Toward A Practical Multi-party Private Set Union

Jiahui Gao, Son Nguyen, and Ni Trieu

Arizona State University

Abstract. This paper studies a multi-party private set union (mPSU), a fundamental cryptographic problem that allows multiple parties to compute the union of their respective datasets without revealing any additional information. We propose an efficient mPSU protocol which is secure in the presence of any number of colluding semi-honest participants. Our protocol avoids computationally expensive homomorphic operations or generic multi-party computation, thus providing an efficient solution for mPSU.

The crux of our protocol lies in the utilization of new cryptographic tools, namely, Membership Oblivious Transfer (mOT) and Conditional Oblivious Pseudorandom Function (cOPRF). We believe that the mOT and cOPRF may be of independent interest.

We implement our mPSU protocol and evaluate their performance. Our protocol shows an improvement of up to 55× and 776.18× bandwidth cost compared to the existing state-of-the-art protocols.

1 Introduction

Secure multi-party computation (MPC) enables multiple parties to compute an arbitrary function on their private input without revealing additional information. A special case of MPC is the private set operation, which provides a secure means for joining data distributed across disparate databases. Private set intersection (PSI) and private set union (PSU) are two common set operations in this category. PSI finds applications in a variety of privacy-sensitive scenarios such as measuring the effectiveness of online advertising [22], contract tracing [43,3], and contact discovery [20], and cache sharing in IoT [32]. Similarly, PSU has numerous practical use cases. For example, PSU can be used to implement Private-ID functionality [8], cyber risk assessment and management via joint IP blacklists and joint vulnerability data [21], private database supporting full join [27], association rule learning [24], joint graph computation [7], and aggregation of multi-domain network events [10].

Over the last decade, a substantial body of research [37,39,9] has focused on PSI, whereas PSU has received relatively little attention. The majority of present practical PSU protocols [27,17,46,1] have only been optimized for the two-party setting. In this study, we investigate multi-party PSU (mPSU) in the semi-honest model, which allows more than two parties to compute the union of their private data sets without revealing additional information.

1.1 Multi-Party PSU vs 2-Party PSU

Multi-party PSU is a natural extension of the two-party PSU and enables much richer data sharing than a two-party PSU. However, designing a multi-party protocol in secure computation is challenging as it usually requires a dishonest majority (e.g. provides security in the presence of a number of dishonest, colluding participants). Existing mPSU protocols in generic MPC [5,44], or homomorphic encryption[25,16,41,19], are considerably more complex and expensive in the multiparty case than in the two-party case.

A possible solution for computing mPSU is leveraging efficient multi-party PSI protocols. Given their recent PSI improvements [11,31] with practical implementations, one might think that mPSU can be computed directly from multi-party PSI using DeMorgan's Law as $\bigcup_{i=1}^{n} X_i = U \setminus \left(\bigcap_{i=1}^{n}(U \setminus X_i)\right)$, where U is a universe of input items. While this approach correctly and securely computes the set union, it is inefficient when U is significantly larger than $\bigcup_{i=1}^{n} X_i$. Thus, this solution is still far from practical.

Another potential approach is to extend the aforementioned practical two-party PSU protocols [37,39,9] to the multi-party case. However, it remains unclear how to achieve a secure mPSU protocol through this extension since the intermediate result would leak information like the intersection or union of a subset of parties' inputs which violate the mPSU functionality. In general, the state-of-the-art 2-party PSU protocols [17,46,1,4] follow the framework of [27] based on oblivious transfer (OT), which consists of two main stages:

1. Reverse Private Membership Test (RPMT): The receiver learns the bit representing the membership of each element in the sender's set. (e.g. for an element x in sender's set, the receiver with set Y learns a bit $b = 1$ if $x \in Y$ and $b = 0$ otherwise.). Note that the bit b revelas no additional information about the sender's set X, apart from the intersection cardinality $|X \cap Y|$, which is already revealed by the final PSU output.
2. Oblivious Transfer (OT): The senders obliviously sends each item x in its set X to the receiver using OT. Concretely, the sender and the receiver invoke an OT functionality in which the sender possesses messages $\{\perp, x\}$ while the receiver holds the choice bit b, where $\perp$ represents a predefined special character. The bit b is the membership indicator bit which is derived from the preceding stage. The result of the OT provides the receiver with either $\perp$ or the sender's item x which is not the intersection item. By merging this outcome with its set Y, the receiver can produce the set union. This OT step prevents the receiver from deducing the intersection set, thereby fulfilling the functionality of a two-party PSU.

