

# Fast Database Joins for Secret Shared Data

Payman Mohassel, Peter Rindal, Mike Rosulek

May 17, 2019

## Abstract

We present a scalable database join protocol for secret shared data in the honest majority three party setting. The key features of our protocol are a rich set of SQL-like join/select queries and the ability to compose join operations together due to the inputs and outputs being generically secret shared between the parties. Given that the keys being joined on are unique, no information is revealed to any party during the protocol. In particular, not even the sizes of intermediate joins are revealed. All of our protocols are constant-round and achieve  $O(n)$  communication and computation overhead for joining two tables of  $n$  rows.

In addition to performing database joins our protocol, we implement two applications on top of our framework. The first performs joins between different governmental agencies to identify voter registration errors in a privacy-preserving manner. The second application considers the scenario where several organizations wish to compare network security logs to more accurately identify common security threats, e.g. the IP addresses of a bot net. In both, cases the practicality of these applications depends on efficiently performing joins on millions of secret shared records. For example, our three party protocol can perform a join on two sets of 1 million records in 4.9 seconds or, alternatively, compute the cardinality of this join in just 3.1 seconds.

## 1 Introduction

We consider the problem of performing SQL-style join operations on secret shared tables with three parties and an honest majority. In particular, the proposed protocol takes two or more arbitrarily secret shared database tables and constructs another secret shared table containing a join of the two tables, without revealing *any* information. Our protocol is constant round and has  $O(n)$  computation and communication overhead to join two tables with  $n$  records. Simulation-based security is achieved in the semi-honest setting. Our protocol can perform inner, left and full joins along with union and arbitrary circuit computation on the resulting table. A central requirement in achieving high efficiency and no leakage is that the join-keys must be unique.

New techniques [PSZ14, PSSZ15, PSZ16, KKRT16, PSWW18, CLR17, CHLR18, IKN<sup>+</sup>17, RA18, KLS<sup>+</sup>17, OOS17, KMP<sup>+</sup>17] for performing set intersection, inner join and related functionalities have shown great promise for practical deployment. To name a few, Ion et al. at Google recently deployed a private set intersection sum protocol [IKN<sup>+</sup>17] to allow customers of Google Adwords to correlate the impact online advertising has on generating offline sales while preserving user privacy. Pinkas et al. [PSWW18] also introduced a practical protocol that can compute any (symmetric) function of the intersection and associated data. In both cases these protocols can be framed in terms of SQL queries where an inner join is computed followed by an aggregation on the resulting table, e.g. summing a column.

The majority of these protocols consider the two party setting and are based on various cryptographic primitives, e.g. exponentiation [IKN<sup>+</sup>17], oblivious transfer [PSWW18], or fully homomorphic encryption [CLR17]. However, in this work we alter the security model to consider three parties

with an honest majority. The motivation is that typical protocols in this setting (e.g.[AFL+16]) require less computation and communication than similar two party protocols, by a factor of the security parameter  $\kappa = 128$ .

Given this observation we investigate how to leverage the efficiency gains in the three party setting to construct practical protocols for performing set intersection and other SQL like operations where both the inputs and outputs are secret shared. One critical aspect of this input/output requirement is that join operations can then be *composed* together, where the output of a join can be the input to another. Allowing this composability greatly increases the ability to perform highly complex queries and enables external parties to contribute data simply by secret sharing it between the primary parties which participate in the protocol.

## 1.1 Functionality

Our protocol offers a wide variety of functionality including set intersection, set union, set difference and a variety of SQL-like joins with complex boolean queries. Generally speaking, our protocol works on tables of secret shared data which are functionally similar to SQL tables. This is contrasted by traditional PSI and PSU protocols[PSZ14, PSSZ15, PSZ16, KKRT16] in that each record is now a tuple of values as opposed to a single value.

We define our database tables in the natural way. Each table can be viewed as a collection of rows or as a vector of columns. For a table  $X$ , we denote the  $i$ th row as  $X[i]$  and the  $j$ th column as  $X_j$ , for  $j = 1, 2, \dots$ . Note that each column of a table has the same length but can contain different data types, e.g.  $X_1$  is a column of 32 bit fixed point decimal values and  $X_2$  is a column of 1024 bit strings.

Our core protocol requires each table to contain unique values in the column defining the join (i.e., we can only join on “unique primary keys”). For example, if we consider the following SQL styled join/intersection query

```
select X2 from X inner join Y on X1 = Y1
```

then the join-keys are  $X_1$  and  $Y_1$ . This uniqueness condition can be extended to the setting where multiple columns are being compared for equality. In this case the tuple of these columns must be unique. Later on we will discuss the case when such a uniqueness property does not hold. Our protocols also support a **where** clause that filters the selection using an arbitrary predicate of the  $X$  and  $Y$  rows. Furthermore, the **select** clause can also return a function of the two rows. For example,

```
select X1, max(X2, Y2) from X inner join Y
on X1 = Y1 where Y2 > 23.3
```

In general, the supported join operations can be characterized in three parts: 1) The select function  $S(\cdot)$  that defines how the rows of  $X, Y$  are used to construct each output row, e.g.  $S(X, Y) = (X_1, \max(X_2, Y_2))$  in the example above. 2) The predicate  $P(\cdot)$  that defines the **where** clause, and 3) which columns are being joined on. We require that both  $S$  and  $P$  must be expressible in the framework of [MR18] which we provide a custom implementation of.

So far we have discussed inner join (intersection) between two tables. Several other types of joins are also supported including left and right joins, set union and set minus (difference) and full joins. A left join takes the inner join and includes all of the missing records from the left table. For the records solely from the left table, the resulting table contains NULL for the columns from

the right table. Right join is defined symmetrically. A full join is a natural extension where all the missing rows from  $X$  and  $Y$  are added to the output table.

We define the union of two tables to contain all records from the left table, along with all the records from the right table which are not in the intersection with respect to the join-keys. Note that this definition is not symmetric with respect to the left and right tables. Table minus is similarly defined as all of the left table whose join-column value is not present in the right table.

Beyond these various join operations, our framework supports two broad classes of operations which are a function of a single table. The first is a general SQL select statement which can perform computation on each row (e.g. compute the max of two columns) and filter the results using a `where` clause predicate. The second class is referred as an aggregation which perform an operation across all of the rows of a table. For example, computing the sum, counts or the max of a given column.

## 1.2 Our Results

We present the first practical secure multi-party computation protocol for performing SQL styled database joins with linear overhead and constant round. Our protocol is fully composable in that the input and output tables are generically secret shared between the parties. No partial information is revealed at any point. We achieve this result by combining various techniques from private set intersection and secure computation more broadly while requiring the joined on keys be unique. We build on the the binary secret sharing technique of [AFL<sup>+</sup>16] with enhancements described by [MR18]. We then combine this secret sharing scheme with cuckoo hashing[PSZ14], an MPC friendly PRF[ARS<sup>+</sup>15] and a custom protocol for evaluating an oblivious switching network[MS13]. Using these building blocks our protocol is capable of computing the intersection of two tables of  $n = 2^{20}$  rows in 4.9 seconds. Alternatively, the cardinality of the intersection can be computed in just 3.1 seconds. Beyond these two specific functionalities, our protocol allows arbitrary computation applied to a shared table. Compared to existing three party protocols with similar functionality (composable), our implementation is roughly 1000× faster. When compared with *non-composable* two party protocol, we observe a larger difference ranging from our protocol being 1.25× slower to 4000× faster depending on what functionality is being computed.

Building on our proposed protocol we demonstrate it’s utility by showcasing two potential applications. The first prototype would involve running our protocol between and within the states of the United States to validate the accuracy of the voter registration data in a privacy preserving way. The Pew Charitable Trust[Smi14] reported 1 in 8 voter registration records in the United States contains a serious error while 1 in 4 eligible citizens remain unregistered. Our privacy preserving protocol identifies when an individual’s address is out of date or more seriously if someone is registered to vote in more than one state which could allow them to cast two votes. Additionally, our protocol can help register eligible citizens. Our protocol ensures that only the minimum amount of information is revealed, namely the identities of individuals with serious registration errors. Due to how the data is distributed between different governmental agencies, it is critical that our protocol allows for composable operations. We implement this application and demonstrate that it is practical to run at a national scale (quarter billion records) and low cost.

The second application that we consider allows multiple organizations to compare computer security incidents and logs to more accurately identify unwanted activities, e.g. a bot net. Several companies already offer this service including Facebook’s ThreatExchange[thr18] and an open source alternative[alt18]. One of the primary limitations of these existing solutions is the requirement that each organization send their security logs to a central party, e.g. Facebook. We propose using our protocol to distribute the trust of this central party between three (or more) parties such

that privacy is guaranteed so long as there is an honest majority.

### 1.3 Related Work

We now review several related works that use secure computation techniques.

With respect to functionality the closest related work is that of Blanton and Aguiar[BA12] which describes a relatively complete set of protocols for performing intersections, unions, set difference, etc. and the corresponding SQL-like operations. Moreover, these operations are composable in that the inputs and outputs are secret shared between the parties. At the core of their technique is the use of a generic MPC protocol and an oblivious sorting algorithm that merges the two sets. This is followed by a linear pass over the sorted data where a relation is performed on adjacent items. Their technique has the advantage of being very general and flexible. However, the proposed sorting algorithm has complexity  $O(n \log^2 n)$  and is not constant round<sup>1</sup>. As a result, the implementation from 2011 performed poorly by current standards, intersecting  $2^{10}$  items in 12 seconds. The modern protocol [KKRT16] which is *not composable* can perform intersections of  $2^{20}$  items in 4 seconds. While this difference of three orders of magnitude would narrow if reimplemented using modern techniques, the gap would remain large.

Huang, Evans and Katz[HEK12] also described a set intersection protocol based on sorting. Unlike [BA12], this work considers the two party setting where each party holds a set in the clear. This requirement prevents the protocol from being composable but allows the complexity to be reduced to  $O(n \log n)$ . The key idea is that each party locally sorts their set followed by merging the sets within MPC. The protocol can then perform a single pass over the sorted data to construct the intersection. While this results in performance improvements the overall protocol requires  $O(n \log n)$  operations and is not composable.

Another line of work was begun by Kissner and Song[KS05] and improved on by [MF06]. Their approach is based on the observation that set intersection and multi-set union have a correspondence to operations on polynomials. A set  $S$  can be encoded as the polynomial  $\hat{S}(x) = \prod_{s \in S} (x - s) \in \mathbb{F}[x]$ . That is, the polynomial  $\hat{S}(x)$  has a root at all  $s \in S$ . Given two such polynomials,  $\hat{S}(x), \hat{T}(x)$ , the polynomial encoding the intersection is  $\hat{S}(x) + \hat{T}(x)$  with overwhelming probability given a sufficiently large field  $\mathbb{F}$ .

Multi-set union can similarly be performed by multiplying the two polynomials together. Unlike with normal union, if an item  $y$  is contained in  $S$  and  $T$  then  $\hat{S}(x)\hat{T}(x)$  will contain two roots at  $y$  which is often not the desired functionality. This general idea can be transformed into a secure multi-party protocol using oblivious polynomial evaluation[NP99] along with randomizing the result polynomial. The original computational overhead was  $O(n^2)$  which can be reduced to the cost of polynomial interpolation  $O(n \log n)$  using techniques from [MF06]. The communication complexity is linear. In addition, this scheme assumes an ideal functionality to generate a shared Paillier key pair. We are unaware of any efficient protocol to realize this functionality except for [HMRT12] in the two party setting.

This general approach is also composable. However, due to randomization that is performed the degree of the polynomial after each operation doubles. This limits the practical ability of the protocol to compose more than a few operations. Moreover, it is not clear how this protocol can be extended to support SQL-like queries where elements are key-value tuples. This general approach is also composable but incurs a 2 times overhead for each successive operations and cannot be extended to SQL-like queries.

---

<sup>1</sup>It is not constant round when the underlying MPC protocol is not constant round which is typically required for high throughput.

Hazay and Nissim[HN12] introduce a pair of protocols computing set intersection and union which are also based on oblivious polynomial evaluation where the roots of the polynomial encode a set. However, these protocols are restricted to the two party case and are not composable. The non-composability comes from the fact that only one party constructs a polynomial  $\hat{S}(x)$  encoding their set  $S$  while the other party obliviously evaluates it on each element in their set. The result of these evaluations are compared with zero<sup>2</sup>. These protocols have linear overhead and can achieve security in the malicious setting.

Pinkas, Schneider and Zohner [PSZ14] introduced a paradigm for set intersection that combines a hash table technique known as cuckoo hashing with a randomized encoding technique using oblivious transfer. Due to the hashing technique, the problem is reduced to comparing a single item  $x$  to a small set  $\{y_1, \dots, y_m\}$ . Oblivious transfer is then used to interactively compute the randomized encoding  $\llbracket x \rrbracket$  while the other party locally computes the encodings  $\{\llbracket y_1 \rrbracket, \dots, \llbracket y_m \rrbracket\}$ . A plaintext intersection can then be performed directly on these encodings. With the use of several optimizations[PSSZ15, PSZ16, KKRT16, OOS17] this paradigm is extremely efficient and can perform a set intersection using  $O(n)$  calls to a random oracle and  $O(n)$  communication. These protocols are not composable. More recently this approach has also been extended to the malicious setting [RR17] and separately to have sublinear communication when one set is much larger than the other[CLR17].

Laur, Talvita and Willemson[LTW13] present techniques in the honest majority setting for composable joins, unions and many other operations at the expense of information leakage. Consider two parties each with a sets  $X, Y$ . The parties first generate secret shares of the sets and then use a generic MPC protocol to apply a pseudorandom function (PRF)  $F$  to the shared sets to compute  $X' = \{F_k(x) \mid x \in X\}, Y' = \{F_k(y) \mid y \in Y\}$  where the key  $k$  is uniformly sampled by the MPC protocol (i.e. neither party knows  $k$ ).  $X'$  and  $Y'$  are then revealed to both parties who use this information to infer the intersection, union and many other SQL-like operations. This basic approach dates back to the first PSI protocols [Mea86, HFH99] where the (oblivious) PRF was implemented using a special purpose Diffie-Hellman protocol. [LTW13] extended this paradigm to allow the input sets to be secret shared as opposed to being known in the clear.

