in D_2\}$ where $r_i = r_j$ indicates that r_i and r_j refer to the same entity in the real world. With **PPER**, we still deal with D_1 , D_2 , and M , however, we additionally say the probability that the data owners of D_1 , D_2 learn any information outside of what record pairs exist in M is negligible, and the probability that an adversary A learns anything, including M or any part of M , is negligible.

Additionally, we define a **full comparison** between D_1 and D_2 to be $T = \{(r_i, r_j) \mid r_i \in D_1, r_j \in D_2\}$, and the number of candidate pairs to be $|T|$, the Cartesian product between sizes of D_1 and D_2 (i.e. $|D_1| \times |D_2|$). We focus on token-based entity resolution algorithms for finding M , where each record is represented by a set of unique tokens $r = (t_1, t_2, \dots)$. **Blocking** algorithms prune T by removing pairs that are unlikely to be part of the solution M . *Inverted-index* blocking algorithms are popular methods that use a function β to encode each record into blocking keys $\beta(r) = (k_1, k_2, \dots)$, preserving *candidate record pairs* that share at least one k . Therefore, the pruned candidate pair set is defined as $T' = \{(r_i, r_j) \mid |\beta(r_i) \cap \beta(r_j)| > 0, r_i \in D_1, r_j \in D_2\}$. A good blocking algorithm reduces as many candidate pairs as possible without losing truly matched pairs [26]. We use **pairs completeness** (PC) to measure the percentage of true pairs that are blocked,

and **reduction ratio** (RR) to measure how well the method reduces the number of candidate pairs. These are formally defined as: $PC = \frac{M \cap T'}{M}$, $RR = 1 - \frac{|T'|}{|T|}$, and $F\text{-score} = \frac{2 * PC * RR}{PC + RR}$.

We now introduce our main algorithmic tools for entity resolution: the set similarity metric of Jaccard similarity [19] and its derivative MinHashLSH [24], which is a MinHash-based variant of Local Sensitive Hashing (LSH) designed for searching similar documents from a large corpus. We also introduce the two main cryptographic and privacy tools that we leverage in our implementations of **AMPPERE**: secure multi-party computation (MPC) [42] and homomorphic encryption (HE) [1].

2.2 Jaccard similarity

Jaccard similarity [19] is a well-known and widely used set-based similarity metric, commonly used for feature generation and similarity measurement in data mining and information retrieval tasks. Given two sets of tokens x and y , Jaccard similarity refers to the size of the intersection divided by the size of the union of the token-sets, that is, $Sim_{Jaccard}(x, y) = \frac{|x \cap y|}{|x \cup y|} = \frac{|x \cap y|}{|x| + |y| - |x \cap y|}$.

Some existing PPER approaches [34] extend the Jaccard similarity metric to privacy-preserving settings by encrypting input sets into bit arrays (or fingerprints), and comparing them using Tanimoto similarity [36], which computes similarity in a manner similar to Jaccard: $Sim_{Tanimoto}(x, y) = \frac{|x \wedge y|}{|x \vee y|} = \frac{|x \wedge y|}{|x| + |y| - |x \wedge y|}$, where $|x|$ and $|y|$ denotes the number of set bits (cardinality) in the bit array x and y respectively. We use Jaccard similarity to perform ER by selecting a threshold t such that $(r_1, r_2) \in M$, if $Sim_{Jaccard}(r_1, r_2) > t$.

2.3 MinHashLSH

MinHash exploits *min-wise independent* permutations to generate similarity-preserving signatures for each record. The probability of getting two similar MinHash signatures from two permutations of sets is identical to applying Jaccard directly over the them; formally, $Pr[h_{min}(perm(x)), h_{min}(perm(y))] = Sim_{Jaccard}(x, y)$ where h_{min} denotes MinHash and $perm$ denotes the random permutation function. The LSH algorithm then splits the signature into smaller bands (chunks) of a fixed length where only the records that share at least one band are considered to be potentially similar (i.e. they have partially common signatures). MinHashLSH Blocking [24] is an inverted-index inspired method utilizing MinHashLSH bands as blocking keys and associated record ids as values.

MinHashLSH is statistically robust and efficient for finding similar records or documents; since blocking keys are signatures, they do not reveal information relating to original records. For these reasons, we use it as our blocking method of choice in section 3.1.3.

2.4 Privacy tools

Secure multi-party computation (MPC) [42] ensures that a set of mutually distrusting parties involved in the same distributed computation correctly derive the result of the computation, even when a subset of parties are dishonest, without revealing any new information other than the result of computation. Formally, if the parties are $p_1, p_2, \dots, p_n$ and their private data is $d_1, d_2, \dots, d_n$, a MPC protocol defines the public function F which computes the

result over the private data, i.e., $F(d_1, d_2, \dots, d_3)$ while preserving the secrecy of $d_1, d_2, \dots, d_n$.

Homomorphic Encryption (HE) is a public-key cryptographic scheme that enables certain computations (e.g. addition, multiplication) to be performed directly over encrypted data, without a need for decryption [1]. Conceptually, it can be defined as an algorithmic quadruple of $(KeyGen, E, D, f)$: $KeyGen$ produces a public, private key pair (p_k, s_k) ; E is an encryption algorithm that uses p_k to encrypt a given plaintext p_i , outputting a corresponding ciphertext c_i ; f is an evaluation function to be directly evaluated over ciphertexts $c_1 \dots c_n$, outputting c_f representing the encrypted computation; D is a deterministic algorithm that takes in c_f , decrypts it using s_k , and outputs the result: $f(p_1 \dots p_n)$. The most important component of HE is its *homomorphic property*: suppose an f that operates on inputs p_k & $c_1 \dots c_n$, and outputs c_f . For all $p_1 \dots p_n$, it must hold that if $p_f = p_1 \circ p_2$, $c_1 = E(p_k, p_1)$, and $c_2 = E(p_k, p_2)$, then $Prob(D(f(s_k, c_1, c_2)) \neq p_f)$ is *negligible* (where $\circ$ refers to a homomorphic operation) [35].

Threshold HE Threshold homomorphic encryption combines the use of threshold cryptography [14] and HE to provide functionality via MPC. Here, n different parties interact, initially, to compute a joint public and private key, where the private key is split into n secret shares, one per party. Each party encrypts their own data using the public key, and the encrypted result is then decrypted, collectively, by some subset of parties assuming that a majority of them are not malicious [11].