To summarize, in the two-party protocol, the parties initially establish the sender's element membership, followed by the receiver obliviously obtaining only the set difference from the sender. While this framework functions effectively and securely for the 2-party PSU, it cannot be directly extended to multi-party settings due to various sources of information leakage. To be more precise, assume there are n parties, each with a set X_i, and P_1 is the one who receives the final

output. Therefore, the element $x \in \bigcup_{i=2}^{n} X_i \setminus X_1$ must be included in the final output. The initial potential leakage arises from the origin of x as we cannot apply the OT on a pairwise basis between the parties (note that this information does not constitute a leakage in the two-party setting). Another potential leakage is the count of element x. In a 2-party setting, this count is consistently one, as the sender is the sole provider of new elements to the receiver (assuming that the X_2 is not a multi-set). In a multi-party setting, the count of element x can range from 1 to $n-1$. Thus, in the multi-party setting, the definition and execution of RPMT must differ from that in the two-party setting. Furthermore, another main challenge in designing mPSU is to prevent leakages in the event of collusion among a subset of parties

1.2 Related Work

In this section, we focus on the state-of-the-art of multi-party PSU protocols. The earliest construction of such a protocol was proposed by Kissner and Song [25], which relied heavily on homomorphic encryption (the Paillier encryption) and the idea of polynomial representation. Input sets are represented as polynomials where each party $P_{i\in[n]}$ represents an input set $X_i = \{x_{i,1}, \ldots, x_{i,m}\}$ as a polynomial whose roots are its elements, which we denote $f_i(x) = (x - x_{i,1})(x - x_{i,2}) \ldots (x - x_{i,m})$. All parties together compute the encryption of polynomial $p = \prod_{i=1}^{n} f_i$ which presents the polynomial of the union $\bigcup_{i=1}^{n} X_i$. Using polynomial evaluation on the encrypted p, all parties are able to extract the union items without disclosing additional information. The protocol proposed in [25] has $O(n^3 m^2)$ computation complexity. Relying on the polynomial presentation technique, Frikken [16] proposed an efficient mPSU protocol that requires the $O(n^2 m \log(m))$ number of multiplications. [41] presented input sets using rational polynomial functions and reversed Laurent series. As a result, it showed a more efficient protocol than previous works [25,16], but the protocol is secure up to $n/2$ corrupted parties.

Blanton and Aguiar [5] presented a new direction to compute mPSU that avoids expensive homomorphic encryption but heavily relies on MPC. Their idea is to combine the input sets of all parties under a secret-shared form, perform an oblivious sort on the resulting set, and then remove the duplications by comparing the adjacent elements. In the context of MPC, a more practical sorting algorithm is Batcher's network which requires $O(mn \log(mn))$ comparisons to sort the union sets. Due to the underlying MPC techniques, the protocol of [5] is inefficient when the m and n are large.

Recently, [42,19] compute the mPSU using Bloom filter (BF). Specifically, each party $P_{i\in[1,n]}$ inserts its input items into a local BF and transmits the encrypted version of the resulting BF to a designated leader party P_1. Subsequently, the P_1 aggregates the encrypted local BFs from all parties to generate a global BF, denoted by G, from which the union items are computed. While the protocol presented in [42] makes use of an outsourcing server to compute G, [19] is built on homomorphic encryption (HE), which requires a homomorphic