The primary limitation of this approach is that all operations require all parties to know  $X'$  and  $Y'$ . This prevents the protocol from being composable without significant information leakage. In particular, the cardinality of  $X' \cap Y'$  and the result of the **where** clause for each row is revealed. This is of particular concern when several datasets are being combined. Learning the size of the intersection or the union can represent significant information. For instance, in the threat log application the union of many sets are taken. Each of these unions would reveal how many unique logs the new set has. Alternatively, taking the join between a set of hospital patients and a set of HIV positive people would reveal how many patients have HIV. When combined with other information it could lead to the ability to identify some or all of these patients. Beyond this, the provided three party implementation achieved relatively poor performance. A join between two tables of a million records is estimated to require one hour on their three benchmark machines[LTW13]. Looking forward, our protocol can perform a similar join operation in 4 seconds while preventing all information leakage.

---

<sup>2</sup>The real protocol is slightly more complicated than this.

## 2 Preliminaries

### 2.1 Security Model

Our protocols are presented in the semi-honest three-party setting with an honest majority. That is, our protocols are computationally secure conditioned on the adversary corrupting at most one of the three parties. See [AFL<sup>+</sup>16, MR18] for details.

Throughout the exposition we will assume these three parties, which we sometimes call *servers*, provide the sets which are computed on. However, in the general case the sets being computed on can be privately input by an arbitrary party which does not participate in the computation. These *client* parties will secret share their set between the servers and reconstruct the output shares that are intended for them. This setting is often referred to as the client server model [MR18, MZ17]. Later on we will also consider a setting with five servers that can tolerate the adversary corrupting any two of them. However, we will explicitly state when this alternative model is being considered.

### 2.2 Notation

Let  $[m]$  denote the set  $\{1, 2, \dots, m\}$ . Vector indices start at 1. We define a permutation of size  $m$  as an injective function  $\pi : [m] \rightarrow [m]$ . We extend this definition such that when  $\pi$  is applied to a vector  $V$  of  $m$  elements, then  $\pi(V) = \{V_{\pi(1)}, \dots, V_{\pi(m)}\}$ . The image of a function  $f : X \rightarrow Y$  is defined as  $image(f) := \{y \in Y : \exists x \in X, f(x) = y\}$ . Preimage of a pair  $(f, y)$  is defined as  $preimage(f, y) := \{x \in X : f(x) = y\}$ . We use  $n$  to represent the number of rows a table has. Parties are referred to as  $P_0, P_1, \dots, P_{N-1}$ .

### 2.3 Secret Sharing Framework

Our protocol builds on the ABY<sup>3</sup> framework of Rindal and Mohassel [MR18] for secure computation of circuits. That is, we use their binary/arithmetic addition and multiplication protocols along with their share conversion protocols. We will use the notation that  $\llbracket x \rrbracket$  is a 2-out-of-3 *binary replicated secret sharing* of the value  $x$ . That is,  $(x_0, x_1, x_2)$  are sampled uniformly s.t.  $x = x_0 \oplus x_1 \oplus x_2$ . Party  $P_i$  holds the shares  $x_i, x_{i+1 \bmod 3}$ . We use the notation  $\llbracket x \rrbracket_i$  to refer to share  $x_i$ .  $\llbracket x \rrbracket$  can locally be converted to a 2-out-of-2 sharing  $\langle\langle x \rangle\rangle$  where  $P_i$  holds  $x'_0$  and  $P_j$  holds  $x'_1$  s.t.  $x = x'_0 \oplus x'_1$ , e.g.  $i = 0, j = 1$ .  $\langle\langle x \rangle\rangle_k$  refers to  $x'_k$ .  $\langle\langle x \rangle\rangle$  can also be converted back to  $\llbracket x \rrbracket$  using one round of communication.

Note that the framework of [MR18] provides efficient facilities to convert between different types of shares, e.g. binary, arithmetic and fixed point.

### 2.4 Cuckoo Hash Tables

The core data structure that our protocols employ is a cuckoo hash table which is parameterized by a capacity  $n$ , two (or more) hash functions  $h_0, h_1$  and a vector  $T$  which has  $m = O(n)$  slots,  $T[1], \dots, T[m]$ . For any  $x$  that has been added to the hash table, there is an invariant that  $x$  will be located at  $T[h_0(x)]$  or  $T[h_1(x)]$ . Testing if an  $x$  is in the hash table therefore only requires inspecting these two locations.  $x$  is added to the hash table by inserting  $x$  into slot  $T[h_i(x)]$  where  $i \in \{0, 1\}$  is picked at random. If there is an existing item at this slot, the old item  $y$  is removed and reinserted at its other hash function location. See [DRRT18] for details. Given a hash table with  $m \approx 1.6n$  slots and three hash functions, then with overwhelming probability  $n$  items can be inserted using  $O(n)$  insertions [DRRT18]. For technical reasons we require  $h_i(x) \neq h_j(x)$  for all  $x$  and  $i \neq j$ . This can be achieved by defining  $h_j(x)$  over the range  $[m] \setminus \{h_i(x)\}_{i < j}$ .

## 3 Our Construction

### 3.1 Overview

Our core protocol is a technique for obliviously mapping together rows with equal join-keys. For the  $i$ th secret shared row  $\llbracket X \rrbracket[i]$  and for some  $j_0, j_1$ , our protocol obliviously maps rows  $\llbracket Y[j_0] \rrbracket, \llbracket Y[j_1] \rrbracket$  to the  $i$ th row of two new tables  $\llbracket \hat{Y}^0 \rrbracket, \llbracket \hat{Y}^1 \rrbracket$  such that if  $Y$  contains a row with matching keys then either  $\llbracket \hat{Y}^0 \rrbracket[i]$  or  $\llbracket \hat{Y}^1 \rrbracket[i]$  will be this row. If no such key exists in  $Y$  then an arbitrary rows from  $Y$  are mapped to these locations. Once the mapping is performed the final output table can be constructed by an MPC protocol [MR18] that directly compares the keys for the row of  $\llbracket X \rrbracket[i]$  with  $\llbracket \hat{Y}^0 \rrbracket[i]$  and  $\llbracket \hat{Y}^1 \rrbracket[i]$ . If the keys match then the output row is constructed and otherwise a dummy NULL row is constructed.

Without loss of generality let us assume that the columns  $X_1$  and  $Y_1$  are the join-keys. Our protocol begins by generating a *randomized encoding* for each of the secret shared join-key  $\llbracket x \rrbracket \in \llbracket X_1 \rrbracket$  and  $\llbracket y \rrbracket \in \llbracket Y_1 \rrbracket$ . Figure 2 contains the ideal functionality for this encoding which takes secret shares from the parties, apply a PRF  $F_k$  to the reconstructed value using a internally sampled key  $k$ , and returns the resulting value to one of the three parties. For  $\llbracket x_i \rrbracket := \llbracket X_1 \rrbracket[i]$ ,  $P_0$  will learn  $F_k(x_i)$  while  $P_1$  will learn  $F_k(y_i)$  for  $\llbracket y_i \rrbracket := \llbracket Y_1 \rrbracket[i]$ . Since they join-keys  $x_i$  (resp.  $y_i$ ) are unique and  $k$  is not known, this reveals no information to  $P_0$  (resp.  $P_1$ ).

Party  $P_1$  proceeds by constructing a *secret shared cuckoo hash table*  $\langle\langle T \rangle\rangle$  from the rows of  $\llbracket Y \rrbracket$  where the hash function values for row  $i$  are defined as  $h_j(y_i) = H(j || F_k(y_i))$ . Note that  $P_1$  knows only the randomized encodings  $F_k(y_i)$  of each row  $Y[i]$ , and not the contents of the row itself. The goal in this step is to construct a secret shared cuckoo table  $\langle\langle T \rangle\rangle$  such that row  $Y[i]$  is located at  $T[h_j(y_i)]$  for some  $j$ . We construct  $\langle\langle T \rangle\rangle$  using a three-party *oblivious permutation protocol* where  $P_1$  inputs a permutation  $\pi$ , all parties input secret shares of  $Y$ , and the result is secret shares of “ $Y$  permuted according to  $\pi$ ” which forms  $T$  (details follow later). This is the first transformation shown in Figure 1.

It is now the case that  $\langle\langle T \rangle\rangle$  is a valid cuckoo hash table of  $\llbracket Y \rrbracket$  which is secret shared between  $P_0$  and  $P_1$ . Party  $P_0$ , who knows the randomized encodings  $F_k(x_i)$  for all  $\llbracket x_i \rrbracket := \llbracket X_1 \rrbracket[i]$ , now must compare the rows of  $\langle\langle T \rangle\rangle$  indexed by  $h_j(x_i) = H(j || F_k(x_i))$  with the row  $\llbracket X \rrbracket[i]$ . In particular, assuming we use two cuckoo hash functions, then  $P_0$  constructs two *oblivious switching networks* that maps the shares  $\langle\langle T[h_0(x_i)] \rangle\rangle$  and  $\langle\langle T[h_1(x_i)] \rangle\rangle$  to be “aligned” with  $\llbracket X \rrbracket[i]$ . Exactly how such a network operates is discussed later but the result is two new tables  $\langle\langle Y'_0 \rangle\rangle, \langle\langle Y'_1 \rangle\rangle$  such that  $T[h_j(x_i)] = Y'_j[i]$ . This is the second transformation shown in Figure 1.

Once the shares of  $Y'_0[i] = T[h_0(x_i)]$ ,  $Y'_1[i] = T[h_1(x_i)]$  and obtained using the switching network, the parties employ an MPC protocol to directly compare these rows with  $\llbracket X \rrbracket[i]$ . That is, they compute a bit  $\llbracket b \rrbracket$  which equals one if the join-keys are equal and the **where** clause  $P(\langle\langle Y'_j \rangle\rangle[i], \llbracket X \rrbracket[i])$  outputs one. For each row, a new output row is constructed as  $S(\langle\langle Y'_j \rangle\rangle[i], \llbracket X \rrbracket[i])$  using MPC where  $S$  is the user defined selection circuit. In addition, the MPC circuit outputs the secret shared flag  $\llbracket b \rrbracket$  indicating whether this row should be NULL.

Left joins work in a similar way except that all rows of  $X$  are output and marked not NULL. Finally, unions can be computed by including all of  $Y$  in the output and all of the rows of  $X$  where the comparison bit  $\llbracket b \rrbracket$  is zero. Regardless of the type of join, the protocols do not reveal any information about the tables. In particular, not even the cardinality of the join is revealed due to the use of NULL rows.

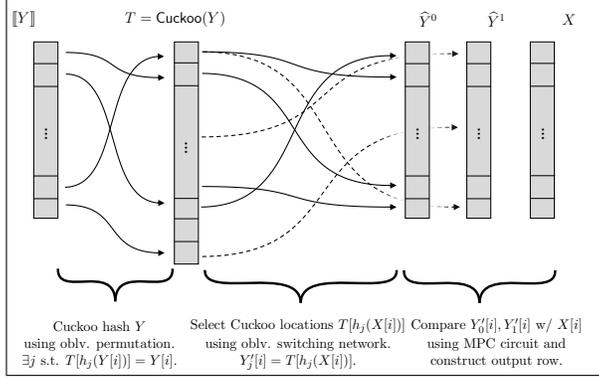


Figure 1: Overview of the join protocol using oblivious switching network.

### 3.2 Randomized Encodings

The randomized encoding functionality  $\mathcal{F}_{\text{ENCODE}}$  of Figure 2 enables the parties to coordinate their secret shares without revealing the underlying values. In particular, the parties will construct a cuckoo hash table using these encodings. The functionality takes as input several tuples  $(\llbracket B_i \rrbracket, \llbracket X_i \rrbracket, P_i)$  where  $B_i \in \{0, 1\}^d$  is an array of  $d$  bits,  $X_i \in (\{0, 1\}^\sigma)^d$  is an array of  $d$  strings and  $P_i$  that denotes that party  $P_i$  should be output the encodings for this tuple. The functionality assigns a random  $\ell$  bit encoding for each input  $x \in \{0, 1\}^\sigma$ . For  $j \in [d]$ , if the bit  $B_i[j] = 0$  then the functionality outputs the encoding for  $X_i[j]$  and otherwise a random  $\ell$  bit string.

Parameters: Input string size of  $\sigma$  bits and output encoding size of  $\ell$  bits.

**[Encode]** Upon receiving command  $(\text{ENCODE}, \{(\llbracket B_i \rrbracket, \llbracket X_i \rrbracket, P_i)\})$  from all parties where  $X_i \in (\{0, 1\}^\sigma)^{d_i}$ ,  $B_i \in \{0, 1\}^{d_i}$  for some  $d_i \in \mathbb{Z}^*$ .

1. Sample a uniformly random  $F : \{0, 1\}^\sigma \rightarrow \{0, 1\}^\ell$ . Define  $F' : \{0, 1\} \times \{0, 1\}^\sigma \rightarrow \{0, 1\}^\ell$  as  $F'(b, x) = \bar{b}F(x) + br$  where  $r \leftarrow \{0, 1\}^\ell$  is sampled each call.
2. For each  $(\llbracket B_i \rrbracket, \llbracket X_i \rrbracket, P_i)$ , send  $\{F'(b, x) \mid (b, x) \in \text{ZIP}(B_i, X_i)\}$  to  $P_i$ .

Figure 2: The Randomized Encoding ideal functionality  $\mathcal{F}_{\text{ENCODE}}$

**LowMC Encodings** We realize this functionality using the LowMC block cipher[ARS<sup>+</sup>15]. When implemented with the honest majority MPC protocols[MR18, AFL<sup>+</sup>16], this approach results in extremely high throughput, computing up to one million encodings per second. Once the parties have their secret shared inputs, they sample a secret shared LowMC key uniformly and encrypt each input under that key using the MPC protocol. These encryptions are revealed as the encodings to the appropriate party.

The LowMC cipher is parameterized by a block size  $\ell$ , keys size  $\kappa$ , s-boxes per layer  $m$  and the desired data complexity  $d$ . To set these parameters, observe that the adversary only sees a bounded number of block cipher outputs (encodings) per key. As such, the data complexity can be bounded by this value. For our implementation we upper bound the number of outputs by  $d = 2^{30}$ . The remaining parameters are set to be  $\ell \in \{80, 100\}$  and  $m = 14$  which results in  $r = 13$  rounds and computational security of  $\kappa = 128$  bits[ARS<sup>+</sup>15]. The circuit for  $\ell = 80$  contains 546 AND gates

(bits of communication).