In our two implementations of **AMPPERE**, we use Sharemind [13], a client-server platform for black-box general-purpose MPC computation, and PALISADE [37], an open-source lattice-based HE (& threshold-HE) library.

Figure 1: Abstraction model

3 AMPPERE

We propose **AMPPERE**, a platform-agnostic abstract machine for privacy-preserving entity resolution. We first introduce the specific challenges regarding PPER, present our computation model, and then two concrete implementations.

Challenges of PPER Privacy preservation adds another dimension to the entity resolution problem, making it much more challenging and computationally difficult. Keeping in mind our original example, hospital 1 & 2, contrary to regular ER approaches, cannot directly compare their records to determine that records for Peter and Bruce overlap between their databases. In the privacy-preserving setting, their records need to firstly be encrypted, and

distributed to a trusted computation host who can perform the data matching. Once encrypted, however, there is limited support in computing equality, similarity or joins, in terms of general-purpose functionality and scalability. Furthermore, one-way hashing techniques are rendered ineffective as they are susceptible to statistical-based attacks, and different data sources may have different representations making it hard to apply them.

3.1 Abstraction model

Our model is presented such that it can be instantiated using only simple platform-agnostic, privacy-preserving primitive operations. As shown in fig. 1, both dataset providers (i.e. P_1, P_2) negotiate on parameters first (e.g., Jaccard, blocking threshold), then apply encoding on each of their records, individually, while also computing blocks. Encoded records and blocks are then uploaded and stored in a cloud owned by a third party (i.e. P_3), where blocks are merged and deduplicated, and candidate pairs are filtered. ER (privately) is performed on the remaining qualified pairs. The accuracy, in performing PPER, of our abstractions and algorithms is 100% given correct implementations, as evidenced by section 3.2 & section 3.3.

Table 1: Abstraction Model Operators. (E) indicates the primitive has an element-wise version.

Category	Name
Primitive operators	(E){Add, Sub, Mul}, DotProduct, {L, R}Shift, Size, Transpose, Enc/Dec
Set intersection size	Pairwise join, Vector rotation, Vector extension, Sorting, Matrix join
Ternary operations	Choose{, Vec, VecExt, Mat, MatExt}
Private matrix manipulation	MaskGen, Vector{Lookup, Update}, Matrix{Lookup, Update}

3.1.1 Core operations. Our abstraction model relies on several core operations, listed in table 1, that should be defined first. The model assumes that the following **primitive operators** for vector and matrix data structures are available: Add for addition, Sub for subtraction, Mul for multiplication, and their corresponding element-wise versions: EAdd/ESub/EMul. Additionally, DotProduct for computing dot products, LShift/RShift for shifting vectors left and right, Size for getting the size of a vector or matrix, and Transpose for transposing a vector or matrix. Enc and Dec are encryption & decryption operators. Additional parameters (e.g. cryptographic keys, etc.) may be required, but are omitted here for simplicity. Optionally, logic and bit operators may be used wherever applicable, but we provide workarounds using only arithmetic operations.

The first essential step in computing Jaccard similarity is calculating **set intersection size**: the intersection size of two tokenized-record sets. We list and compare several possible approaches using the primitive operators mentioned above; choosing the appropriate approach depends on the implementation of the operations that it relies on. We use v_1, v_2 to indicate two input token-sets (vectors) that contain only unique entries.

- **Pairwise join (PJ)** Compare entries in v_1 and v_2 in a pairwise manner while accumulating the results into a counter.

Figure 2: Example of data flow. Original records with selected attributes after (1) pre-processing and tokenization to construct token sets. The token sets are then (2) ASCII-encoded, blocking keys are generated and both are encrypted and (3) sent to computational unit. Blocks with same keys are merged and (4) de-duplicated candidate pairs are generated. (5) Jaccard similarity is applied to the candidate pairs and (6) matches are found.

Since this method must be computed serially, it is inefficient and does not allow the use of optimized primitive operators.

- **Vector rotation (VR)** Fix the longer of the two input vectors (e.g. v_2), pad the smaller (e.g. v_1) with arbitrary, non-integer elements not a part of v_1 or v_2 until the sizes are equal (if necessary), and count the number of overlapping elements represented by zeroes in $v_1 - v_2$. This calculation must be done $Size(v_1)$ steps as each iteration requires a right shift of v_2 by one entry. For example, say $v_1 = [1, 2]$ and $v_2 = [2, 2, 3]$, then $v'_1 = [1, 2, *]$ and v_2 remains unchanged. Computing $v'_1 - v_2$ in the first iteration returns $[-1, 0, -3]$, giving us one 0, repeat this process for $Size(v_1)$ iterations. The total number of 0's represents the number of overlapping elements. Note, this method is an adaptation of basic PSI (private set intersection) protocol presented by Chen et al [7].
- **Vector extension (VE)** Extend v_1 to v'_1 by repeating the it $Size(v_2)$ times, and extend v_2 to v'_2 by repeating each element of v_2 $Size(v_1)$ times; the total number of 0's in $v'_1 - v'_2$ is the number of overlapping elements. For example, say $v_1 = [1, 2, 4]$ and $v_2 = [2, 3]$, then $v'_1 = [1, 2, 4, 1, 2, 4]$ and $v'_2 = [2, 2, 2, 3, 3, 3]$, hence, $v'_1 - v'_2 = [-1, 0, 2, -2, -1, 1]$ and the overlap is 1. This approach offers an alternative to VR (in terms of space vs. time complexity) by pre-allocating memory for all possible entry-wise comparisons and only needing one direct comparison (subtraction) between the two resulting vectors.
- **Sorting (SO)** Merge v_1 and v_2 into v_3 , sort v_3 , and then count overlaps between two derived vectors from v_3 : the first not including the first entry, and the second not including the last. For example, say $v_1 = [1, 2, 4]$, $v_2 = [2, 3]$, and $Sorted(v_3) = [1, 2, 2, 3, 4]$; then computing $[1, 2, 2, 3] - [2, 2, 3, 4] = [-1, 0, -1, -1]$ gives us an overlap of 1.
- **Matrix join (MJ)** Compute the join between v_1^T and v_2^T and count the size of the result.