Protocol	Overall Communication	Computation			Round
		Overall	# HE		
			Leader P_1	Party $P_{i\in[2,n]}$	
[25]	$O(n^2m)$	$O(n^3m^2)$	$O(nm^2)$	$O(nm^2)$	$O(n)$
[16]	$O(n^2m)$	$O(n^2m\log(m))$	$O(nm)$	$O(nm\log(m))$	$O(n)$
[5]	$O(n^2m\log^2(mn))$	$O(n^2m\log^2(mn))$	0	0	$O(\log(nm))$
[19]	$O(n^2m\lambda)$	$O(n^2m\lambda)$	$O(nm\lambda)$	$O(nm\lambda)$	$O(1)$
[44]	$O(n^2m\log\lvert\mathcal{U}\rvert)$	$O(n^2m\log\lvert\mathcal{U}\rvert)$	$O(\lvert\mathcal{U}\rvert)$	$O(\lvert\mathcal{U}\rvert)$	$O(1)$
Ours	$O(n^2m\log m/\log\log m)$	$O(n^2m\log m/\log\log m)$	$O(nm)$	$O(nm)$	$O(n)$

Table 1: Communication (overall), computation (overall and number of homomorphic operation), and round complexities of n-party PSU protocols which are secure in *the presence of any number of colluding semi-honest participants*. #HE represents the number of additive homomorphic operations. n is number of parties, each with set size m; λ is the statistical security parameter; $\mathcal{U}$ is the universal domain of the input; σ is the bit length of input element;t is the number of AND gates in the SKE decryption circuit. Notably, the complexity of [19] consists of λ due to the usage of the Bloom filter. All [25,16,19] use Paillier encryption to compute addition on the encryptions (#HE) while our protocol uses Elgamal encryption scheme to re-randomize the ciphertexts.

computation per each entry of G, and might need the expensive multi-key HE. Moreover, the BF-based approach is associated with a high false positive rate.

In another work, Vos *at. el.*[44] proposed private OR protocols and build mPSU protocols upon it. They consider relatively small universe (e.g. up to 32-bit long element). At the high-level idea, their approach presents the input set in a bit vector of length $|\mathcal{U}|$. The bit is set to 1 if its corresponding element belongs to the given input set and 0 otherwise. By invoke the proposed private OR protocol, the leader learns the bit vector of the union. While optimization is given by applying divide-and-conquer so that the long vector can be divided into small ones, it is still inefficient especially for the standard input of 128-bit elements. Concurrently with our work, Liu and Gao [29] presented an efficient mPSU protocol, but requires a weak security assumption wherein *the leader is not in collusion with any other participating parties.*

For a comprehensive analysis of representative multi-party PSU protocols that are resilient to *the presence of any number of colluding semi-honest participants*, we provide a summary of their theoretical complexity in Table 1. Additionally, in Section 5, we present a numerical performance comparison of our proposed protocols with prior works [16,5,19].

1.3 Technical Overview of Our Protocols

We present an efficient protocol for mPSU that guarantees security in the semi-honest setting. We demonstrate the practicality of our mPSU protocol with an implementation. It is shown to be efficient even for large sets with 2^{20} items distributed among 8 parties. The main reason for our protocol's high performance

is its reliance on fast symmetric-key primitives, Diffie-Hellman-based PRF [13] from the elliptic curve, and Elgamal encryption. This is in contrast with prior protocols, which require expensive Paillier encryption on the polynomial set representation [25,16] or each entry of the Bloom filter [19]. Additionally, our approach eliminates the need for the inefficient oblivious sort/OR operations and generic MPC of [5,44].

Technical Overview. In our protocol, we assume the existence of a leader party denoted as P_1. This party collects the union for each of the following sets, specifically $\bigcup_{i=1}^{t} X_i, \forall t \in [2, n]$. This is achieved by interacting sequentially with each party P_t, and the sets are encrypted to prevent the leader from learning the partial union. Moreover, the leader P_1 acquires the encryption of zero for duplicate items from P_t to conceal the size of each union set, as well as the multiplicity of each element in the union. To give the P_1 these encryptions, we use our new building block, Membership Oblivious Transfer (mOT), which is introduced in Section 3.1.