One issue with the LowMC approach alone is that the input size is fixed to be at most  $\ell \in \{80, 100\}$  bits. However, we will see that the larger join protocol requires an arbitrary input size  $\sigma$ . This is accommodated by applying a universal hash function to the input shares. Specifically, the parties jointly pick a random matrix  $E \leftarrow \{0, 1\}^{\sigma \times \ell}$ . The parties can then locally multiply each secret shared input before it is sent into the LowMC block cipher.

The security of this transformation follows from  $xE \neq x'E$  with overwhelming probability if  $x \neq x'$ . In particular,  $f(x) = xE$  is a universal hash function given that  $E$  is independent of  $x$ . As such the probability that  $f(x) = f(x')$  for any  $x \neq x'$  is  $2^{-\ell}$ . Applying the birthday bound we obtain that probability of any collisions among the tuples is  $2^{-\ell+p}$  where  $p = \log_2 D^2/2 = 2 \log_2(D) - 1$  and  $D = \sum_i d_i$ .

Conditioned on the inputs to the block cipher being unique, the outputs of the block cipher is also distinct and indistinguishable from random  $\ell$  bit strings. As such, in the simulation the real outputs can be replaced with that of the ideal functionality so long as  $2^{-\ell+p}$  is statistically negligible, i.e.  $\ell - p \geq \lambda$ .

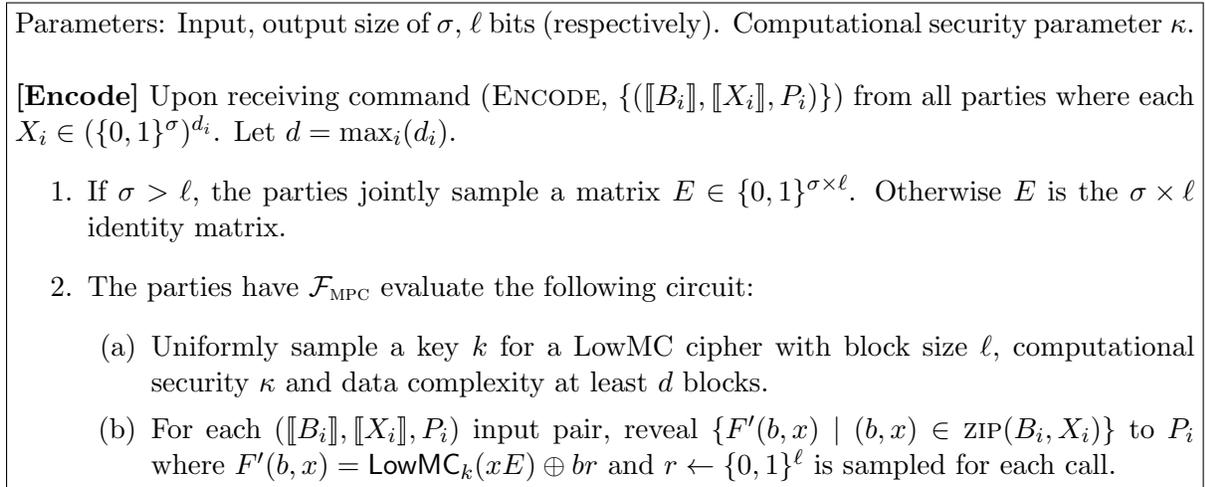


Figure 3: The randomized encoding LowMC protocol.

### 3.3 Oblivious Switching Network

The ideal functionality of a switching network was introduced by Mohassel and Sadeghian[MS13]. It obviously transform a vector  $A = \{A_1, \dots, A_n\}$  such that the output is  $A' = \{A_{\pi(1)}, \dots, A_{\pi(m)}\}$  for an arbitrary function  $\pi : [m] \rightarrow [n]$ . The accompanying protocol of [MS13] was designed in the two party setting where the first party inputs  $A$  while the second party inputs a description of  $\pi$ . This switching network require  $O(n \log n)$  cryptographic operations. Building on this general paradigm, we introduce a *new* oblivious switching network protocol tailored for the honest majority setting with significantly efficiency improves. Our protocol has  $O(n)$  overhead and is constant round. [MS13] requires  $O(n \log n)$  communication/computation and log rounds. Moreover, our protocol can be instantiated with information theoretic security.

The ideal functionality of our protocol is given in Figure 4 with three parties, a programmer  $P_p$ , a sender  $P_s$  and a receiver  $P_r$ .  $P_p$  has a description of  $\pi$  while  $P_s$  has a vector  $A$  containing  $n$  elements each consisting of  $\sigma$  bits.  $P_p$  and  $P_r$  are each output a share of  $\langle\langle A' \rangle\rangle$  s.t.  $A' = \{A_{\pi(1)}, \dots, A_{\pi(m)}\}$ . Later we will discuss the case where  $A$  is secret shared.

Parameters: 3 parties denoted as the  $P_p$ ,  $P_s$  and  $P_r$ . Elements are strings in  $\{0, 1\}^\sigma$ . An input vector size of  $n$  and output size of  $m$ .

[Switch] Upon the command (SWITCH,  $\pi$ ) from the  $P_p$  and (SWITCH,  $A$ ) from the  $P_s$ :

1. Interpret  $\pi : [m] \rightarrow [n]$  and  $A \in (\{0, 1\}^\sigma)^n$ .
2. Compute  $B$  s.t.  $\forall i \in [m], A_{\pi(i)} = B_i$ .
3. Generate  $\langle\langle B \rangle\rangle$  and send  $\langle\langle B \rangle\rangle_0$  to  $P_p$  and  $\langle\langle B \rangle\rangle_1$  to  $P_r$ .

Figure 4: The Oblivious Switching Network ideal functionality  $\mathcal{F}_{\text{SWITCH}}$

**Permutation Network** We begin with a restricted class of switching networks where the programming function  $\pi$  is injective. That is, each input element  $A_i$  will be mapped to a maximum of one location in the output.  $P_p$  samples two random functions  $\pi_0, \pi_1$  such that  $\pi_1 \circ \pi_0 = \pi$ ,  $\pi_0 : [n] \rightarrow [n]$  is bijective and  $\pi_1 : [m] \rightarrow [n]$  is injective.  $P_p$  sends  $\pi_1$  to  $P_r$  and  $\pi_0, S \leftarrow \{0, 1\}^{\sigma \times n}$  to  $P_s$  who sends  $B := \{A_{\pi(1)} \oplus S_0, \dots, A_{\pi(n)} \oplus S_n\}$  to  $P_r$ . The final shares of  $A' = \pi(A)$  are defined as  $P_p$  holding  $\langle\langle A' \rangle\rangle_0 := \{S_{\pi_1(1)}, \dots, S_{\pi_1(m)}\}$  and the  $P_r$  holding  $\langle\langle A' \rangle\rangle_1 := \{B_{\pi_1(1)}, \dots, B_{\pi_1(m)}\}$ .

The simulation of this protocol is perfect. The view of  $P_s$  contains a uniform permutation  $\pi_0$  and vector  $S$ . Similarly, the view of  $P_r$  contains  $\pi_1$  which is uniformly distributed (when  $\pi_0$  is unobserved) and the uniform vector  $B$ . A simulator can sample these directly on behalf of the honest parties.

In the computational secure setting,  $\pi_0, S$  can be generated locally by  $P_0$  and  $P_1$  using a common source of randomness, e.g. a seeded PRG. This reduces the rounds to 1.

**Duplication Network** Next we consider a second type of network where  $\pi : [n] \rightarrow [n]$ ,  $\pi(1) = 1$  and  $\pi(i) \in \{i, \pi(i-1)\}$  for  $i = 2, \dots, n$ . That is, each output position is either a copy of the same input position or is a duplicate of the previous output position, e.g.  $A' = \{A_1, A_1, A_3, A_4, A_4, A_4\}$  where  $A_1, A_4$  were duplicated into the next position(s). This transformation can be characterized by a vector  $b \in \{0, 1\}^n$  where  $b_i = 1$  denotes that the output position  $i$  should be a copy of output position  $i-1$ .

This observation gives rise to a natural protocol: for  $i = 1, \dots, m$ , if  $b_i = 1$  then use MPC to copy  $A_{i-1}$  into  $A_i$ . The primary challenge is to achieve this while using a constant number of communication rounds which prevents the use of a generic (secret sharing) MPC protocol such as [MR18, AFL<sup>+</sup>16].

As a warm-up, suppose that the  $P_s$  inputs  $A_{i-1}, A_i$  and that  $P_p$  &  $P_s$  should receive a share of  $\langle\langle B_i \rangle\rangle$  such that  $B_i := A_{i-b_i}$ .  $P_s$  samples three uniform strings  $\langle\langle B_i \rangle\rangle_1, w_0, w_1 \leftarrow \{0, 1\}^\sigma$  and a uniform bit  $\phi \leftarrow \{0, 1\}$ .  $P_s$  constructs two messages  $m_0 = A_i^1 \oplus \langle\langle B_i \rangle\rangle_1 \oplus w_\phi$  and  $m_1 = A_{i-1} \oplus \langle\langle B_i \rangle\rangle_1 \oplus w_{\phi \oplus 1}$ .  $P_s$  sends  $w_0, w_1$  to the  $P_r$  and sends  $m_0, m_1, \phi$  to  $P_p$  who sends  $\rho = \phi \oplus b_i$  to  $P_r$ . The final shares are constructed by having  $P_r$  send  $w_\rho$  to  $P_p$  who computes  $\langle\langle B_i \rangle\rangle_0 := m_{b_i} \oplus w_\rho$ .

The simulation of this protocol is also perfect.  $P_p$  learns  $m_0 = A_i \oplus \langle\langle B_i \rangle\rangle_1 \oplus w_\phi, m_1 = A_{i-1} \oplus \langle\langle B_i \rangle\rangle_1 \oplus w_{1 \oplus \phi}, w_{b_i \oplus \phi}$ . As such, they can compute  $\langle\langle B_i \rangle\rangle_0 = A_{i-b_i} \oplus \langle\langle B_i \rangle\rangle_1$  but not  $A_{i-b_i} \oplus \langle\langle B_i \rangle\rangle_1$ .  $P_r$  only learns  $w_0, w_1$  and  $\rho = b_i \oplus \phi$  which can be trivially simulated. If computational security is sufficient, then observe that  $w_0, w_1, \phi$  are uniformly random and sending them can be optimized away with a pre-shared PRG seed.

The protocol just described considers the setting where the messages  $A_{i-1}, A_i$  are the private input of  $P_s$ . However, we require that at each iteration the messages being selected is either  $B_{i-1}$  or  $A_i$  where  $\langle\langle B_{i-1} \rangle\rangle$  was computed in the previous iteration. In this case  $P_s$  inputs their share of

$B_{i-1}$  instead of  $A_{i-1}$  while the  $P_p$  computes  $\langle\langle B_i \rangle\rangle_0 := m_{b_i} \oplus w_\rho \oplus b_i \langle\langle B_{i-1} \rangle\rangle_0$ .

Note that this protocol outputs shares to  $P_s$  as opposed to  $P_r$ . This can be corrected by having  $P_s$  secret share  $\langle\langle B \rangle\rangle_1$  between  $P_r$  &  $P_p$  but we prefer to leave the protocol as is for composability reasons.

**Shared Inputs** The protocols and functionality described above assume the vector being transformed is the private input of the  $P_s$ . However, our larger protocols will require the transformations to be applied to secret shared vectors. In particular, the parties hold  $\llbracket A \rrbracket$ . As described in [shared Switch] of 5, the parties first locally convert  $\llbracket A \rrbracket$  into  $\langle\langle A \rangle\rangle$  between  $P_s$  and  $P_p$ . They run  $\mathcal{F}_{\text{SWITCH}}$  where  $P_s$  inputs their share  $B = \langle\langle A \rangle\rangle_0$ .  $P_p$  and  $P_r$  receive  $\langle\langle \pi(B) \rangle\rangle$  from  $\mathcal{F}_{\text{SWITCH}}$  where  $P_p$  holds  $\langle\langle \pi(B) \rangle\rangle_0$ .  $P_p$  locally define  $\langle\langle C \rangle\rangle_0 := \langle\langle \pi(B) \rangle\rangle_0 \oplus \pi(\langle\langle A \rangle\rangle_1)$  and  $P_r$  defines  $\langle\langle C \rangle\rangle_1 := \langle\langle \pi(B) \rangle\rangle_1$ . It is easy to verify that  $C = \pi(A)$ . Simulation of this protocol essentially equivalent to simulating the call to  $\mathcal{F}_{\text{SWITCH}}$  since there is no added communication.

**Universal Switching Network** A universal switching network supporting an *arbitrary*  $\pi : [m] \rightarrow [n]$  can be constructed in three phases[MS13]:  $A \xrightarrow{\pi_1} B \xrightarrow{\pi_2} C \xrightarrow{\pi_3} D = \pi(A)$ .

1.  $B := \pi_1(A)$ : The input vector  $A$  is permuted by the injective function  $\pi_1 : [m] \rightarrow [n]$  such that if  $\pi$  maps an input position  $i$  to  $k$  outputs positions (i.e.  $k = |\text{preimage}(\pi, i)| = |\{j : \pi(j) = i\}|$ ), then there exists a  $j$  such that  $\pi_1(j) = i$  and  $\{\pi_1(j) + 1, \dots, \pi_1(j) + k\} \cap \text{image}(\pi) = \emptyset$ . That is, wherever position  $i$  is mapped by  $\pi_1$ , it should be followed by  $k - 1$  positions that do not appear in the final output.
2.  $C := \pi_2(B)$ : The intermediate vector  $B$  is transformed by a *duplication network*  $\pi_2 : [m] \rightarrow [m]$  which is defined as follows. If position  $A_i$  is mapped to  $k$  positions in  $\pi(A)$ , then  $\{C_j, \dots, C_{j+k}\} = \{A_i\} = \{B_j\}$  where  $\pi_1(j) = i$ . That is, copies  $B_j$  into the next  $k - 1$  positions.
3.  $D := \pi_3(C)$ : The final transformation  $\pi_3 : [m] \rightarrow [m]$  permutes  $C$  to have the same ordering as  $\pi(A)$ . That is, the elements  $\{C_j, \dots, C_{j+k}\}$  which all have the value  $A_i$  are arbitrary mapped to the  $k$  positions  $\{j : \pi(j) = i\}$ .