Control flow statements are involved in almost all programs. Choice (e.g. if) and loop (e.g. for, while) based flows require conditional expressions for execution; in privacy-preserving settings, however, encrypted Boolean values cannot be used in conditional

Algorithm 1: Ternary Operators

Input: $cond$: bool, n_1, n_2 : numbers, v_1, v_2 : vectors, m_1, m_2 : matrices

- Function** Choose($cond, n_1, n_2$):
2 | **return** $cond * (n_1 - n_2) + n_2$;
- Function** ChooseVec($cond, v_1, v_2$):
4 | **return** EAdd(EMul($cond, ESub(v_1, v_2)$), v_2);
- Function** ChooseVecExt($cond, v_1, v_2$):
6 | $cond : vector_{Size(v_1)} = cond$;
7 | **return** ChooseVec($cond, v_1, v_2$);
- Function** ChooseMat($cond, m_1, m_2$):
9 | **return** EAdd(EMul($cond, ESub(m_1, m_2)$), m_2);
- Function** ChooseMatExt($cond, m_1, m_2$):
11 | $cond : matrix_{Size(m_1)} = cond$;
12 | **return** ChooseMat($cond, m_1, m_2$);

operations in order to avoid exposing the execution path, unless they are decrypted. **Ternary operations** can be designed to bypass this limitation: they take a condition followed by two values; if the condition is true, the first value is returned, otherwise the latter (e.g. if/else or $?:$). If the two values have the same data type and dimension, the return value can be determined by pure arithmetic. We design a set of oblivious-style ternary operators [9, 28] which work with encrypted conditions and inputs in various forms (numbers, vectors and matrices) in algorithm 1.

We also define functionality for getting & setting values for encrypted vectors and matrices. We refer to the set of these operations as **private matrix manipulation**, described in algorithm 2. To privately access vectors and matrices, arithmetic operations need to run over a special mask generated from an encrypted index (idx , rol , col). In MaskGen (Line 1-3), EEq element-wisely compares if seq and idx are equal, and sets the corresponding element in the result to be 1 if equal and 0 if not. This procedure generally uses a logical operator, but can also be achieved via arithmetic operations as follows: $(seq - idx) * \frac{-1}{seq - idx + \xi} + 1$ where ξ is a small offset added to prevent dividing by zero (seq, idx must be integers). If a

Algorithm 2: Private matrix manipulation

Input: *size, idx, row, col*: integers, *vec*: vector, *mat*: matrix, *val*: number

```

1 Function MaskGen(size, idx):
2   seq : [0, 1, ..., size], idx : vectorsize = idx;
3   return EEq(seq, idx);

4 Function VectorLookup(vec, idx):
5   mask ← MaskGen(Size(vec), idx);
6   return InnerProduct(vec, Transpose(mask));

7 Function VectorUpdate(vec, idx, val):
8   val : vectorSize(vec) = val;
9   mask ← MaskGen(Size(vec), idx);
10  return ChooseVec(mask, val, vec);

11 Function MatrixLookup(mat, row, col):
12  mask ← MaskGen(Size(mat)[0], row);
13  row_vec ← Mul(mask, mat);
14  return VectorLookup(row_vec, col);

15 Function MatrixUpdate(mat, row, col, val):
16  val : matrixSize(mat) = val;
17  row_mask ← MaskGen(Size(mat)[0], row);
18  col_mask ← MaskGen(Size(mat)[1], col);
19  mask ← Mul(Transpose(row_mask), col_mask);
20  return ChooseMat(mask, val, mat);

```

division operator is not available, an element n can be multiplied by a randomly-sampled encrypted plaintext element r , generated by P_1 or P_2 , i.e. $enc(r) * enc(nx) = enc(r * n)$. P_3 can decrypt this product and compute its reciprocal, $\frac{1}{r*n}$, in clear. Finally, it can encrypt this value and multiply with $enc(r)$ to output the reciprocal of n : $enc(\frac{1}{r*n}) * enc(r) = enc(\frac{1}{n})$.

3.1.2 Entity resolution. We model two datasets, D_1 and D_2 (e.g. $D_i = [Enc([t_1, t_2, \dots]), \dots]$), each containing encrypted records (e.g. $Enc(r) \in D_1$) where each record is formed by a set of tokens (e.g. $t \in r_1$). To represent records as token-sets, we use n -gram [3] tokenization, a general-purpose, fault-tolerant method to tokenize each record with ASCII encoding, represented in 8 bits. We choose a $n = 2$ (bi-gram) encoding, therefore each token occupies 16 bits allowing us to use any data type that has more than 16 bits as a container. In particular, we use the 64 bit container (e.g. `uint64`), so 4 encoded tokens to be placed in one, as 4 falls $\in [2, 8]$, the commonly-used range of n -grams [23].

The input for entity resolution, besides D_1 and D_2 , is a threshold t , and the output is a $|D_1| \times |D_2|$ sized encrypted, Boolean matrix *results* whose row-indices indicate record ids from D_1 , and column-indices indicate record ids from D_2 ; a cell's value indicates if the corresponding records are in the solution M . We use Jaccard similarity to compute how *similar* two records are based on their tokens, and only consider a match if the similarity is greater than t . We thereby define the similarity for a candidate record pair (r_1, r_2) to be $\frac{inter}{Size(r_1) + Size(r_2) - inter} + \epsilon > t$ where *inter* indicates the set intersection size of two token-sets and ϵ is a value added for tolerating the accuracy of floating point representation.

3.1.3 Blocking. Blocks are represented using a hashmap where each key is a blocking key, and each value is an encrypted vector

containing record ids; that is, $\{ "bkey1" : Enc([rid_1, rid_2, \dots]), \dots \}$. Blocks are generated by each party (P_1, P_2) individually, and are subsequently merged by P_3 if they have the same blocking key. Record ids in a block coming from different datasets are grouped into candidate pairs.