Briefly, the mOT is a two-party protocol, in which a leader P_1 (also referred to as the receiver) holding a set X_1 interacts with the sender P_t who possesses an input item $x_{t,j}, j \in [m]$, and two associated values $\{\mathsf{Enc}(\mathsf{pk}, 0), \mathsf{Enc}(\mathsf{pk}, x_{t,j})\}$). Here, pk represents the encryption's public key, which requires P_t for the decryption process. Similar to the traditional OT [38], the result is that the sender P_t learns nothing whereas the receiver P_1 obtains one of the two sender's associated values depending on whether $x_{t,j} \in X_1$. More specifically, if $x_{t,j} \in X_1$, P_1 learns $\mathsf{Enc}(\mathsf{pk}, 0)$; otherwise it learns $\mathsf{Enc}(\mathsf{pk}, x_{t,j})$. By executing the mOT multiple times with $P_{t \in [2,n]}$ for each item $x_{t,j} \in X_t$, the leader P_1 obtains a set E of encryptions $\mathsf{Enc}(\mathsf{pk}, x_{i,j})$ for $x_{i,j} \in \bigcup_{i=1}^{n} X_i$ and τ encryptions of zero, where $\tau = \sum_{i=1}^{n} |X_i| - |\bigcup_{i=1}^{n} X_i|$ indicates the number of duplicate items. At this point, the union set can be obtained by decrypting E and removing the zeros.

For the dishonest majority setting in which the protocol is secure against an arbitrary number of colluding parties, the decryption should be executed by all the parties. Thus, we employ the multi-key cryptosystem based on Elgamal encryption scheme (ref. Section 2.6). The decryption process involves a partial decryption that requires the individual party's secret key. In our protocol, each party is required to perform its own private permutation on the partial decryption result before sending it to another party. This step aims to prevent a coalition of corrupt parties (including the leader P_1) from learning which parties hold which elements. We implement the permutation and decryption using our simple building block "Oblivious Shuffle and Decryption" ($\mathsf{Shuffle\&Decrypt}$), which is described in Section 3.3.

The brief overview of mOT executions above ignores many important concerns — in particular, how the P_1 obtains encryption of zero from P_i for a duplicated item x which $P_{i \in [2,n]}$ also has. We propose utilizing an Oblivious PRF (OPRF), wherein P_1 obliviously learns and maintains a list of PRF values for the union sets $\bigcup_{i=1}^{t} X_i, \forall t \in [2, n]$. The PRF values hide the actual input items and are used as the input into mOT, rather than the input sets as pre-

viously described. In our protocol, we require an extended variant of OPRF to prevent a coalition of corrupt parties from learning the partial union (we describe the attack explicitly in Section 4). Thus, we introduce a new gadget called Conditional OPRF (cOPRF). Similar to the classical OPRF, the cOPRF has the additional feature that the sender P_i has a set of elements X_i, and the receiver $P_{t>i}$ only obtains the correct PRF value if its input query x_t is not present in the sender's set X_i, and random "fake" value otherwise. We describe the cOPRF ideal functionality and its instantiation in Section 3.2.

In brief, our contributions can be summarized as follows:

- We present an efficient mPSU construction, which eliminates the need for computationally expensive homomorphic operations or generic multi-party computation, and is secure in the presence of any number of colluding semi-honest participants.
- We introduce new building blocks, namely Membership Oblivious Transfer (mOT), Conditional OPRF (cOPRF), Oblivious Shuffle and Decryption (Shuffle&Decrypt), which may be of independent interest and can be used in other related protocols.
- We show that our protocol is significantly faster than previous work [16,5,19]. For example, for four parties with data-set of 2^{16} item each, our mPSU protocol shows an improvement up to 37.82× in terms of running time and up to 776.18× less bandwidth requirement when compared to the state-of-the-art protocols. Our implementation will be made publicly available on GitHub.

2 Preliminaries

In this work, the computational and statistical security parameters are denoted by κ, λ, respectively. We use $[m]$ to refer to the set $\{1, \dots, m\}$, and $[i, j]$ to denote the set $\{i, \dots, j\}$. We denote the concatenation of two strings x and y by $x||y$. We use $f \circ g$ to denote the composition of the functions f and g.

2.1 Multi-party Private Set Union

The ideal functionality of multi-party PSU (mPSU) is given in Figure 1. The mPSU allows n parties, each holding a set X_i of the input items, to learn the union set $\bigcup_{i=1}^{n} X_i$ and nothing else. For simplicity, we assume that all parties have the same set size. Note that our protocol can be easily extended to accommodate varying set sizes, while also revealing the set size of each party. To avoid this leakage, all parties can agree on an upper bound set size m, and use it as the input set size. Before initiating the protocol, each party can pad their set with a particular item, such as zero, to reach the size of m. It is customary in private set operation literature to assume that all parties have the same set size.