Observe that steps  $\pi_1, \pi_3$  can both be implemented using the oblivious permutation protocol while  $\pi_2$  can be implemented with a duplication network. Figure 5 provides a formal description of the full switching network protocol.

The simulation of the full protocol ([switch] of Figure 5) essentially follows from the simulation of the permutation and duplication protocols. That is, simulation of step b) and d) of [switch] follows the simulation of Section 3.3 and step c) follows from the simulation of Section 3.3.

### 3.4 Join Protocols

Our join protocol can be divided into four phases:

1. Compute randomized encodings of the join-columns/keys.
2. Party  $P_1$  constructs a cuckoo table  $T$  for table  $Y$  and arranges the secret shares using a permutation protocol.
3. For each row  $x$  in  $X$ ,  $P_0$  uses an oblivious switching network to map the corresponding location  $i_1, i_2$  of the cuckoo hash table to a secret shared tuple  $(x, T[i_1], T[i_2])$ .

Parameters: 3 parties denoted as  $P_p$ ,  $P_s$  and  $P_r$ . Elements are strings in  $\{0, 1\}^\sigma$ . An input, output vector size of  $n, m$ .

**[Permute]** Upon the command (PERMUTE,  $\pi$ ) from  $P_p$  and (PERMUTE,  $A$ ) from  $P_s$ .  $\pi : [m] \rightarrow [n]$  is parsed as a *injective* function and  $A \in \{0, 1\}^{n \times \sigma}$  as a vector of  $n$  elements. Then:

1. If  $n < m$ ,  $P_s$  redefines  $A$  to be  $A := A || \{0\}^{(m-n) \times \sigma}$  and all parties redefine  $n := m$ .
2.  $P_p$  samples a uniformly random bijective function  $\pi_0 : [n] \rightarrow [n]$  and computes the injective function  $\pi_1 : [n] \rightarrow [m]$  such that  $\pi_1 \circ \pi_0 = \pi$ .  $\pi_0$  and a random vector  $S \leftarrow \{0, 1\}^{n \times \sigma}$  are sent to  $P_s$ .
3.  $P_s$  computes and sends  $B := \{A_{\pi_0(1)} \oplus S_1, \dots, A_{\pi_0(n)} \oplus S_n\}$  to  $P_r$ .
4.  $P_p$  sends  $\pi_1$  and a random vector  $T \leftarrow \{0, 1\}^{m \times \sigma}$  to  $P_r$  who outputs  $C^0 := \{B_{\pi_1(1)} \oplus T_1, \dots, B_{\pi_1(m)} \oplus T_m\}$ .  $P_p$  outputs  $C^1 := \{S_{\pi_1(1)} \oplus T_1, \dots, S_{\pi_1(m)} \oplus T_m\}$ .

**[Switch]** Upon the command (SWITCH,  $\pi$ ) from  $P_p$  and (SWITCH,  $A$ ) from  $P_s$ .  $\pi : [m] \rightarrow [n]$  is parsed as a function and  $A \in \{0, 1\}^{n \times \sigma}$  as a vector of  $n$  elements. Then:

1.  $P_p$  samples an injective function  $\pi_1 : [m] \rightarrow [n]$  such that for  $i \in \text{image}(\pi)$  and  $k = |\text{preimage}(\pi, i)|$ , there exists a  $j$  where  $\pi_1(j) = i$  and  $\{\pi_1(j+1), \dots, \pi_1(j+k)\} \cap \text{image}(\pi) = \emptyset$ .  
 $P_p$  sends (PERMUTE,  $\pi_1$ ) to  $\Pi_{\text{SWITCH}}$  and  $P_s$  sends (PERMUTE,  $A$ ).  $P_p$  receives  $B^0 \in \{0, 1\}^{m \times \sigma}$  in response and  $P_r$  receives  $B^1 \in \{0, 1\}^{m \times \sigma}$ .
2.  $P_p$  computes the vector  $b \in \{0, 1\}^m$  such that for  $i \in \text{image}(\pi)$  and  $k = |\text{preimage}(\pi, i)|$ ,  $b_j = 0$  and  $b_{j+1} = \dots = b_{j+k} = 1$  where  $\pi_1(j) = i$ .  
 $P_r$  samples three  $m$  element vectors  $C^1, W^0, W^1 \leftarrow \{0, 1\}^{m \times \sigma}$  and  $\phi \leftarrow \{0, 1\}^m$ . They set  $C_1^1 := B_1^1$  and computes

$$\begin{aligned} M_i^0 &:= B_i^1 \oplus C_i^1 \oplus W_i^{\phi_i} \\ M_i^1 &:= C_{i-1}^1 \oplus C_i^1 \oplus W_i^{\phi_i \oplus 1} \end{aligned}$$

for  $i \in \{2, \dots, m\}$ .  $P_r$  sends  $M, \phi$  to  $P_p$  and  $C^1, W$  to  $P_s$ .  $P_p$  sends  $\rho := \phi \oplus b$  to  $P_s$  who responds with  $\{W_i^{\rho_i} : i \in [m]\}$ .  $P_p$  defines  $C_1^0 := B_1^0$  and computes

$$C_i^0 := M_i^{b_i} \oplus W_i^{\rho_i} \oplus b_i C_{i-1}^0$$

for  $i \in \{2, \dots, m\}$ .

3.  $P_p$  computes the permutation  $\pi_3$  such that for  $i \in \text{image}(\pi)$  and  $k = |\text{preimage}(\pi, i)|$ ,  $\{\pi_3(\ell) : \ell \in \text{preimage}(\pi, i)\} = \{j, \dots, j+k\}$  where  $i = \pi_1(j)$ .  $P_p$  sends (PERMUTE,  $\pi_3$ ) to  $\Pi_{\text{SWITCH}}$  and  $P_s$  sends (PERMUTE,  $C^1$ ).  $P_p$  receives  $S \in \{0, 1\}^{m \times \sigma}$  in response.  $P_r$  receives and outputs  $D^1 \in \{0, 1\}^{m \times \sigma}$ .  
 $P_p$  outputs  $D_i^0 := S_i \oplus C_{\pi_3(i)}^0$  for  $i \in [m]$ .

**[Shared Switch]** Upon the command (SHARED SWITCH,  $\pi, \langle A \rangle_0$ ) from  $P_p$  and (SHARED SWITCH,  $\langle A \rangle_1$ ) from  $P_s$ .  $\pi : [m] \rightarrow [n]$  is parsed as a function and  $A \in \{0, 1\}^{n \times \sigma}$  as a vector of  $n$  elements.

1.  $P_p$  sends (SWITCH,  $\pi$ ) to  $\Pi_{\text{SWITCH}}$  and  $P_r$  sends (SWITCH,  $\langle A \rangle_1$ ).
2. In response,  $P_p$  and  $P_r$  respectively receive  $\langle B \rangle_0$  and  $\langle B \rangle_1$ .  $P_p$  outputs  $\langle B \rangle_0 \oplus \pi(\langle A \rangle_0)$  and  $P_r$  outputs  $\langle B \rangle_1$ .

Figure 5: The Oblivious Switching Network protocol  $\Pi_{\text{SWITCH}}$ .

4. The join-key(s) of  $x$  is compared to that of  $T[i_1], T[i_2]$ . If one of them match then the

corresponding output row is populated; otherwise the output row is set to NULL.

Steps 1 through 3 are performed by the Map routine of [Figure 6](#) while step 4 is performed in [Figure 7](#). [Figure 8](#) contains the ideal functionality of the join protocol.

**Randomized Encodings** We begin by generating randomized encodings of the columns being used for the join-keys. For example,

```
select * from X inner join Y on X1 = Y1 and X2 = Y3
```

In this case there are two join-keys,  $X_1, X_2$  from  $X$  and  $Y_1, Y_3$  from  $Y$ . The protocol has  $P_0$  learn the randomized encoding for each row of  $X$  and  $P_1$  learn them for  $Y$ . Importantly, is that after a previous join operation, some (or all) of the rows being joined can be NULL. We require that the randomized encodings of these rows not reveal that they are NULL. For table  $X$ , a special column  $X_{\text{NULL}}$  encodes if for each row is logically NULL. The  $\mathcal{F}_{\text{ENCODE}}$  functionality will then return a random encoding for all NULL rows. Specifically, the parties will send  $(\text{ENCODE}, \{(\llbracket X_{\text{NULL}} \rrbracket, \llbracket X_{j_1} \rrbracket \dots \llbracket X_{j_l} \rrbracket), P_0\}, (\llbracket Y_{\text{NULL}} \rrbracket, \llbracket Y_{k_1} \rrbracket \dots \llbracket Y_{k_l} \rrbracket), P_1\})$  to  $\mathcal{F}_{\text{ENCODE}}$  where  $j_1, \dots, j_l$  and  $k_1, \dots, k_l$  index the join-keys of  $X$  and  $Y$ . Let  $\mathbb{E}_x, \mathbb{E}_y \in (\{0, 1\}^\ell)^n$  be the encodings that  $P_0$  and  $P_1$  respectively receive from  $\mathcal{F}_{\text{ENCODE}}$ .

For correctness, we require the encoding bit-length  $\ell$  to be sufficiently large such that the probability of a collision between encodings is statistically negligible. Given that there are a total of  $D = 2n$  encodings, the probability of this is at most  $2^{-\ell + 2 \log_2 D - 1}$  which we require to be less than  $2^{-\lambda}$ , therefore  $\ell \geq \lambda + 2 \log_2 D - 1$ . Our implementation uses  $\lambda = 40$  and  $\ell \in \{80, 100\}$  depending on  $D$ .

**Constructing the Cuckoo Table** The next phase of the protocol is for  $P_1$  to construct a secret shared cuckoo table for  $Y$  where each row is inserted based on its encoding in  $\mathbb{E}_y$ .  $P_1$  locally inserts the encodings  $\mathbb{E}_y$  into a plain cuckoo hash table  $t$  with  $m$  slots using the algorithm specified in [Section 2](#). We assume two hash functions are used.  $P_1$  samples an injective function  $\pi : m \rightarrow m$  such that  $t[j] = \mathbb{E}_y[i]$ , then  $\pi(j) = i$ . That is,  $\pi$  defines the mapping from each row's original position in the table  $Y$  to the corresponding position in the cuckoo table  $t$ .

Parties  $P_0$  and  $P_1$  convert  $\llbracket Y \rrbracket$  to  $\langle\langle Y \rangle\rangle$  such that  $P_0$  holds  $\langle\langle Y \rangle\rangle_0$ .  $P_1$  sends  $(\text{SHAREDSWITCH}, \pi, \langle\langle Y \rangle\rangle_1)$  to  $\mathcal{F}_{\text{SWITCH}}$  and  $P_0$  sends  $(\text{SHAREDSWITCH}, \langle\langle Y \rangle\rangle_0)$ . In response  $\mathcal{F}_{\text{SWITCH}}$  sends  $\langle\langle T \rangle\rangle_1$  to  $P_1$  and  $\langle\langle T \rangle\rangle_0$  to  $P_2$ . It is now the case that  $T$  is a valid secret shared cuckoo hash table of  $Y$ . In particular, for a given row  $Y[i]$  with encoding  $e = \mathbb{E}_y[i]$ , there exists a  $j \in \{h_1(e), h_2(e)\}$  such that  $T[j] = Y[i]$ . Here, the  $h_0, h_1$  functions are hash functions used to construct the cuckoo table  $T$ . Another important observation is that  $\pi$  is a permutation and therefore the more efficient permutation protocol can be used in place of the universal switching protocol.

We note that some of the columns of the tables may be secret shared in arithmetic group as opposed to binary shares. In this case the switching network will use the appropriate arithmetic operation as note in [Section 3.3](#).

**Selecting from the Cuckoo Table** The next phase of the protocol is for each row of  $X$ , select the appropriate rows of  $T$  so the keys can be compared.  $P_0$  knows that if the join-keys of the  $X[i]$  row will match with a row from  $Y$ , then this row will be at  $T[j]$  for some  $j \in \{h_1(e), h_2(e)\}$  where  $e = \mathbb{E}_x[i]$ .

To obviously compare these rows,  $P_0$  will construct two switching networks with programming  $\pi_1, \pi_2 : n \rightarrow m$  such that if  $h_l(\mathbb{E}_x[i]) = j$  then  $\pi_l(i) = j$ . Each of these will be used to construct

the tables  $\langle\langle \tilde{Y}^1 \rangle\rangle, \langle\langle \tilde{Y}^2 \rangle\rangle$  which are the result of applying the switching networks  $\pi_1, \pi_2$  to  $\langle\langle T \rangle\rangle$ . In particular, for the  $i$ th row  $X[i]$  it is now the case that if  $X[i]$  has a matching row in  $Y$  then it will be contained at  $\tilde{Y}^1[i]$  or  $\tilde{Y}^2[i]$ .

**Inner Join** Given the three secret shared tables  $\llbracket X \rrbracket, \langle\langle \tilde{Y}^1 \rangle\rangle, \langle\langle \tilde{Y}^2 \rangle\rangle$  as described above, the parties do a linear pass over the  $n$  rows to construct the join between  $X$  and  $Y$ . Recall that the inner join consists of all the selected columns from the rows  $X[i], Y[j]$  where the join-keys of the rows  $X[i]$  and  $Y[j]$  are equal.

If row  $X[i]$  has a matching row in  $Y$  then this row will occupy either  $\tilde{Y}^1[i]$  or  $\tilde{Y}^2[i]$ . To determine which, the parties input the secret shares of these rows to an MPC protocol where the join-keys are compared. For each  $i$  and rows  $\tilde{Y}^1[i], \tilde{Y}^2[i]$  the bits  $c_1[i], c_2[i]$  are generated where  $c_l[i] = 1$  iff the join-keys of  $\tilde{Y}^l[i]$  are equal to that of  $X[i]$ . The MPC circuit then computes  $Y'[i] := c_1[i]\tilde{Y}^1[i] \oplus c_2[i]\tilde{Y}^2[i]$  and  $Y'_{\text{NULL}}[i] := Y'_{\text{NULL}}[i] \vee \neg(c_1[i] \oplus c_2[i])$ . That is,  $Y'[i]$  is a NULL row if it was already NULL or none of the comparisons were equal.