Algorithm 3 shows how blocks are used in our entity resolution pipeline. Lines 1-10 represent block merging and deduplication. Block deduplication refers to the process of determining which encrypted record id pairs have the same blocking key, but are from different datasets; this is done to avoid duplicated candidate pairs and duplicate comparisons. Our approach uses an encrypted Boolean matrix *cand_pairs* with the same dimensions & indices as *results*, where True values denote candidate pairs. This provides a simple solution to the difficult problem of constructing privacy-preserving inverted indices for deduplication [40]. VectorLookup retrieves record ids in encrypted form, and MatrixUpdate privately labels candidate pairs. In order to execute entity resolution on the candidate pairs, the program needs to get the record ids of each pair in clear. Even if there exists a method that is able to fetch records and invoke filtering with encrypted record ids, it is inevitable that an attacker could sniff memory or time the code block to detect if certain record pairs are faked. Our solution to not leak the total number of candidate pairs and certain combinations of record pairs is to obfuscate the candidate pair matrix before decryption. In Line 11, the *cand_pairs* matrix is converted to *obfu_pairs* by adding (union-ing) random truthy values with AddNoise before it is decrypted. The *results* matrix is then created along with the *dummies* matrix (Line 12-13). After applying filtering and ER (Line 14-17), the results are privately *chosen* from *results* or *dummies* via ChooseMatExt with the help of *cand_pairs* (Line 18), and the added noise is automatically eliminated. If a random number generator for AddNoise is not supported on P_3 , a random matrix can be pre-generated and encrypted by P_1 or P_2 for obfuscation.

3.1.4 Privacy analysis. This section analyzes the privacy properties of AMPPERE from end to end. Firstly, we note that AMPPERE provides cryptographic protections for privacy-preservation: the MPC-based implementation provides privacy of computation to all parties subject to a *majority* being corrupted [22] whereas the HE-based implementation provides privacy against a *honest-but-curious* adversary. Next, we note that P_1, P_2 first encode their input records and perform blocking on their end, and then send their corresponding encrypted data (records & record ids) to P_3 . Blocking keys are sent to P_3 in clear. The output from P_3 is an encrypted matrix where P_1, P_2 are responsible for finally decrypting it to retrieve the *matching* record pairs. P_3 needs the vector sizes to calculate the size of the set intersection, resulting in disclosure of the encoded record sizes. Block merging, deduplication, candidate pair generation and filtering use ternary & private matrix operations, all which need to invoke the Size primitive. Consequently, the total number of blocks, the size of each block, and the size of each dataset are disclosed. Therefore, after a complete run, P_1, P_2 learn the size of each dataset and the matched record id pairs only. P_3 learns the size of each dataset, the size of each encoded record, the size of each block and all blocking keys. No record content is disclosed and because of that no useful information is revealed to P_3 . Additionally, the blocking key is the partial MinHash signature,

Algorithm 3: PPER with blocking

Input: D_1 and D_2 are two datasets, b_1 and b_2 are two blocks, t is the Jaccard similarity threshold.

Output: *results* is a matrix.

```

// Merging and de-duplication
1 cand_pairs : matrix(Size( $D_1$ ),Size( $D_2$ )) = Enc(false) ;
2 for bkey  $\in b_1.keys$  do
3   if bkey  $\notin b_2.keys$  then
4     continue;
5    $r_1\_ids \leftarrow b_1[bkey], r_2\_ids \leftarrow b_2[bkey]$ ;
6   for  $r_1\_id \leftarrow 0$  to Size( $r_1\_ids$ ) do
7      $enc\_r1\_id = \text{VectorLookup}(r_1\_ids, r_1\_id)$ ;
8     for  $r_2\_id \leftarrow 0$  to Size( $r_2\_ids$ ) do
9        $enc\_r2\_id = \text{VectorLookup}(r_2\_ids, r_2\_id)$ ;
10       $cand\_pairs = \text{MatrixUpdate}(cand\_pairs,$ 
         $enc\_r1\_id, enc\_r2\_id, \text{Enc}(true))$ ;

// Obfuscation
11 obfu_pairs  $\leftarrow \text{Dec}(\text{AddNoise}(cand\_pairs))$  ;

// Filtering and entity resolution
12 results : matrix(Size( $D_1$ ),Size( $D_2$ )) = Enc(false) ;
13 dummies : matrix(Size( $D_1$ ),Size( $D_2$ )) = Enc(false) ;
14 for  $i \leftarrow 0$  to Size( $D_1$ ) do
15   for  $j \leftarrow 0$  to Size( $D_2$ ) do
16     if obfu_pairs[ $i, j$ ] then
17        $results[i, j] \leftarrow \text{Jaccard}(D_1[i], D_2[j], t)$ ;

18 results  $\leftarrow \text{ChooseMatExt}(cand\_pairs, results, dummies)$  ;

```

which is irreversible; the total number of blocks as well as the size of each block also do not provide any valuable information. Strictly, two blocks with the same key will be merged, but P_3 only knows the size of each block before merging. If two typical blocks with the same key k are b_1^k and b_2^k , the maximum possible matches could be found in merged block is $\min(|b_1^k|, |b_2^k|)$. Records often appear in multiple blocks, making it difficult for P_3 to estimate the number of matches. The total number of candidate pairs is also protected. Because of obfuscation, its value is somewhere between $|T'|$ to $|T|$, making it such that an attacker is not able to identify which pairs are noisy and which are real.

3.2 Sharemind implementation

SecreC [12] is a C-like domain-specific programming language designed for the Sharemind platform application [30]. SecreC provides generic programming facilities, hiding details of its underlying secure protocol; it has many differences as compared to general-purpose languages: (i) SecreC distinguishes between public and private data types at the system level via a security type (separate from the data type). The programmer must indicate whether certain data is public or private, and then the implementation can automatically convert one-way from public to private if necessary. (ii) Operations on private data types are limited. For example, *bool*, *int*, *uint* and *float* do not support bitwise and and bitwise or operators. (iii) Private values cannot be used in conditional statements. (iv) No file system access. The host machine and running environment can exchange data only via input arguments, via keydb, a key

value store based on Redis, or via a vector-based table_database; the data has to be batched and pre-uploaded through command line arguments to these containers.

Our implementation has two components: upload and link. upload loads encoded records and blocks to the Sharemind keydb instance and link is used for entity resolution. The key for a record is a public string in the form {dataset id}_{record id}, and the value is a private uint64 array which stores encoded tokens. The format of a block is similar to a record, but the key has an additional prefix as a namespace. For a record, if the associated blocking key exists, the program will fetch the corresponding vector back and merge the record id into it with the cat function;

The blocking keys belonging to D_1 are acquired from the keydb instance, and are constructed with the blocking keys with D_2 's identifier; the blocks belonging to different datasets, but with the same blocking key are fetched, candidate pairs are generated and deduplicated, and finally stored into *cand_pairs* with *matrixUpdate*. After this, the link program obfuscates the candidate pairs, and runs filtering and ER. Utilizing Sharemind's API, we implemented the following methods for computing set intersection size: Pairwise join, Vector extension (with internal function *gather_uint64_vec*), Sorting (with *quicksort*) and Matrix join (with *tableJoinAes128*). In the end, *results* is reshaped to a 1D array (as Sharemind does not allow matrices to be returned) and is shipped back to the host.