PARAMETERS: n parties $P_1, ..., P_n$, and the set size m.
FUNCTIONALITY:

- Wait for input set X_i of size m from P_i.
- Give P_1 the union $\bigcup_{i=1}^{n} X_i$.

Fig. 1: Multi-party Private Set Union Ideal Functionality

2.2 Oblivious PRF

An oblivious pseudorandom function (OPRF) [15] is a 2-party protocol in which the sender learns a PRF key k and the receiver learns $F(k, q_1), \ldots, F(k, q_m)$. Here, F is a PRF, and $(q_1, \ldots, q_m)$ are query inputs chosen by the receiver. Figure 9 formally presents a variant of the OPRF where the receiver obtains outputs of multiple chosen queries.

2.3 Oblivious Transfer

Oblivious Transfer (OT) is a fundamental primitive of secure computation, and introduced by Rabin [38]. It refers to the problem where a sender with two input strings (x_0, x_1) interacts with a receiver who has an input choice bit b. The OT gives the receiver x_b and nothing to the sender. Figure 10 presents the OT functionality.

2.4 Secret-shared Private Membership Test

Secret-shared Private Membership Test (SS-PMT) is the main building block in different applications [35,28,11,36,29,46]. It refers to the two-party setting where a P_0 with input a set of items $X = \{x_1, \ldots, x_n\}$ interacts with a P_1 who has an input single item y. SS-PMT gives both parties a secret-share of a membership bit, i.e. the two parties obtain XOR shares of 1 if $y \in X$ and 0 otherwise. Figure 11 presents the SS-PMT functionality.

2.5 Bin-and-ball Scheme

Our protocols employ hashing schemes such as the Cuckoo and Simple hashing schemes [34,37] to allocate items into bins. We review the basics of the Cuckoo hashing and Simple hashing schemes [34,37] as follows.

Cuckoo hashing. In basic Cuckoo hashing, there are μ bins denoted $B[1 \ldots \mu]$, a stash, and h random hash functions $H_1, \ldots, H_h : \{0,1\}^\star \rightarrow [\mu]$. One can use a variant of Cuckoo hashing such that each item $x \in X$ is placed in exactly one of μ bins. Using the Cuckoo analysis [34,12] based on the set size $|X|$, the parameters μ, h are chosen so that with high probability $(1 - 2^{-\lambda})$ every bin contains at most one item, and no item has to place in the stash during the Cuckoo eviction (.i.e. no stash is required).

Simple hashing. One can map its input set Y into μ bins using the same set of h Cuckoo hash functions (i.e, each item $y \in Y$ appears h times in the hash table). Using a standard ball-and-bin analysis based on h, μ, and $|X|$, one can deduce an upper bound η such that no bin contains more than β items with high probability $(1 - 2^{-\lambda})$.

2.6 Multi-key Cryptosystem

We revise the multi-key cryptosystem that needs for our mPSU protocol. We first give an overview of each component of the cryptosystem. We then present a construction based on the ElGamal scheme. A **multi-key cryptosystem** is defined as a tuple of PPT algorithm $(\mathsf{KeyGen}, \mathsf{Enc}, \mathsf{ParDec}, \mathsf{FulDec}, \mathsf{ReRand})$ with properties as follows:

- **Key Generation**: $\mathsf{KeyGen}(1^\kappa, n)$. In a setting with n parties, a key generation algorithm takes security parameter κ as input and gives each party P_i a secret key sk_i and a joint public key $\mathsf{pk} = \mathsf{Combine}(\mathsf{sk}_1, \mathsf{sk}_2, \ldots, \mathsf{sk}_n)$, where $\mathsf{Combine}$ is an algorithm to generate the public key from the input secret keys depending on the construction.
- **Encryption**: $\mathsf{ct} \leftarrow \mathsf{Enc}(\mathsf{pk}; \mathsf{m})$. Given a joint public key pk and a message $\mathsf{m} \leftarrow \mathcal{M}$ from the plaintext space $\mathcal{M}$, an encryption algorithm outputs a ciphertext ct.
- **Decryption**: There are two types of decryption:
  - Partial decryption $\mathsf{ct}' \leftarrow \mathsf{PartDec}(\mathsf{sk}_i, \mathsf{ct}, \mathsf{A})$. A partially decryption algorithm takes a secret key sk_i and a ciphertext $\mathsf{ct} \leftarrow \mathcal{C}$ encrypted under the partial pubic key $\mathsf{pk}_A = \mathsf{Combine}(\{\mathsf{sk}_j \mid j \in A\})$ and outputs a ciphertext $\mathsf{ct}' \leftarrow \mathcal{C}$ which is encrypted under the partial pubic key $\mathsf{pk}_{A \setminus \{i\}} = \mathsf{Combine}(\{\mathsf{sk}_j \mid j \in A, j \neq i\})$. Note that in the context of the multi-key encryption system, we utilize set A to represent the collection of public keys belonging to the parties within A.
  - Full decryption: $\mathsf{m} \leftarrow \mathsf{FullDec}(\mathsf{sk}_1, \mathsf{sk}_2, ..., \mathsf{sk}_n; \mathsf{ct})$. A full decryption algorithm takes a ciphertext $\mathsf{ct} \leftarrow \mathcal{C}$ encrypted under pk and all the secret keys and outputs a message $\mathsf{m} \leftarrow \mathcal{M}$.
- **Re-randomization**: $\mathsf{ct}' \leftarrow \mathsf{ReRand}(\mathsf{ct}, \mathsf{pk})$. A re-randomization algorithm takes a ciphertext $\mathsf{ct} \leftarrow \mathcal{C}$ encrypted under pk and gives a new ciphertext $\mathsf{ct}' \leftarrow \mathcal{C}$ encrypted under the same pk such that they are both encryptions of the same message $\mathsf{m} \leftarrow \mathcal{M}$.

The multi-key cryptosystem should satisfy correctness and security as defined in [18,2,6]. Informally, the multi-key cryptosystem satisfies correctness if $\mathsf{m} = \mathsf{FullDec}(\mathsf{sk}_1, ..., \mathsf{sk}_n, \mathsf{ct})$ or $\mathsf{m} = \mathsf{FullDec}(\mathsf{sk}_1, \ldots, \mathsf{sk}_{i-1}, \mathsf{sk}_{i+1}, \ldots, \mathsf{sk}_n, \mathsf{ct}')$ for $\mathsf{ct} = \mathsf{Enc}(\mathsf{pk}, m)$ and $\mathsf{ct}' = \mathsf{PartDec}(\mathsf{sk}_i, \mathsf{ct}_{\{1,\ldots,i-1,i+1,\ldots,n\}})$. For security, the ciphertext ct or ct' is random and reveals nothing about the plaintext. When $n = 1$, we have a single-key encryption scheme which is indeed the traditional Elgamal system[14].

A Construction While there are many multi-key cryptosystems, we choose Elgamal system[14] as it is easy to implement and efficient (we do not perform any arithmetic computation on the encryption). In the following, we present the Elgamal scheme in the multi-key setting with n parties $P_1, ..., P_n$.