Next, the **where** clause of the query. The **where** clause further filters the output table as a function of  $Y'[i]$  and  $X[i]$ . For example, the query may specify that only rows where  $Y_2'[i] + X_3[i] > 22$  are to be selected. Regardless of the exact where clause, the MPC protocol sets the NULL-bit of the final output table  $Z$  as  $Z_{\text{NULL}}[i] := X_{\text{NULL}}[i] \vee Y'_{\text{NULL}}[i] \vee \neg P(Y'[i], X[i])$  where  $P$  is the predicate function specified by the **where** clause. Finally, the computation specified by the **select** query is performed, e.g. copying the columns of  $X, Y$  or computing a function of them.

**Optimizations** Several optimizations can be applied to this protocol. First, observe that only columns of  $Y$  which explicitly appear in the query need to be input to the switching networks. This reduces the amount of data to be sent and improves performance. Secondly, when comparing the join-columns, instead of computing the equality circuit between all of these columns it suffices to compare the randomized encodings. In the event that the join-column(s) contain many bits, comparing the encodings can reduce the size of the equality circuit. In addition, observe that including columns from  $X$  in the output table is essentially free due to these secret share columns simply being copied from  $X$ . Leveraging this the queries can be optimized by ensuring that the majority of the output columns are taken from  $X$ . Moreover, if a join-column from  $Y$  is in the **select** clause, this output column can be replaced with the matching column in  $X$ .

Also observe that the computation perform heavily lends itself to SIMD instructions. That is, the same computation is repeatedly applied to each row of the output table. Modern MPC protocol such as the ABY<sup>3</sup> framework [MR18, AFL<sup>+</sup>16] are optimized for this setting and can process billions of binary circuit gates per second[AFL<sup>+</sup>16]. In addition the ABY<sup>3</sup> framework can switch between using binary and arithmetic circuits based on which is most efficient for the given computation.

**Left/Right Join** A left join query is similar to an inner join except that all of the rows from the left table  $X$  are included. All rows that are in the inner join are computed as before. For rows only in  $X$ , the bit  $Y'_{\text{NULL}}[i]$  will equal one and is used to initialize the missing columns from  $Y$  to a default, typically NULL. A right join can be implemented symmetrically.

**Union and Set Minus** Our framework is also capable of computing the union of two tables with respect to the join-keys. Specifically, we define the union operator as taking all of the rows from the left table and all of the rows from the right table that would not be present in the inner join. First we compute  $Y \setminus X$  by only including  $X[i]$  if  $Y'[i]$  is NULL, e.i.  $X[i]$  has no matching row in  $Y$ .

The union of  $X$  and  $Y$  is then constructed as  $(Y \setminus X) \parallel X$  where the  $\parallel$  operator denotes the row-wise concatenation of  $X$  to the end of  $Y \setminus X$ .

**Full Join** We construct a full join as  $(X \text{ left join } Y) \text{ union } Y$ . The left join merge the rows in the intersection and the union includes the missing rows of  $Y$ . The overhead of this protocol is effectually twice that of the other protocols.

We note that under some restrictions on the tables being joined, a more efficient protocol for full joins can be achieved. We defer an explanation of this technique to [Section 5.2](#).

**Security** The simulation of these protocols directly follow from the composibility of the subroutines  $\mathcal{F}_{\text{ENCODE}}$ ,  $\mathcal{F}_{\text{SWITCH}}$  and  $\mathcal{F}_{\text{MPC}}$ . First, the output of  $\mathcal{F}_{\text{ENCODE}}$  simply outputs random strings and it is therefore straightforward to simulate.  $\mathcal{F}_{\text{SWITCH}}$  and  $\mathcal{F}_{\text{MPC}}$  both output secret shared values. Finally, correctness is straight forward to analysis and holds so long as there is no encoding collisions and cuckoo hashing succeeds. Parameters are chosen appropriately so these events happen with probability at least  $1 - 2^{-\lambda}$ .

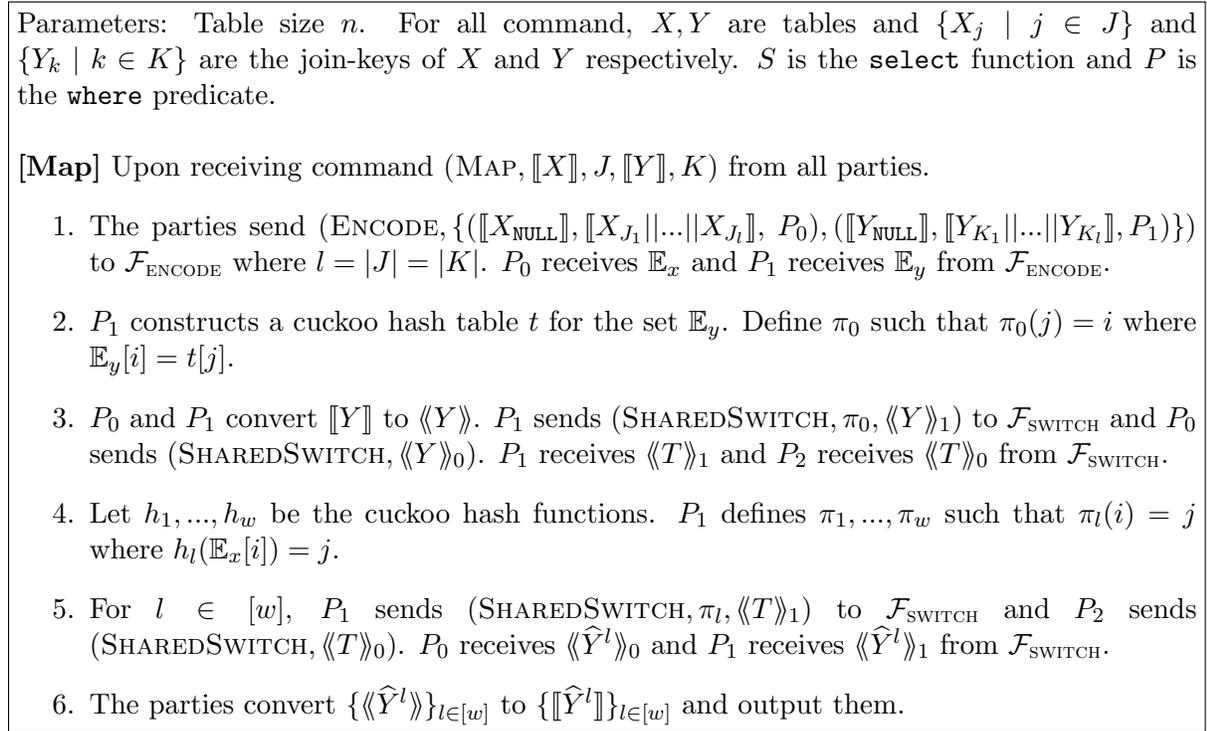


Figure 6: Join protocols  $\Pi_{\text{JOIN}}$ .

### 3.5 Non-unique Join on Column

When values in the join-column are not unique within a single table, the security guarantees begin to erode. Recall that the randomized encodings for  $X, Y$  are revealed to  $P_0, P_1$  respectively. Repeated values in the join-columns will lead to duplicate randomized encodings and therefore reveal their location. Learning the distribution of these duplicates reveals that the underlying table has the same distribution. In the event that only one of the tables contains duplicates, the core protocol can naturally be extended to compute the various join operations subject to  $P_0$  learning

- [Join]** Upon receiving command  $(\text{JOIN}, \text{type}, \llbracket X \rrbracket, J, \llbracket Y \rrbracket, K, S, P)$  from all parties.
1. The parties send  $(\text{MAP}, \llbracket X \rrbracket, J, \llbracket Y \rrbracket, K)$  to  $\Pi_{\text{JOIN}}$  and receive  $\{\llbracket \widehat{Y}^l \rrbracket\}_{l \in [w]}$ .
  2. The parties have  $\mathcal{F}_{\text{MPC}}$  evaluate the following circuit: For  $l \in [w]$  and  $i \in [n]$ ,  $\llbracket c_l \rrbracket[i] := \wedge_j (\llbracket X_{J_j} \rrbracket[i] \stackrel{?}{=} \llbracket \widehat{Y}_{K_j}^l \rrbracket[i])$ . For  $i \in [n]$ ,  $\llbracket Y' \rrbracket[i] := \oplus_l \llbracket c_l \rrbracket[i] \cdot \llbracket \widehat{Y}^l \rrbracket[i]$  and  $\llbracket Y'_{\text{NULL}} \rrbracket[i] := \llbracket Y'_{\text{NULL}} \rrbracket[i] \wedge \neg(\oplus_l \llbracket c_l \rrbracket[i])$ .
  3. If  $\text{type} = \text{INNERJOIN}$ , define the output table  $Z$  by having  $\mathcal{F}_{\text{MPC}}$  evaluate: For  $i \in [n]$ ,  $\llbracket Z_{\text{NULL}} \rrbracket[i] := \llbracket X_{\text{NULL}} \rrbracket[i] \vee \llbracket Y'_{\text{NULL}} \rrbracket[i] \vee \neg P(\llbracket X \rrbracket[i], \llbracket Y' \rrbracket[i])$  and  $\llbracket Z \rrbracket[i] := S(\llbracket X \rrbracket[i], \llbracket Y' \rrbracket[i])$ .
  4. If  $\text{type} = \text{LEFTJOIN}$ , define the output table  $Z$  by having  $\mathcal{F}_{\text{MPC}}$  evaluate: For  $i \in [n]$ ,  $\llbracket Z_{\text{NULL}} \rrbracket[i] := \llbracket X_{\text{NULL}} \rrbracket[i] \vee \neg P(\llbracket X \rrbracket[i], \llbracket Y' \rrbracket[i])$  and  $\llbracket Z \rrbracket[i] := S(\llbracket X \rrbracket[i], \llbracket Y' \rrbracket[i])$ .
  5. If  $\text{type} = \text{UNION}$ , define the output table  $Z$  by having  $\mathcal{F}_{\text{MPC}}$  evaluate: For  $i \in [n]$ ,  $\llbracket Z_{\text{NULL}} \rrbracket[i] := \llbracket X_{\text{NULL}} \rrbracket[i] \vee \neg P(\llbracket X \rrbracket[i], \text{NULL})$  and  $\llbracket Z \rrbracket[i] := S(\llbracket X \rrbracket[i], \text{NULL})$ .  
For  $i \in [n+1, 2n]$ ,  $\llbracket Z_{\text{NULL}} \rrbracket[i] := \llbracket Y_{\text{NULL}} \rrbracket[i] \vee \neg \llbracket X_{\text{NULL}} \rrbracket[i] \vee \neg P(\text{NULL}, \llbracket Y' \rrbracket[i])$  and  $\llbracket Z \rrbracket[i] := S(\text{NULL}, \llbracket Y' \rrbracket[i])$ .
  6. If  $\text{type} = \text{FULLJOIN}$ , all parties sending  $(\text{JOIN}, \text{LEFTJOIN}, \llbracket X \rrbracket, J, \llbracket Y \rrbracket, K, S', P)$  to  $\Pi_{\text{JOIN}}$  and receiving  $\llbracket X' \rrbracket$  in response. They then send  $(\text{JOIN}, \text{UNION}, \llbracket X' \rrbracket, J, \llbracket Y \rrbracket, K, S'', P')$  to  $\Pi_{\text{JOIN}}$  and output the response, where  $S', S''$  and  $P'$  are appropriately updated version of  $S, P$ .

Figure 7: Join protocols  $\Pi_{\text{JOIN}}$  continued.

the duplicate distribution. This is achieved by requiring the left table  $X$  contain the duplicate rows. After learning the randomized encodings for this table  $P_0$  can program the switching networks appropriately to query the duplicate locations in the cuckoo hash table.

When both tables contain duplicates we fall back to a less secure protocol architecture. This is required due to the cuckoo table not supporting duplicates. First,  $P_1$  samples two random permutations  $\pi_0, \pi_1$  and computes  $X' = \pi_1(X), Y' = \pi_2(Y)$  using the oblivious permutation protocol.  $P_0$  then learns all of the randomized encodings for the permuted tables  $X'$  and  $Y'$ . Given this,  $P_0$  can compute the size of the output table and inform the other two parties of it. Alternatively, an upper bound on the output table size can be communicated. Let  $n'$  denote this value.  $P_0$  can then construct two switching networks which map the rows of  $X'$  and  $Y'$  to the appropriate rows of the output table. The main disadvantage of this approach is that  $P_0$  learns the size of the output, the distribution of duplicate rows and how these duplicate rows are multiplied together. However, unlike [LTW13] which takes a conceptually similar approach, our protocol does not leak any information to  $P_1$  and  $P_2$ , besides the value  $n'$ .

### 3.6 Revealing Results

Revealing a secret shared table  $\llbracket X \rrbracket$  requires two operations. First observe that the data in the NULL rows is not cleared out by the join protocols. This is done as an optimization. As such naively reconstructing these rows would lead to significant leakage. Instead  $X[i]$  is updated as  $X[i] = (\neg X_{\text{NULL}}[i]) \cdot X[i]$ . The second operation is to perform an oblivious shuffle of the rows. This operation randomly reorders all the rows without revealing the ordering to any of the parties. This step is necessary since the original ordering of the result table is input-

**[Join]** Upon receiving command  $(\text{JOIN}, \text{type}, \llbracket X \rrbracket, J, \llbracket Y \rrbracket, K, S, P)$  from all parties. Let  $n_X$  and  $n_Y$  denote the number of rows in  $X, Y$  respectively.