3.3 PALISADE implementation

Our PALISADE implementations work with the Brakersi-Gentry-Vaikuntanathan (BGV) [6] (for integer arithmetic) and the Cheon-Kim-Kim-Song (CKKS) [21] (for approximate numbers) constructions, in the form of their respective RNS (residue number system) variants [10][20]. As a first step in utilizing any scheme within PALISADE, one must instantiate a corresponding *CryptoContext* object for encryption functionality. Each context must be configured with relevant parameters that provide certain bit-security levels and performance; in our implementations we configured the following: *multiplicative depth*, *scale factor bits*, *plaintext modulus*, *batch size*, *key switching technique*, and *security level*. Both BGV and CKKS are *fully* homomorphic (FHE) in the sense that they both support (an arbitrary, unbounded number of) additive and multiplicative homomorphic operations [17]. In PALISADE, these schemes work as parametrized, *leveled* versions, meaning computations are computed up to a pre-determined depth. To establish a fair comparison between both schemes, we keep their corresponding ring dimensions and CRT moduli [2] the same. Table 2 displays our value choices for the parameters above; further fine-tuning can be done for potentially more efficient runtimes.

We developed 4 specific implementations in PALISADE: a two party protocol in both BGV and CKKS respectively, and their corresponding threshold (MPC) variants. In the two party protocol, P_2 encrypts their data, sends it to P_1 who then operates with P_2 's encrypted data, and their own data (in clear). In the MPC version, we apply threshold HE where P_1 and P_2 must first agree on a public/private key-pair to be used for encryption, and distributed decryption. Then, both parties encrypt and send their data to P_3 to carry out the computation over both of their (encrypted) data. Parties are modeled via individual data-structures, each referencing

Table 2: BGV & CKKS Parameters. (B) indicates parameter for blocking.

Parameter	BGV	CKKS
ring dim.	4096	4096, 8192 (B)
mult. depth	1	1, 2 (B)
scale factor bits	-	40
plaintext mod.	65537	-
batch size	-	16
key switching	BV	BV
security level	128 bits	128 bits

public/private keys generated via the KeyGen function and are realized as separate entities through different threads that take on the evaluation compute. We implemented the Pairwise join, Vector extension, and Vector rotation set intersection size approaches within PALISADE. PJ encodes each token of each record as a separate Ciphertext object with the value of the integer located at slot 0, whereas VR and VE are able to take advantage of *plaintext packing* (i.e. how many plaintexts you can fit into a single ciphertext slot) by encoding an entire record of tokens together as one ciphertext (Single Instruction Multiple Data (SIMD) computations). PJ computes the ER process in serial, and only relies on the PALISADE function EvalSub to homomorphically compare each token of one record, pairwise with each token of another record. VR and VE both utilize EvalAtIndex to homomorphically rotate (right shift) along with EvalSub for subtractions, while VE additionally uses EvalSum to repeat elements for a certain *batchSize* over the ring dimension.

We implemented blocking in the CKKS scheme. A discussion around why we chose CKKS vs. BGV for blocking is located in section 4.2.2. We used our own custom HE implementation of ternary operators (algorithm 1) and private matrix manipulation (algorithm 2) for deduplication and filtering (algorithm 3).

4 EXPERIMENTS

4.1 Datasets and settings

This section describes the dataset used for evaluation and the configurations used for our two implementations.

Febrl dataset Febrl is a synthetic dataset generated using *dsgen* [8], a tool to produce census records including a customizable amount of noise. We generated it using the following settings: Maximal number of duplicates for one original record: 5, Maximum number of modifications per field: 5, Maximum number of modifications per record: 5, Probability distribution for duplicates: zipf, Type of modification: typo, ocr, phonetic. Our dataset used for experiments contains 100 records, split into a 20% component for D_1 and 80% component for D_2 . Since our purpose of experimentation is to demonstrate the feasibility, not scalability of AMPPERE, we did not generate any larger datasets. We further processed our dataset using a Python script to create bi-gram, integer tokens for each record, and to compute MinHashLSH blocking keys for each record with thresholds ranging from 0.1 to 0.9 with step 0.1.

Expected performance To better understand the dataset, we created a Python script to apply the same blocking and ER methods in a non-privacy preserving setting, computing precision, recall, *PC*, *RR* and *F-score*. Figure 3 shows that ER precision and recall reach

Figure 3: Expected entity resolution and blocking results **Figure 4: Optimal MinHashLSH blocking key size**

1.0 when the threshold is in range 0.4 to 0.6. *PC* remains 1.0 when the threshold is $\leq$ to 0.5 and *RR* climbs sharply in that range.

We used Datasketch [44] to calculate MinHash signatures for records and to generate LSH blocking keys. We set the number of permutations to 128, and the relative importance of false positive/negative probability to 0.5 each. Datasketch calculates range and band to minimize the weighted sum of probabilities of false positives/negatives with all range and band combinations under a given threshold. The optimal blocking key sizes (the same as range) in different thresholds are shown in fig. 4. The blocking key size is correlated with the *b*: the blocking threshold.

Environment settings The experiments on our Sharemind and PALISADE implementations are conducted on three virtual Ubuntu 18.04.4 LTS servers each with 2 CPUs from Intel Xeon CPU E5-2690 v4 @ 2.60GHz and 4GB memory. All servers are in the same network and the average PING latency is around 0.12-0.23ms.

The Sharemind runtime environment [13] and implementations are deployed on all three servers, as for the SecreC program the shared3p private domain requires distributed across 3 different nodes. For the fairest comparison of evaluation results with Sharemind, we used the following configurations in PALISADE: compiled the PALISADE cmake project with the `-DWITH-NATIVEOPT` flag set to on; used single-threaded runtimes to establish baselines between systems, configured via `export OMP_NUM_THREADS=1`; turned CPU scaling off. We also disabled obfuscation (for both platforms) to achieve the most optimal running time possible.

4.2 Entity resolution performance

We tested the efficiency of executing Jaccard similarity with the different approaches for computing set intersection size described in section 3.1.1. For each approach, we recorded the execution time of Jaccard similarity on 100 randomly sampled single pairs.

4.2.1 Sharemind performance. In Sharemind, we adopted PJ, VE, SO and MJ respectively. Since PJ is extremely slow compared with other approaches, the time cost of all its runs have been divided by 5 in order to be shown properly in one diagram. Figure 5a shows that except for PJ, which takes on average 8k ms per record pair, the other three approaches have around the same performance: around 1.8k-4k ms per record pair. PJ executes the two loops serially and does not benefit from any parallelism of vector operations. VE achieved the best efficiency, likely a result of Sharemind's `gather_uint64_vec` internal method which saves DSL level operations. The shortcoming of VE, however, is that if the input records were long, overhead due to memory allocation could be a problem. The results of SO show more variance because quick sort is not a stable sorting algorithm; MJ is almost identical to SO. Notice that MJ uses Sharemind's `tableJoinAes128` function which requires the input records to be encoded in `xor_uint32` and this limitation halves the capacity of input containers (which is `uint64`). Overall, if the average length of records is not long, VE and MJ can be good choices, otherwise SO is better suited.

4.2.2 PALISADE performance. In PALISADE, we adopted the PJ, VR, and VE method in both BGV and CKKS. The runtime efficiency here is better than Sharemind's because these experiments are not distributed across three machines. We perform the experiments in two modes: sender-receiver or SR (2 party protocol) and three-party computation or 3PC (threshold variant). The VR and VE approaches both performed significantly faster than PJ, as reflected in fig. 5b and fig. 5c, with VR being the fastest of the three. VR is able to reach speeds of 125-175 ms per record pair for both BGV and CKKS in the sender-receiver variation and between 1.5 and 4 seconds in the corresponding threshold (3PC) version. VR and VE are much faster than PJ because they use *plaintext packing* as described in section 3.3. Additionally, VR is faster than VE because VR only relies on primitive operations (rotations and subtractions), while VE requires extending records using an expensive operation for repeating elements in encrypted form (`EvalSum`). Consistent with Sharemind, the PJ approach computes ER serially and does not benefit from any parallelization or packing. There are other ways to implement VE and VR in PALISADE that might perform faster or slower, depending on what you try to optimize (time or space). **BGV vs. CKKS Performance** BGV generally had equal or better runtime efficiency as compared to CKKS, which is expected. The main difference between the two schemes is that CKKS is an approximation algorithm that supports fast rescaling while dealing with fixed-precision arithmetic. BGV, depending on implementation, either does not support this type of rescaling or requires bootstrapping, which is a much more expensive operation. Since BGV works with exact numbers, when numbers are multiplied, they should not wrap around the *plaintext modulus*, whereas in CKKS, rescaling is supported after every multiplication. Since our non-blocking approaches did not require a *multiplicative depth* greater than 1 (because there was no need for sequential homomorphic multiplications), BGV becomes an ideal choice; however, when nested-multiplications are involved, BGV becomes quickly inefficient as its ciphertext modulus for depth k , must be at least p^k , where p is also a power of 2^k . CKKS is the ideal choice in these situations as it can rescale each time, keeping the parameters small.

Takeaway Determining the optimal approach for privacy-preserving entity resolution may vary with the computational platform and the concrete implementation, but our experiments above show many options are applicable and reasonably efficient.

4.3 Blocking performance

In this experiment, we show the effectiveness of blocking in performing (private) entity resolution with the two implementations. We used the non-blocking approach as the baseline (performs all 1,600 pairwise comparisons) and ran blocking with a MinHashLSH threshold t of $t=0.2, 0.5$, and 0.8 . The blocking key sizes corresponding to these threshold values are 28, 25 and 9 (resp.) and the numbers of distinct comparisons after blocking are limited to 1,314, 118 and 36 (resp.) as compared to the overall 1,600.

Figure 6b demonstrates the total and step time cost of PPER with blocking in Sharemind. The baseline takes around 6,000 seconds to make all 1,600 pair-wise similarity comparisons, whereas, for the blocking version, the dataset upload cost depends on the size of blocking keys. Hence, when $t = 0.8$, it takes the least amount of time to upload data as there are only 9 keys per record and block merging, deduplication and candidate pair generation costs are minor. In the entity resolution step, when $t = 0.2$, it needs 1,314 comparisons, which while smaller than the baseline, shows an overall longer runtime due to extra overhead incurred via data upload and block operations. When $t = 0.5$, the number of comparisons along with execution time drops drastically, and this trend continues when $t = 0.8$, but with recall loss as a trade-off cost.

As with Sharemind, blocking in PALISADE shows great time cost savings as it filters out unnecessary pairwise record comparisons. Figure 6c shows the overall time cost resulting in PALISADE with the various blocking thresholds against the non-blocking baseline. We implement blocking in both the SR and 3PC variants with CKKS as opposed to BGV due to the optimized performance in dealing with rescaling after multiplications, and because of good support for scalar/plaintext & ciphertext operations in implementing ternary operators presented in algorithm 1. Unlike Sharemind, we perform blocking in PALISADE we three main steps: preprocessing (generating block data structures), block merging + deduplication + candidate pair generation, and filtering + ER calculations. In fig. 6c, we group the runtime of steps 1 and 2 into one bar as, for all blocking threshold levels, preprocessing remains < 1 second. For both SR and 3PC variants, when $t = 0.2$, the overall time-cost is actually longer than the baseline (just like with Sharemind), because while filtering is applied through deduplication and candidate pair generation, it only removes around 300 overall comparisons, which does not outweigh the overhead incurred by those steps. When $t = 0.5, 0.8$, blocking and filtering is able to brunt much of the computation, by limiting the number of comparisons needed drastically, allowing for the ER calculation to be as minimal as possible.

Takeaway The incurred cost of privacy preservation is significant and evident as shown in fig. 6a. Filtering and ER takes between 0 and 0.013 seconds as compared to our PPER approaches: between 800 and 7,000 seconds in Sharemind, between 100 and 3,000 seconds in the SR variant, and between 200 and 7,000 seconds in the 3PC variant in PALISADE.