1. Define  $\text{KEYS}(X, J, i) = (X_j[i])_{j \in J}$ .
2. If the collections  $\{\text{KEYS}(X, J, i) \mid i \in n \wedge X_{\text{NULL}}[i] = 0\}$  or  $\{\text{KEYS}(Y, K, i) \mid i \in n \wedge Y_{\text{NULL}}[i] = 0\}$  contains duplicates, output  $\perp$ .
3. If  $\text{type} = \text{INNERJOIN}$ , let the rows of  $Z$  be  $\{S(X[i], Y[j]) \mid \exists i, j \text{ s.t. } \neg X_{\text{NULL}}[i] \wedge \neg Y_{\text{NULL}}[j] \wedge \text{KEYS}(X, J, i) \stackrel{?}{=} \text{KEYS}(Y, K, j) \wedge P(X[i], Y[i])\}$  along with zero or more NULL rows s.t.  $Z$  has  $n_X$  rows.
4. If  $\text{type} = \text{LEFTJOIN}$ , let the rows of  $Z$  be  $\{S(X[i], Y[j]) \mid \exists i, j \text{ s.t. } \neg X_{\text{NULL}}[i] \wedge \neg Y_{\text{NULL}}[j] \wedge \text{KEYS}(X, J, i) \stackrel{?}{=} \text{KEYS}(Y, K, j) \wedge P(X[i], Y[i])\} \cup \{S(X[i], \text{NULL}) \mid \exists i, \forall j \text{ s.t. } \neg X_{\text{NULL}}[i] \wedge \text{KEYS}(X, J, i) \neq \text{KEYS}(Y, K, j) \wedge P(X[i], \text{NULL})\}$  along with zero or more NULL rows s.t.  $Z$  has  $n_X$  rows.
5. If  $\text{type} = \text{UNION}$ , let the rows of  $Z$  be  $\{S(X[i], \text{NULL}) \mid \exists i \text{ s.t. } \neg X_{\text{NULL}}[i] \wedge P(X[i], \text{NULL})\} \cup \{S(\text{NULL}, Y[i]) \mid \exists i, \forall j \text{ s.t. } \neg Y_{\text{NULL}}[i] \wedge \text{KEYS}(X, J, j) \neq \text{KEYS}(Y, K, i) \wedge P(\text{NULL}, Y[i])\}$  along with zero or more NULL rows s.t.  $Z$  has  $n_X + n_Y$  rows.
6. If  $\text{type} = \text{FULLJOIN}$ , let the rows of  $Z$  be  $\{S(X[i], Y[j]) \mid \exists i, j \text{ s.t. } \neg X_{\text{NULL}}[i] \wedge \neg Y_{\text{NULL}}[j] \wedge \text{KEYS}(X, J, i) \stackrel{?}{=} \text{KEYS}(Y, K, j) \wedge P(X[i], Y[i])\} \cup \{S(\text{NULL}, Y[i]) \mid \exists i, \forall j \text{ s.t. } \neg Y_{\text{NULL}}[i] \wedge \text{KEYS}(X, J, j) \neq \text{KEYS}(Y, K, i) \wedge P(\text{NULL}, Y[i])\} \cup \{S(X[i], \text{NULL}) \mid \exists i, \forall j \text{ s.t. } \neg X_{\text{NULL}}[i] \wedge \text{KEYS}(X, J, i) \neq \text{KEYS}(Y, K, j) \wedge P(X[i], \text{NULL})\}$
7. Output  $\llbracket Z \rrbracket$  to the parties.

Figure 8: Join functionality  $\mathcal{F}_{\text{JOIN}}$ .

dependent. For example, say  $X$  is a list of patents info,  $Y$  is patent billing status, and  $Z$  is a list of patent diseases. Say we reveal `select X.name, Y.balance from X, Y on X.id = Y.id` and `select X.gender, Z.disease from X, Z on X.id = Z.id`. Without reordering you could connect  $X.name, X.gender, Y.balance$  and  $Z.disease$  by the row index and infer secret information. However, by randomly shuffling this connection is destroyed and the reveal can be simulated.

## 4 Computing a Function of a Table

In addition to join queries, our framework can perform computation on a single secret shared table. For example, selecting  $X_1 + X_2$  where  $X_3 > 42$ . For each row  $i$  we generate the corresponding output row  $Z[i]$  by computing the new NULL-bit as  $Z_{\text{NULL}}[i] := X_{\text{NULL}}[i] \vee P(X[i])$  where  $P(\cdot)$  is the **where** predicate. The new column(s), e.g.  $Z_1 = X_1 + X_2$ , can then be constructed in a straightforward MPC protocol, e.g. [MR18, AFL<sup>+</sup>16]. The key property is that all of the operations are with respect to a single row of  $X$ , allowing them to be evaluated in parallel.

Our framework also considers a second class of functions on a table that allow computation between rows. For example, computing the sum of a column. We refer to this broad class of operations as an aggregation function. Depending on the exact computation, various levels of efficiencies can be achieved. Our primary approach is to employ the ABY<sup>3</sup> framework [MR18] to

express the desired computation in an efficient way and then to evaluate the resulting circuit. Next we highlight a sampling of some important aggregation operations:

- **Sum:** For a column  $\llbracket X_j \rrbracket$ , compute  $\llbracket s \rrbracket = \sum_i \llbracket X_j \rrbracket[i]$  where  $X_j[i] \in \mathbb{Z}_{2^\ell}$  and  $i$  indexes only non-NULL rows. The parties compute  $\llbracket s \rrbracket^A := \sum_{i \in [n]} \text{B2A}(\neg \llbracket X_{\text{NULL}} \rrbracket[i] \cdot \llbracket X_j \rrbracket[i])$  where B2A is the boolean to arithmetic share conversion of [MR18]. In total this requires  $2n\ell$  binary gates and  $\ell + 1$  rounds [MR18]. The parties can then convert  $\llbracket s \rrbracket^A$  back to  $\llbracket s \rrbracket$  if desired.
- **Count/Cardinality:** Here, we consider two cases. 1) In the general case there is an arbitrary table over which the count is being computed. This is performed by computing  $\llbracket s \rrbracket^A := \sum_{i \in [n]} \text{B2A}(\neg \llbracket X_{\text{NULL}} \rrbracket[i])$ . The more efficient bit injection protocol [MR18] can be used to convert each bit to an arithmetic sharing in a constant number of rounds. In particular, the malicious secure bit injection protocol provided in [MR18] is suggested due to it reducing the overall communication.
 

2) Consider case where some of the parties should learn the cardinality of a join without a **where** clause. First, w.l.o.g. let us assume that  $P_2$  should learn the cardinality. The randomized encodings  $\mathbb{E}_x, \mathbb{E}_y$  are respectively revealed  $P_0$  and  $P_1$  as done in the standard join protocol. These encodings are then sent to  $P_2$  in a random order.  $P_2$  outputs  $|\mathbb{E}_x \cap \mathbb{E}_y|$  as the count/cardinality. In the event that  $P_0$  or  $P_1$  should also learn the cardinality,  $P_2$  sends  $|\mathbb{E}_x \cap \mathbb{E}_y|$  to them.
- **Min/Max:** We propose a recursive algorithm where the min/max of the first and second half of the rows is recursively computed. The final result is then the min/max of these two values. Concerning NULL rows, the corresponding value can be initialized to a maximum or minimum sentential value which guarantee that the other value will be propagated. The overall complexity of this approach is  $O(n\ell)$  binary gates and  $O(\ell \log n)$  rounds when using a basic comparison circuit [MR18].

More generally, any polynomial time function can generically be expressed using the ABY<sup>3</sup> framework [MR18]. However, the resulting efficiency may not be adequate for practical deployment.

## 5 Applications

### 5.1 Voter Registration

Improving the privacy and integrity of the United States voter registration system was a primary motivation of the developed protocols. In the United States Electoral College, each state has the responsibility of maintaining their own list of registered citizens. A shortcoming of this distributed process is that without coordination between states it is possible for a voter to register in more than one state. If this person then went on to cast more than one vote the integrity of the system would be compromised. In the case of double registering, it is often a result of a person moving to a new state and failing to unregister from the old state. Alternatively, when a voter moves to a new state it may take them some time to register in the new state, and as such their vote may go uncast. The Pew Charitable Trust [Smi14] reported 1 in 8 voter registration records in the United States contains a serious error while 1 in 4 eligible citizens remain unregistered. The goal in this application of our framework is to improve the accuracy of the voting registration data and help register eligible voters.

A naive solution to this problem is to construct a centralized database of all the registered voters and citizen records. It is then a relatively straightforward process to identify persons with inaccurate

records, attempt to double register or are simply not register at all. However, the construction of such a centralized repository of information has long had strong opposition in the United States due to concerns of data privacy and excessive government overreach. As a compromise many states have volunteered to join the Electronic Registration Information Center (ERIC)[[eri18](#)] which is a non-profit organization with the mission of assisting states to improve the accuracy of America’s voter rolls and increase access to voter registration for all eligible citizens. This organization acts as a semi-trusted third party which maintains a centralized database containing hashes of the relevant information, e.g. names, addresses, drivers license number and social security number.

In particular, instead of storing this sensitive information in plaintext, all records are randomized using two cryptographically strong salted hash functions. Roughly speaking, before this sensitive information is sent to ERIC, each state is provided with the first salt value  $salt_1$  and updates each value  $v$  as  $v := H(salt_1||v)$ . This hashed data is then sent to ERIC where the data is hashed a second time by ERIC which possesses the other salt value. The desired comparisons can then be applied to the hashed data inside ERIC’s secure data center. When compared with existing alternative, this approach provides a moderate degree of protection. In particular, so long as the salt values remain inaccessible by the adversary, deanatomized any given record is likely non-trivial. However, a long series of works, e.g. [[NS06](#), [Mer12](#), [DSS12](#), [OGE16](#), [ZW18](#)], have shown that a significant amount of information can be extracted with sophisticated statistical techniques. Moreover, should the adversary possess the salt values a straightforward dictionary attack can be applied.

We propose adding another layer of security with the deployment of our secure database join framework. In particular, two or more of the states and ERIC will participate in the MPC protocol. From here we consider two possible solutions. The first option is to maintain the existing repository but now have it secret shared between the computational parties. Alternatively, each state could be the long-term holder of their own data and the states perform all pairwise comparison amongst themselves. For reason of preferring the more distributed setting we further explore the pairwise comparison approach.

The situation is further complicated by how this data is distributed within and between states. In the typical setting no single state organization has sufficient information to identify individuals which are incorrectly or double registered. For example, typical voter registration forms requires a name, home address and state ID/driver’s license number. If two states compared this information there would be no reliable attribute for joining the two records. The name of the voter could be used but names are far from a unique identifier. The solution taken by ERIC is to first perform a join between a state’s registered voters and their Department of Motor Vehicles (DMV) records, using the state ID/driver’s license number as the join-key. Since the DMV typically possesses an individual’s Social Security Number (SSN), this can now be used as a unique identifier across all states. However, due to regulations within some states this join is only allowed to be performed on the hashed data or, presumably, on secret shared data.

In addition to identifying individuals that are double registered, the mission of ERIC is to generally improve the accuracy of all voter records. This includes identifying individuals that have moved and not yet registered in their new state or that have simply moved within a state and not updated their current address. In this case the joins between/within states should also include an indicator denoting that an individual has updated their address at a DMV which is different than the voter registration record. There are likely other scenarios which ERIC also identifies but we leave the exploration of them to future work.

Given the building blocks of [Section 3](#) it is a relatively straightforward task to perform the required joins. First a state performs a left join between their DMV data and the voter registration data. Within this join the addresses in the inner join are compared. In the event of a discrepancy, the date of when these addresses were obtained can be compared to identify the most up to date

address. Moreover, the agency with the older address can be notified and initial a procedure to determine which, if any, of the addresses should be updated.

Once this join is performed, each state holds a secret shared table of all their citizens that possess a state ID and their current registration status. Each pair of states can then run an inner join protocol using the social security number as the key. There are several cases that a result record can be in. First it is possible for a person to have a DMV record in two states and be registered in neither. The identity of these persons should not be revealed as this does not effect the voting process. The next case is that a person is registered in both states. We wish to reveal this group to both states so that the appropriate action can be taken. The final case that we are interested in is when a person is registered in state *A* and has a newer DMV address in state *B*. In this case we want to reveal the identity of the person to the state that they are registered to. This state can then contact the person to inquire whether they wish to switch their registration to the new state.

This approach has additional advantages over the hashing technique of ERIC. First, all of the highly sensitive information such as a persons address, state ID number and SSN can still be hashed before being added to the database<sup>3</sup>. However, now that the data is secret shared less sensitive information such as dates need not be hashed. This allows for the more expressive query described above which uses a numerical comparison. To achieve the the same functionality using the current ERIC approach these dates would have to be stored in plaintext which leaks significant information. In addition, when the ERIC approach performs these comparison the truth value for each party of the predicate is revealed. Our approach reveals no information about any intermediate value.

```

stateA = select DMV.name,
               DMV.ID,
               DMV.SSN,
               DVM.date > Voter.date ?
               DMV.date : Voter.date as date,
               DVM.date > Voter.date ?
               DMV.address : Voter.address as address,
               DVM.address ≠ Voter.address as mixedAddress,
               Voter.name ≠ NULL as registered
from DMV left join Voter
  on DMV.ID = Voter.ID
stateB = select ...
resultA = select stateA.SSN
              stateA.address as addressA
              stateB.address as addressB
              stateA.registered
              stateB.registered
from stateA inner join stateB
  on stateA.SSN = stateB.SSN
where (stateA.date < stateB.date and stateA.registered)
      or (stateA.registered          and stateB.registered)
resultB = select ...

```

Figure 9: SQL styled join query for the ERIC voter registration application.

---

<sup>3</sup>The hashing originally performed by ERIC can be replaced with the randomized encoding protocol.

Once the parties construct the tables in [Figure 9](#), state  $A$  can query the table  $stateA$  to reveal all IDs and addresses where the  $mixedAddress$  attribute is set to `true`. This reveals exactly the people who have conflicting addresses between that state’s voter and DMV databases. When comparing voter registration data between states, state  $B$  should define  $stateB$  in a symmetric manner as  $stateA$ . The table  $resultA$  contains all of the records which are revealed to state  $A$  and  $resultB$ , which is symmetrically defined, contains the results for state  $B$ . We note that  $resultA$  and  $resultB$  can be constructed with only one join.

Both types of these queries can easily be performed in our secure framework. All of the conditional logic for the select and where clauses are implemented using a binary circuit immediately after the primary join protocol is performed. This has the effect that overhead of these operation is simply the size of the circuit which implements the logic times the number of potential rows contained in the output.

The average US state has an approximate population of 5 million with about 4 million of that being of voting age. For this set size, our protocol is capable of performing the specified query in 30 seconds and 6GB of total communication. If we consider running the same query where one of the states is California with a voting population of 30 million, our protocol can identify the relevant records in five minutes. For a more extensive performance evaluation, see [Section 6](#).

## 5.2 Threat Log Comparison

Another motivating application is referred to as threat log comparison where multiple organizations share data about current attacks on their computer networks. The goal of sharing this data is to allow the participating parties to identify and stop threats in a more timely manner. Facebook has a service called ThreatExchange[[thr18](#)] which provides this functionality. One drawback of the Facebook approach is that all of the data is collected on their servers and is often viewable by the other participants. This architecture inherently relies on trusting Facebook with this data.