Figure 5: Time cost (per candidate pair) for PPER

Figure 6: Overall time cost for PPER with blocking

5 RELATED WORK

ER has been a well-studied area for many decades, from specific matching algorithms including Jaccard to semi/full automatic frameworks. PPER, however, has only been investigated recently. Vatsalan et al. [38, 39] provide a taxonomy of some privacy-preserving record linkage techniques, and suggest a general architecture. Some earlier work has explored [15, 29] exploiting secure one-way hashing and masking algorithms to protect the content of records, but this is limited as it does not allow for *similar* variants of records. Some algorithms that rely on certain probabilistic data structures (e.g. Bloom filters [27, 32]), though, do exist, and provide reasonable performance & privacy protection, while also maintaining relative similarity between records. Additionally, some record linking approaches depend on the modification of string matching algorithms (e.g. private edit distance [41, 45], P4Join [34]). With the increase of word/sentence embedding techniques, some methods including [5, 31] convert original records into multi-dimensional embedding vectors which still preserve some principle characteristics, and can be used in similarity comparison and clustering.

Very limited exploration [25], however, of MPC or HE based approaches exists today, and what does exist only handles the case of *exact matching* [16], due to expensive computation and lack of computational tool support. Currently, there is no general and comprehensive computational model like AMPPERE for privacy-preserving blocking and ER. In recent times, though, many excellent MPC frameworks and HE libraries including Microsoft SEAL [33], HELib [18], Obliv-C [43] have been developed and are becoming

more applicable for large-scale computing purposes. It is clear and evident that a general, precise, and *scalable* solution for ER based on these techniques will shine and break through in the near future.

6 CONCLUSION AND FUTURE WORK

We proposed an abstract machine for PPER and adopted it to two concrete implementations (Sharemind & PALISADE). Our experiments showed that both implementations could produce results with the exact or relatively equivalent correctness to traditional ER methodologies, while *maintaining* privacy. This lends credence to the possibilities of solving ER problems with MPC and HE.

Under this direction, many possible and promising future works can be explored. For instance, how to tune and find the optimal parameters for the matching algorithm without revealing data via a privacy-preserving adaptation of ground truth sampling and/or oracle annotation mechanisms. Moreover, data-set schema and metadata can also be utilized to further improve the accuracy and performance of matching results. And finally, scaling efficiently to larger and more complex data-sets and real-world scenarios is also a non-trivial problem that can be explored, but this will become increasingly computationally feasible as MPC and HE implementations become more efficient.

ACKNOWLEDGMENTS

This project was supported by the DataSafes program with funding from Actuate and the Alfred P. Sloan Foundation.

REFERENCES

- [1] A. Selcuk Uluagac & Mauro Conti Abbas Acar, Hidayet Aksu. 2018. A Survey on Homomorphic Encryption Schemes: Theory and Implementation. (2018). <https://dl.acm.org/doi/pdf/10.1145/3214303>
- [2] Martin Albrecht, Melissa Chase, Hao Chen, Jintai Ding, Shafi Goldwasser, Sergey Gorbunov, Shai Halevi, Jeffrey Hoffstein, Kim Laine, Kristin Lauter, Satya Lokam, Daniele Micciancio, Dustin Moody, Travis Morrison, Amit Sahai, and Vinod Vaikuntanathan. 2019. Homomorphic Encryption Standard. Cryptology ePrint Archive, Report 2019/939. <https://eprint.iacr.org/2019/939>.
- [3] Satanjeev Banerjee and Ted Pedersen. 2003. The design, implementation, and use of the ngram statistics package. In *International Conference on Intelligent Text Processing and Computational Linguistics*. Springer, 370–381.
- [4] D. Bogdanov. 2008. Sharemind: A framework for fast privacy-preserving computations. *ES-ORJCS 2008* (2008), 192–206. <https://ci.nii.ac.jp/naid/20001634171/en/>
- [5] Luca Bonomi, Li Xiong, Rui Chen, and Benjamin CM Fung. 2012. Frequent grams based embedding for privacy preserving record linkage. In *Proceedings of the 21st ACM international conference on Information and knowledge management*. 1597–1601.
- [6] Zvika Brakerski, Craig Gentry, and Vinod Vaikuntanathan. 2011. Fully Homomorphic Encryption without Bootstrapping. Cryptology ePrint Archive, Report 2011/277. <https://eprint.iacr.org/2011/277>.
- [7] Hao Chen, Kim Laine, and Peter Rindal. 2017. Fast Private Set Intersection from Homomorphic Encryption. In *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security* (Dallas, Texas, USA) (CCS '17). Association for Computing Machinery, New York, NY, USA, 1243–1255. <https://doi.org/10.1145/3133956.3134061>
- [8] Peter Christen. 2005. Probabilistic data generation for deduplication and data linkage. In *International Conference on Intelligent Data Engineering and Automated Learning*. Springer, 109–116.
- [9] Scott D Constable and Steve Chapin. 2018. libOblivious: A c++ library for oblivious data structures and algorithms. (2018).
- [10] Nigel P. Smart Craig Gentry, Shai Halevi. 2012. Homomorphic Evaluation of the AES Circuit. (2012). https://link.springer.com/chapter/10.1007%2F978-3-642-32009-5_49
- [11] Ronald Cramer, Ivan Damgård, and Jesper B. Nielsen. 2001. Multiparty Computation from Threshold Homomorphic Encryption. In *Advances in Cryptology — EUROCRYPT 2001*, Birgit Pfitzmann (Ed.). Springer Berlin Heidelberg, Berlin, Heidelberg, 280–300.
- [12] Cybernetica. 2019. *SecreC (version 2019.03)*. <https://sharemind-sdk.github.io/stdlib/reference>
- [13] Cybernetica. 2019. *Sharemind SDK and Academic server (version 2019.03)*. <https://sharemind.cyber.ee/sharemind-mpc>
- [14] Yvo Desmedt. 2011. *Threshold Cryptography*. Springer US, Boston, MA, 1288–1293. https://doi.org/10.1007/978-1-4419-5906-5_330
- [15] L Dusserrer, C Quantin, and H Bouzelat. 1995. A one way public key cryptosystem for the linkage of nominal files in epidemiological studies. *Medinfo. MEDINFO 8* (1995), 644–647.
- [16] Fatih Emekci, Divyakant Agrawal, Amr E Abbadi, and Aziz Gulbeden. 2006. Privacy preserving query processing using third parties. In *22nd International Conference on Data Engineering (ICDE'06)*. IEEE, 27–27.
- [17] Craig Gentry. 2009. A FULLY HOMOMORPHIC ENCRYPTION SCHEME. (2009). <https://crypto.stanford.edu/craig/craig-thesis.pdf>
- [18] Shai Halevi. 2021. *HElib (version 2.1.0)*. <https://github.com/homenc/HElib>
- [19] Paul Jaccard. 1901. Étude comparative de la distribution florale dans une portion des Alpes et des Jura. *Bull Soc Vaudoise Sci Nat* 37 (1901), 547–579.
- [20] Yongsoo Song Jung Hee Cheon & Kyoohyung Han, Andrey Kim & Miran Kim. 2019. A Full RNS Variant of Approximate Homomorphic Encryption. (2019). https://link.springer.com/chapter/10.1007%2F978-3-030-10970-7_16
- [21] Yongsoo Song Jung Hee Cheon, Andrey Kim & Miran Kim. 2017. Homomorphic Encryption for Arithmetic of Approximate Numbers. (2017). https://link.springer.com/chapter/10.1007%2F978-3-319-70694-8_15
- [22] Jonathan Katz and Yehuda Lindell. 2014. *Introduction to Modern Cryptography, Second Edition*. CRC Press. <https://www.crcpress.com/Introduction-to-Modern-Cryptography-Second-Edition/Katz-Lindell/p/book/9781466570269>
- [23] Gregory W Leshner, Bryan J Moulton, D Jeffery Higginbotham, et al. 1999. Effects of ngram order and training text size on word prediction. In *Proceedings of the RESNA '99 Annual Conference*. Citeseer, 52–54.
- [24] Jure Leskovec, Anand Rajaraman, and Jeffrey David Ullman. 2014. *Mining of Massive Datasets* (2nd ed.). Cambridge University Press, USA.
- [25] Yehida Lindell. 2005. Secure multiparty computation for privacy preserving data mining. In *Encyclopedia of Data Warehousing and Mining*. IGI global, 1005–1009.
- [26] Matthew Michelson and Craig A. Knoblock. 2006. Learning Blocking Schemes for Record Linkage. In *Proceedings of the 21st National Conference on Artificial Intelligence - Volume 1* (Boston, Massachusetts) (AAAI'06). AAAI Press, 440–445. <http://dl.acm.org/citation.cfm?id=1597538.1597609>
- [27] Frank Niedermeyer, Simone Steinmetzer, Martin Kroll, and Rainer Schnell. 2014. Cryptanalysis of basic bloom filters used for privacy preserving record linkage. *German Record Linkage Center, Working Paper Series, No. WP-GRLC-2014-04* (2014).
- [28] Olga Ohrimenko, Felix Schuster, Cédric Fournet, Aastha Mehta, Sebastian Nowozin, Kapil Vaswani, and Manuel Costa. 2016. Oblivious multi-party machine learning on trusted processors. In *25th {USENIX} Security Symposium ({USENIX} Security 16)*. 619–636.
- [29] Catherine Quantin, Hocine Bouzelat, FAA Allaert, Anne-Marie Benhamiche, Jean Faivre, and Liliane Dusserrer. 1998. How to ensure data security of an epidemiological follow-up: quality assessment of an anonymous record linkage procedure. *International journal of medical informatics* 49, 1 (1998), 117–122.
- [30] Jaak Randmets et al. 2017. *Programming Languages for Secure Multi-party Computation Application Development*. Ph.D. Dissertation.
- [31] Monica Scannapieco, Ilya Figotin, Elisa Bertino, and Ahmed K Elmagarmid. 2007. Privacy preserving schema and data matching. In *Proceedings of the 2007 ACM SIGMOD international conference on Management of data*. 653–664.
- [32] Rainer Schnell, Tobias Bachteler, and Jörg Reiher. 2009. Privacy-preserving record linkage using Bloom filters. *BMC medical informatics and decision making* 9, 1 (2009), 1–11.
- [33] SEAL. 2020. Microsoft SEAL (release 3.6). <https://github.com/Microsoft/SEAL>. Microsoft Research, Redmond, WA.
- [34] Ziad Sehili, Lars Kolb, Christian Borgs, Rainer Schnell, and Erhard Rahm. 2015. Privacy preserving record linkage with PPJoin. *Datenbanksysteme für Business, Technologie und Web (BTW 2015)* (2015).
- [35] Jaydip Sen. 2013. Homomorphic Encryption: Theory & Application. (2013). <https://arxiv.org/pdf/1305.5886.pdf>
- [36] Taffee T Tanimoto. 1958. Elementary mathematical theory of classification and prediction. (1958).
- [37] PALISADE team. 2020. *PALISADE Lattice Cryptography Library (release 1.10.6)*. <https://palisade-crypto.org>
- [38] Dinusha Vatsalan, Peter Christen, and Vassilios S Verykios. 2013. A taxonomy of privacy-preserving record linkage techniques. *Information Systems* 38, 6 (2013), 946–969.
- [39] Dinusha Vatsalan, Ziad Sehili, Peter Christen, and Erhard Rahm. 2017. Privacy-preserving record linkage for big data: Current approaches and research challenges. In *Handbook of Big Data Technologies*. Springer, 851–895.
- [40] Bing Wang, Wei Song, Wenjing Lou, and Y Thomas Hou. 2015. Inverted index based multi-keyword public-key searchable encryption with strong privacy guarantee. In *2015 IEEE Conference on Computer Communications (INFOCOM)*. IEEE, 2092–2100.
- [41] Xiao Shaun Wang, Yan Huang, Yongan Zhao, Haixu Tang, XiaoFeng Wang, and Diyu Bu. 2015. Efficient genome-wide, privacy-preserving similar patient query based on private edit distance. In *Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security*. 492–503.
- [42] Andrew C Yao. 1982. Protocols for secure computations. In *23rd annual symposium on foundations of computer science (sfcs 1982)*. IEEE, 160–164.
- [43] Samee Zahur and David Evans. 2015. Obliv-C: A Language for Extensible Data-Oblivious Computation. *IACR Cryptol. ePrint Arch.* 2015 (2015), 1153.
- [44] Eric Zhu. 2017. *Datasketch (version 1.2.5)*. <https://github.com/ekzhu/datasketch>
- [45] Ruiyu Zhu and Yan Huang. 2017. Efficient privacy-preserving general edit distance and beyond. *IACR Cryptology ePrint Archive* 2017 (2017), 683.