We propose using our distributed protocol to provide a similar functionality while reducing the amount of trust in any single party, e.g. Facebook. In this setting we consider a moderate number of parties each holding a dataset containing the suspicious events on their network along with possible meta data on that event, e.g. how many times that event occurred. All of the parties input these sets into our join framework where the occurrences of each event type are counted. An example of such an event is the IP address that makes a suspicious request.

There are at least two ways to securely compute the occurrences of these events. One method is to perform a full join of all the events where the counts are added together during each join. The resulting table would contain all of the events and the number of times that each event occurred. The drawback of this approach is that each full joins require performing a left join followed by a union, twice the overhead compared to other join operations.

Now consider a different strategy for this problem. First, the parties can compute and reveal the union of the events. Given this information the parties can locally compute the number of times this event occurred on their network and secret share this information between the parties. The parties then add together this vector of secret shared counts and reveal it.

One shortcoming of this approach is no ability to limit which events are revealed. For example, it can be desirable to only reveal an event if it happens on  $k$  out of the  $n$  networks. This can be achieved by having the parties compute and reveal the randomized encodings for all of the items in the union, instead of the items themselves. Under the same encoding key, each party holding a set employs the three server parties to compute the randomized encodings for the items in their set. These encodings are revealed to the party holding the set. For each encoding in the union, the parties use the MPC protocol to compute the number of occurrences that event had and

Operation	Protocol, # Parties	LAN					WAN				
		$2^8$	$2^{12}$	$n$		$2^{24}$	$2^8$	$2^{12}$	$n$		$2^{24}$
Intersection	This ,3	0.02	0.03	0.2	4.9	117.8	2.3	2.5	6.4	41.4	902.0
	[KKRT16] ,2	0.2	0.2	0.4	3.8	58.6	0.6	0.6	1.3	7.5	106.8
	[BA12]* ,3	2.9	23.4	374.4	5,990.4	95,846.4	–	–	–	–	–
Joins/Union	This ,3	0.02	0.03	0.3	9.1	192.1	2.6	2.9	6.6	61.4	1,337.8
	[LTW13]* ,3	2.0	8.0	128.0	2,048.0	32,768.0	–	–	–	–	–
Cardinality	This ,3	0.01	0.02	0.2	3.1	74.5	1.1	1.1	1.8	15.8	267.4
	[PSWW18]a ,2	–	2.2	9.1	86.6	–	–	10.0	45.3	389.9	–
	[PSWW18]b ,2	–	–	–	–	–	–	13.0	56.2	–	–
	[CGT12] ,2	1.0	16.0	262.0	4190.0	67,100.0	–	–	–	–	–
Sum	This ,3	0.03	0.04	0.3	6.8	158.8	3.7	4.0	7.9	51.0	1,099.6
	[IKN+17] ,2	7.0	115.0	1,860.0	29,700.0	475,000.0	–	–	–	–	–

Figure 10: The running time in seconds for various join operations. The input tables each contain  $n$  rows. The [PSWW18] protocol has two implementation where [PSWW18]b is optimized for the WAN setting. – denotes that the running time is not available. \* denotes that the running times were linearly extrapolated from the values of  $n$  provided by the publication.

conditionally reveal the value. For example, if at least  $k$  of the networks observed the event. Other computation on meta data can also be performed as this stage.

## 6 Experiments

We implemented our full set of protocols and applications along with a performance evaluation of them here. Specifically, we considered set intersection (without associated values), our various join and union operations, intersection cardinality, and intersection sum of key-value pairs along with the two application described in Section 5. We also compare to protocols that offer a similar functionality, i.e. [KKRT16, PSWW18, BA12, CGT12, IKN+17]. Our implementation is written in c++ and building on primitives provided by [Rin]. Crucial to the performance of implementation is the widespread use of SIMD instructions that allow processing 128 binary gates with a throughput of one cycle.

**Experimental Setup** We performed all of our experiments on a single server acquired in 2015 which is equipped with two 18 core CPUs at 2.7 GHz and 256 GB of RAM. Despite having many cores, our implementation restricts each party to a *single thread*. We note this is a limitation of development time/resources and not of the protocols themselves. The parties communicate over a loopback device on the local area network which allows to shape the traffic flow to emulate a LAN and WAN setting. Specifically, the LAN setting allows 10 Gbps throughput with a latency of a quarter millisecond while the WAN setting allows an average 100 Mbps and 40 millisecond latency. Despite having such a fast LAN bandwidth, our protocol only utilizes a peak bandwidth of 1Gbps.

All cryptographic operations are performed with computational security parameter  $\kappa = 128$  and statistical security  $\lambda = 40$ . We consider set/table sizes of  $n \in \{2^8, 2^{12}, 2^{16}, 2^{20}, 2^{24}\}$  and  $n = 2^{26}$  in some cases. Times are reported as the average of several trials.

**Set Intersection** We first consider set intersection. In this case the two tables of our protocol consist of a single column which is used as the join-key. We compare our protocol to [KKRT16] which is a two party set intersection protocol where the input sets each are known in the clear to

Operation	Protocol, # Parties	Total Communication (MB)				
		$2^8$	$2^{12}$	$2^{16}$	$2^{20}$	$2^{24}$
Set Intersection	This ,3	0.2	3.0	48.1	769.4	12,318.2
	[KKRT16] ,2	0.04	0.5	8.1	127.2	1,955.2
Joins/Union	This ,3	0.3	4.9	78.1	1,249.4	19,998.1
Cardinality	This ,3	0.1	2.0	32.6	521.5	8,344.0
	[PSWW18]a ,2	–	52.7	826.1	9,971.4	–
	[PSWW18]b ,2	–	14.3	171.3	–	–
	[CGT12] ,2	0.1	0.4	6.2	99.0	1,584.0
Sum	This ,3	0.3	2.0	33.1	526.5	8,372.0
	[IKN <sup>+</sup> 17] ,2	0.1	1.9	30.2	483.0	7,728.0
Voter Intra-state	This ,3	0.2	3.4	54.1	867.1	13,903.6
Voter Inter-state	This ,3	0.4	5.7	91.3	1,463.9	23,482.8
Threat Log $\sim 2$	This ,3	0.2	3.1	50.2	804.2	12,867.9
Threat Log $\sim 4$	This ,3	0.6	9.7	155.4	2,487.8	39,804.4
Threat Log $\sim 8$	This ,3	1.4	22.8	365.7	5,854.9	93,677.5

Figure 11: The total communication overhead in MB for various join operations and applications. The input tables each contain  $n$  rows. The [PSWW18] protocol has two implementation where [PSWW18]b is optimized for the WAN setting. – denotes that the running time is not available.

one of the parties and one party learns the intersection exactly. This is contrasted by our three party protocol where the input and output sets are secret shared between the parties. That is, our protocol is composable while [KKRT16] is not. Both our protocol and [KKRT16] were benchmarked on the same hardware. We also compare to the three party protocol of [BA12] which is composable and was not benchmarked on the same hardware. Due to the code of the [BA12] protocol not being publicly available, we cite their benchmarks which were performed on three AMD Opteron computers at 2.6GHz connected on a 1Gbps LAN network. Given the relative performance of our machines, we believe this to yield a fair comparison. This protocol first sorts the two input sets/tables which in practice requires  $O(n \log^2 n)$  operations/communication [BA12]. In contrast, our protocol and [KKRT16] requires  $O(n)$  operations/communication and  $O(1)$  rounds.

This asymptotic difference also translates to a large difference in the concrete running time as shown in Figure 10. Out of these three protocol [KKRT16] is the fastest requiring 3.8 seconds in the LAN setting to intersect two sets of size  $n = 2^{20}$  while our protocol requires 4.9 seconds. However, our protocol is fully composable while [KKRT16] is not. Considering this we argue that a slowdown of  $1.28\times$  is acceptable. When compared to [BA12] which provides the same composable functionality, our protocol is estimated<sup>4</sup> to be  $1220\times$  faster.

In the WAN setting our protocol has a relative slowdown compared to [KKRT16]. This can be contributed to our protocol requiring more rounds and communication. For instance, with  $n = 2^{20}$  the protocol of [KKRT16] in the WAN setting requires 7.5 seconds while our protocol requires 41 seconds, a difference of  $5.5\times$ . With respect to the communication overhead, our protocol for  $n = 2^{20}$  requires 769 MB of communication and [KKRT16] requires 127 MB, a difference of  $6\times$ . The WAN running time and communication overhead of [BA12] is not known due to their code not being publicly available.

**Joins/Union** The second point of comparison is performing an inner join protocol on two tables consisting of five columns of 32-bit values. We note that [BA12] is capable of this task but no

<sup>4</sup>We linearly extrapolate the overhead of their protocol, despite having  $O(n \log n)$  complexity.

Application	LAN						WAN					
	$2^8$	$2^{12}$	$2^{16}$	$n$ $2^{20}$	$2^{24}$	$2^{26}$	$2^8$	$2^{12}$	$2^{16}$	$n$ $2^{20}$	$2^{24}$	$2^{26}$
Voter Intra-state	0.01	0.02	0.2	4.7	114.7	2,190.1	1.0	1.0	2.2	27.1	456.1	7,463.9
Voter Inter-state	0.01	0.02	0.3	7.0	134.8	2,546.4	1.6	1.6	4.0	45.4	747.7	12,284.1
Threat Log $N = 2$	0.02	0.03	0.2	5.1	121.4	488.1	2.4	2.5	4.8	34.6	585.6	2,342.4
Threat Log $N = 4$	0.05	0.09	0.9	17.9	388.4	1,553.9	6.6	6.8	13.1	108.7	1,739.2	6,956.8
Threat Log $N = 8$	0.10	0.19	1.7	47.1	1,021.0	16,336.1	14.9	15.3	30.0	264.3	4,228.8	16,915.2

Figure 12: The running time in seconds for the Voter Registration and Threat Log applications. The input tables each contain  $n$  rows.

performance results were available. Instead we compare with the join protocol of [LTW13]. This protocol is composable but requires that the cardinality of the intersection be revealed after each join is performed. As previously discussed, this leakage limits the suitability of the protocol in many applications. The numbers reported for [LTW13] are from their paper and the experiments were performed on three servers each with 12 CPUs at 3GHz in the LAN setting. As can be seen in Figure 10, we estimate<sup>5</sup> our protocol is roughly 200× faster in the LAN setting. For example, with  $n = 2^{20}$  our protocol requires a running time of 9.1 seconds while [LTW13] requires a running time of 2048 seconds. Moreover, our protocol scales quite well with the addition of these extra four columns as compared to a intersection protocol. For example, in the WAN setting an intersection with  $n = 2^{20}$  requires 41 seconds while the addition of the four columns results in a running time of 61 seconds. For both protocols, operations such as left join and unions can be performed with little to no additional computation as compared to inner join.

**Cardinality** The set cardinality protocol presented here also outperforms all previous protocols. As described in Section 4, our cardinality protocol allows the omission of the switching network which reduces the amount of communication and overall running time. We demonstrate the performance by comparing with the two-party protocols of Pinkas et al. [PSWW18] and De Cristofaro et al. [CGT12]. The protocol of [PSWW18] was benchmarked on two multi-core i7 machines at 3.7GHz and 16GB of RAM with similar network settings. For the protocol of [CGT12], we performed rough estimates on the time required for our machine to perform the computation without any communication overhead. For sets of size  $n = 2^{20}$  our protocol requires 3.1 seconds in the LAN setting and 15.8 in the WAN setting. The next fastest protocol is [PSWW18] which requires 86.6 seconds in the LAN setting and 390 seconds in the WAN setting. In both cases this represents more than a 20× difference in running time. [PSWW18] considers a variant of their protocol optimized for the WAN setting which reduces their communication at the expense of increased running time. The protocol of [CGT12] requires the most running time by a large margin due to the protocol being based on exponentiation. Just to locally perform these public key operations requires roughly 4200 seconds of computation on our benchmark machine, a difference of 1350×. However, the protocol of [CGT12] also requires the least amount of communication, consisting of 99MB for  $n = 2^{20}$  while our protocol requires 521MB followed by [PSWW18] with almost 10GB.

**Sum** The last generic comparison we perform is for securely computing the weighted sum of the intersection. Our protocol for performing this task is described in Section 4. We compare to the protocol of Ion et al. [IKN+17] which can be viewed as an extension of the public key based cardinality protocol of [CGT12]. In particular, [IKN+17] also revealed the cardinality of the

<sup>5</sup>Again, we linearly extrapolate the overhead of the protocol.

intersection and then performs a secondary computation using Paillier homomorphic encryption to compute the sum. Although this protocol reveals more information than ours, we still think it is a valuable point of comparison. Not surprisingly, the protocol of [IKN<sup>+</sup>17] requires significantly more computation time than our protocol. For a dataset size of  $n = 2^{20}$ , their protocol requires almost 30000 seconds to just perform the public key operations without any communication. Our protocol requires just 6.8 seconds in the LAN setting and 51 in the WAN setting. Both of these protocols also consume roughly the same amount of communication with [IKN<sup>+</sup>17] requiring 483 MB and our protocol requires 527 MB, a increase of just 9 percent.

**Voter Registration** We now turn our attention to the application of auditing the voter registration data between and within the states of the United States as described in Section 5.1 & Appendix ???. In summary, this application checks that a registered voter is not registered in more than one state and cross validates that their current address is correct. Only the identities of the voters which have conflicting data are revealed to the appropriate state to facilitate a process to contact the individual. In addition, the application can be extended to assist the process of enrolling unregistered citizens. This audit processes is performed using two types of join queries. First, each state computes a left join between the DMV database and the list of registered voters. In Figure 11 and 12 we call this join *Voter Intra-state*. For all pairs of states, these tables are then joined to identify any registration error, e.g. double registered. This join is referred to as *Voter Inter-state*. Performance metrics are reported for each of these joins individually and then we estimate the total cost to perform the computation nation wide.

As shown in Figure 12, our protocol can perform the *Voter Intra-state* join with an input set size of 16 million voters ( $n = 2^{24}$ ) in 115 seconds on a LAN network and in 456 seconds on a WAN network. Considering all but three states have a voting population less than  $n = 2^{24}$ , we consider this a realistic estimate on the running time overhead. Our protocol also achieves relatively good communication overhead of 13.9 GB (Figure 11), where each of the servers sends roughly one third of this. On average, that is 830 bytes for each of the  $n$  records. Given these tables, the *Voter Inter-state* join is performed between all pairs of states. For two states with  $n = 2^{24}$ , the benchmark machine required 135 seconds in the LAN setting and 748 seconds in the WAN setting. The added overhead in this second join protocol is an additional **where** clause which requires a moderate sized binary circuit to be securely evaluated. This join requires 23.4 GB of communication, or 1.4KB for each of the  $n$  records.

Given the high value and low frequency of this computation we argue that these computational overheads are very reasonable. Given the current population estimates of each state, we extrapolate that the overall running time to run the protocol between all pairs of 50 states in a LAN setting would be 53,340 seconds (14.8 hours) or 285,687 seconds (about 80 hours) in the WAN setting. However, the running time in the WAN setting could easily be reduced by running protocols in parallel and increasing the bandwidth above the relatively low 100Mbps per party. The total communication overhead is 9,131 GB which is the main bottleneck. While this amount of communication is non-negligible, the actual dollar amount on a cloud such as AWS[aws18] is relatively low (given the importance of the computation), totaling roughly \$820 at a rate of \$0.09 per GB[Has17].

**Threat Log** In this application  $N$  party secret share their data between the three computational parties and delegate the task of identifying the events that appear in at least  $k$  out of the  $N$  data sets. As described in Section 5.2, the protocol proceeds by taking the union of the sets and then the number of times each event occurred is counted and compared against  $k$ . Each event that appears more than  $k$  times is then revealed to all parties. The union protocol can only function

with respect to two input sets. To compute the union of  $N$  sets we use a binary tree structure where pairs of sets are combined. As such, there are a total of  $N - 1$  union operations and a depth of  $\log N$  protocol instances.

When benchmarking we consider  $N = \{2, 4, 8\}$  input sets each of size  $n \in \{2^8, 2^{12}, 2^{16}, 2^{20}, 2^{24}\}$ . Since we do not reveal the size of the union, the final table will be of size  $nN$ . For  $N = 2$  sets each with  $n = 2^{24}$  items our protocol requires 121 seconds in the LAN setting and 586 seconds in the WAN setting. The total communication is 804MB, or approximately 24 bytes per record. If we increase the number of sets to  $N = 8$  we observe that the LAN running time increases to 1,021 seconds and 4,228 seconds in the WAN setting. Given the the total input size increased by  $4\times$ , we observe roughly an  $8\times$  increase in running time. This difference is due to each successive union operation being twice as big. Theoretically the running time and communication of this protocol is  $O(nN \log N)$ .

## References

- [AFL<sup>+</sup>16] Toshinori Araki, Jun Furukawa, Yehuda Lindell, Ariel Nof, and Kazuma Ohara. High-throughput semi-honest secure three-party computation with an honest majority. In Edgar R. Weippl, Stefan Katzenbeisser, Christopher Kruegel, Andrew C. Myers, and Shai Halevi, editors, *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security, Vienna, Austria, October 24-28, 2016*, pages 805–817. ACM, 2016.
- [alt18] Alienvault open threat exchange (otx). 2018.
- [ARS<sup>+</sup>15] Martin R. Albrecht, Christian Rechberger, Thomas Schneider, Tyge Tiessen, and Michael Zohner. Ciphers for MPC and FHE. In Elisabeth Oswald and Marc Fischlin, editors, *Advances in Cryptology - EUROCRYPT 2015 - 34th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Sofia, Bulgaria, April 26-30, 2015, Proceedings, Part I*, volume 9056 of *Lecture Notes in Computer Science*, pages 430–454. Springer, 2015.
- [aws18] Amazon web services. 2018.
- [BA12] Marina Blanton and Everaldo Aguiar. Private and oblivious set and multiset operations. pages 40–41, 2012.
- [CGT12] Emiliano De Cristofaro, Paolo Gasti, and Gene Tsudik. Fast and private computation of cardinality of set intersection and union. In Josef Pieprzyk, Ahmad-Reza Sadeghi, and Mark Manulis, editors, *Cryptology and Network Security, 11th International Conference, CANS 2012, Darmstadt, Germany, December 12-14, 2012. Proceedings*, volume 7712, pages 218–231. Springer, 2012.
- [CHLR18] Hao Chen, Zhicong Huang, Kim Laine, and Peter Rindal. Labeled psi from fully homomorphic encryption with malicious security. In *Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security, CCS 2018, Toronto, Canada, October 14 - 16, 2018*. ACM, 2018.
- [CLR17] Hao Chen, Kim Laine, and Peter Rindal. Fast private set intersection from homomorphic encryption. In Bhavani M. Thuraisingham, David Evans, Tal Malkin, and Dongyan Xu, editors, *Proceedings of the 2017 ACM SIGSAC Conference on Computer*

and *Communications Security, CCS 2017, Dallas, TX, USA, October 30 - November 03, 2017*, pages 1243–1255. ACM, 2017.

- [DRRT18] Daniel Demmler, Peter Rindal, Mike Rosulek, and Ni Trieu. Pir-psi: Scaling private contact discovery. *Proceedings on Privacy Enhancing Technologies*, 2018(4), 2018.
- [DSS12] Anupam Datta, Divya Sharma, and Arunesh Sinha. Provable de-anonymization of large datasets with sparse dimensions. In Pierpaolo Degano and Joshua D. Guttman, editors, *Principles of Security and Trust*, pages 229–248, Berlin, Heidelberg, 2012. Springer Berlin Heidelberg.
- [eri18] Electronic registration information center, inc. 2018.
- [Has17] Hasham. Aws data transfer costs and how to minimize them. 2017.
- [HEK12] Yan Huang, David Evans, and Jonathan Katz. Private set intersection: Are garbled circuits better than custom protocols? In *19th Annual Network and Distributed System Security Symposium, NDSS 2012, San Diego, California, USA, February 5-8, 2012*. The Internet Society, 2012.
- [HFH99] Bernardo A. Huberman, Matt Franklin, and Tad Hogg. Enhancing privacy and trust in electronic communities. In *Proceedings of the 1st ACM Conference on Electronic Commerce, EC '99*, pages 78–86, New York, NY, USA, 1999. ACM.
- [HMRT12] Carmit Hazay, Gert Læssøe Mikkelsen, Tal Rabin, and Tomas Toft. Efficient rsa key generation and threshold paillier in the two-party setting. In Orr Dunkelman, editor, *Topics in Cryptology – CT-RSA 2012*, pages 313–331, Berlin, Heidelberg, 2012. Springer Berlin Heidelberg.
- [HN12] Carmit Hazay and Kobbi Nissim. Efficient set operations in the presence of malicious adversaries. *J. Cryptology*, 25(3):383–433, 2012.
- [IKN<sup>+</sup>17] Mihaela Ion, Ben Kreuter, Erhan Nergiz, Sarvar Patel, Shobhit Saxena, Karn Seth, David Shanahan, and Moti Yung. Private intersection-sum protocol with applications to attributing aggregate ad conversions. *Cryptology ePrint Archive*, Report 2017/738, 2017. <https://eprint.iacr.org/2017/738>.
- [KKRT16] Vladimir Kolesnikov, Ranjit Kumaresan, Mike Rosulek, and Ni Trieu. Efficient batched oblivious PRF with applications to private set intersection. pages 818–829, 2016.
- [KLS<sup>+</sup>17] Ágnes Kiss, Jian Liu, Thomas Schneider, N. Asokan, and Benny Pinkas. Private set intersection for unequal set sizes with mobile applications. *PoPETs*, 2017(4):177–197, 2017.
- [KMP<sup>+</sup>17] Vladimir Kolesnikov, Naor Matania, Benny Pinkas, Mike Rosulek, and Ni Trieu. Practical multi-party private set intersection from symmetric-key techniques. In Bhavani M. Thuraisingham, David Evans, Tal Malkin, and Dongyan Xu, editors, *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security, CCS 2017, Dallas, TX, USA, October 30 - November 03, 2017*, pages 1257–1272. ACM, 2017.

- [KS05] Lea Kissner and Dawn Xiaodong Song. Privacy-preserving set operations. In Victor Shoup, editor, *Advances in Cryptology - CRYPTO 2005: 25th Annual International Cryptology Conference, Santa Barbara, California, USA, August 14-18, 2005, Proceedings*, volume 3621 of *Lecture Notes in Computer Science*, pages 241–257. Springer, 2005.
- [LTW13] Sven Laur, Riivo Talviste, and Jan Willemsen. From oblivious aes to efficient and secure database join in the multiparty setting. In *Proceedings of the 11th International Conference on Applied Cryptography and Network Security, ACNS’13*, pages 84–101, Berlin, Heidelberg, 2013. Springer-Verlag.
- [Mea86] C. Meadows. A more efficient cryptographic matchmaking protocol for use in the absence of a continuously available third party. In *1986 IEEE Symposium on Security and Privacy*, pages 134–134, April 1986.
- [Mer12] Martin M. Merener. Theoretical results on de-anonymization via linkage attacks. *Trans. Data Privacy*, 5(2):377–402, August 2012.
- [MF06] Payman Mohassel and Matthew K. Franklin. Efficient polynomial operations in the shared-coefficients setting. In Moti Yung, Yevgeniy Dodis, Aggelos Kiayias, and Tal Malkin, editors, *Public Key Cryptography - PKC 2006, 9th International Conference on Theory and Practice of Public-Key Cryptography, New York, NY, USA, April 24-26, 2006, Proceedings*, volume 3958 of *Lecture Notes in Computer Science*, pages 44–57. Springer, 2006.
- [MR18] Payman Mohassel and Peter Rindal. ABY3: A mixed protocol framework for machine learning. *IACR Cryptology ePrint Archive*, 2018:403, 2018.
- [MS13] Payman Mohassel and Seyed Saeed Sadeghian. How to hide circuits in MPC an efficient framework for private function evaluation. In Thomas Johansson and Phong Q. Nguyen, editors, *Advances in Cryptology - EUROCRYPT 2013, 32nd Annual International Conference on the Theory and Applications of Cryptographic Techniques, Athens, Greece, May 26-30, 2013. Proceedings*, volume 7881 of *Lecture Notes in Computer Science*, pages 557–574. Springer, 2013.
- [MZ17] Payman Mohassel and Yupeng Zhang. Secureml: A system for scalable privacy-preserving machine learning. In *2017 IEEE Symposium on Security and Privacy, SP 2017, San Jose, CA, USA, May 22-26, 2017*, pages 19–38. IEEE Computer Society, 2017.
- [NP99] Moni Naor and Benny Pinkas. Oblivious transfer and polynomial evaluation. In Jeffrey Scott Vitter, Lawrence L. Larmore, and Frank Thomson Leighton, editors, *Proceedings of the Thirty-First Annual ACM Symposium on Theory of Computing, May 1-4, 1999, Atlanta, Georgia, USA*, pages 245–254. ACM, 1999.
- [NS06] Arvind Narayanan and Vitaly Shmatikov. How to break anonymity of the netflix prize dataset. *CoRR*, abs/cs/0610105, 2006.
- [OGE16] Efe Onaran, Siddharth Garg, and Elza Erkip. Optimal de-anonymization in random graphs with community structure. In *2016 37th IEEE Sarnoff Symposium, Newark, NJ, USA, September 19-21, 2016*, pages 1–2. IEEE, 2016.

- [OOS17] Michele Orrù, Emmanuela Orsini, and Peter Scholl. Actively secure 1-out-of-n OT extension with application to private set intersection. In Helena Handschuh, editor, *Topics in Cryptology – CT-RSA 2017: The Cryptographers’ Track at the RSA Conference 2017, San Francisco, CA, USA, February 14–17, 2017, Proceedings*, pages 381–396, Cham, 2017. Springer International Publishing.
- [PSSZ15] Benny Pinkas, Thomas Schneider, Gil Segev, and Michael Zohner. Phasing: Private set intersection using permutation-based hashing. In Jaeyeon Jung and Thorsten Holz, editors, *24th USENIX Security Symposium, USENIX Security 15, Washington, D.C., USA, August 12–14, 2015.*, pages 515–530. USENIX Association, 2015.
- [PSWW18] Benny Pinkas, Thomas Schneider, Christian Weinert, and Udi Wieder. Efficient circuit-based PSI via cuckoo hashing. In Jesper Buus Nielsen and Vincent Rijmen, editors, *Advances in Cryptology - EUROCRYPT 2018 - 37th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Tel Aviv, Israel, April 29 - May 3, 2018 Proceedings, Part III*, volume 10822 of *Lecture Notes in Computer Science*, pages 125–157. Springer, 2018.
- [PSZ14] Benny Pinkas, Thomas Schneider, and Michael Zohner. Faster private set intersection based on OT extension. In *23rd USENIX Security Symposium (USENIX Security 14)*, pages 797–812, San Diego, CA, 2014. USENIX Association.
- [PSZ16] Benny Pinkas, Thomas Schneider, and Michael Zohner. Scalable private set intersection based on OT extension. Cryptology ePrint Archive, Report 2016/930, 2016. <http://eprint.iacr.org/2016/930>.
- [RA18] Amanda Cristina Davi Resende and Diego F. Aranha. Faster unbalanced private set intersection. 2018.
- [Rin] Peter Rindal. libOTe: an efficient, portable, and easy to use Oblivious Transfer Library. <https://github.com/osu-crypto/lib0Te>.
- [RR17] Peter Rindal and Mike Rosulek. Malicious-secure private set intersection via dual execution. pages 1229–1242, 2017.
- [Smi14] Aaron Smith. 6 new facts about facebook. Pew Research Center Fact Tank, 2014. <http://www.pewresearch.org/fact-tank/2014/02/03/6-new-facts-about-facebook/>.
- [thr18] Facebook threat exchange. 2018.
- [ZW18] Yuchen Zhao and Isabel Wagner. POSTER: evaluating privacy metrics for graph anonymization and de-anonymization. In Jong Kim, Gail-Joon Ahn, Seungjoo Kim, Yongdae Kim, Javier López, and Taesoo Kim, editors, *Proceedings of the 2018 on Asia Conference on Computer and Communications Security, AsiaCCS 2018, Incheon, Republic of Korea, June 04–08, 2018*, pages 817–819. ACM, 2018.