- **Key Generation**: Given a cyclic group $\mathcal{G}$ of order p, all the parties agree on a common prime g. Each party $P_{i\in[n]}$ chooses a random secret key $\mathsf{sk}_i \leftarrow \{0,1\}^\kappa$ and publishes the value of $h_i = g^{\mathsf{sk}_i}$. We can define the public key $\mathsf{pk} = \mathsf{Combine}(\mathsf{sk}_1, \mathsf{sk}_2, ..., \mathsf{sk}_n) = g^{\Sigma_{i=1}^n \mathsf{sk}_i} = \prod_{i=1}^n h_i$.
- **Encryption**: To encrypt a message m, one can compute $\mathsf{ct} = (\mathsf{ct}_1, \mathsf{ct}_2) = (g^r, m \cdot \mathsf{pk}^r)$ where r is a randomly chosen value from $\{0,1\}^\kappa$.
- **Decryption**: The two decryption algorithms are as follows:
  - Partial decryption: To partially decrypt a ciphertext $\mathsf{ct} = (\mathsf{ct}_1, \mathsf{ct}_2)$ encrypted under the partial pubic key $\mathsf{pk}_A = \prod_{j\in A} h_j$, one can output $\mathsf{ct}' = \mathsf{PartDec}(\mathsf{sk}_i, \mathsf{ct}, A) = (\mathsf{ct}_1', \mathsf{ct}_2')$, where $\mathsf{ct}_1' = \mathsf{ct}_1 \cdot g^{r'}$, $\mathsf{ct}_2' = \mathsf{ct}_2 \cdot \mathsf{ct}_1^{-\mathsf{sk}_i} \cdot (\mathsf{pk}_{A\setminus\{i\}})^{r'}$, the $r' \leftarrow \{0,1\}^\kappa$ is a random value, and $\mathsf{pk}_{A\setminus\{i\}} = \prod_{j\in A\setminus\{i\}} h_j$. Note that the use of the random r' aims to re-randomize the ct_1.
  - Full decryption: To fully decrypt a ciphertext $\mathsf{ct} = (\mathsf{ct}_1, \mathsf{ct}_2)$ encrypted under $\mathsf{pk} = \prod_{i\in[n]} h_i$, one can compute $\mathsf{m} = \mathsf{FullDec}(\mathsf{sk}_1, \mathsf{sk}_2, ..., \mathsf{sk}_n; \mathsf{ct}) = \mathsf{ct}_2 \cdot \mathsf{ct}_1^{-\Sigma_{i=1}^n \mathsf{sk}_i}$.
- **Re-randomization**: To rerandomize a ciphertext ct encrypted under the pk, one can choose a random value $r' \leftarrow \{0,1\}^\kappa$, and compute $(\mathsf{ct}_1', \mathsf{ct}_2') = \mathsf{ReRand}((\mathsf{ct}_1, \mathsf{ct}_2), \mathsf{pk})$ where $\mathsf{ct}_1' = \mathsf{ct}_1 \cdot g^{r'}$ and $\mathsf{ct}_2' = \mathsf{ct}_2 \cdot \mathsf{pk}^{r'}$.

3 Our mPSU Building Blocks

We introduce three simple cryptographic gadgets that will serve as the fundamental building blocks in our mPSU protocol.

- The first gadget is called "Membership Oblivious Transfer" (mOT) which enables a receiver to obtain one of two associated values from the sender based on the set membership. The mOT allows a leader party in our mPSU protocol to obliviously retrieve the items of other parties that are not in the intersection while maintaining privacy.
- The second gadget is called "Conditional OPRF" (cOPRF) which allows the receiver to obtain a PRF value of its query x if and only if x satisfies the predefined condition (i.e., membership test). The cOPRF eliminates duplicated items from the final mPSU output by giving the incorrect (or fake) PRF value of the intersection item.
- In our mPSU protocol, the union result is stored under the multi-key encryption until the final step, which requires all parties to decrypt the ciphertexts together. The encryption protects against corrupted parties from learning partial union. We revise a multi-key cryptosystem in Section 2.6, and introduce a simple tool called "Shuffle and Decryption" (Shuffle&Decrypt) to implement the last step of our mPSU construction.

PARAMETERS: Sender $\mathcal{S}$ and Receiver $\mathcal{R}$, the receiver set size m, the length ℓ.
FUNCTIONALITY:

- Wait for input keyword y and a pair $(v_0, v_1) \subset (\{0,1\}^\ell)^2$ from $\mathcal{S}$.
- Wait for input set $X = \{x_1, ..., x_m\}$ from $\mathcal{R}$.
- Give $\mathcal{R}$ the value v where v equals to v_0 if $y \in X$, and v_1 otherwise.

Fig. 2: Membership Oblivious Transfer (mOT) Ideal Functionality

3.1 Membership Oblivious Transfer (mOT)

Membership Oblivious Transfer (mOT) is a two-party protocol, in which a sender $\mathcal{S}$ with a keyword and two associated values as $(y, \{v_0, v_1\})$ interacts with a receiver $\mathcal{R}$ who has a set of keywords $X = \{x_1, ..., x_m\}$. The result is that the receiver learns the value v_b without learning anything about others $v_{1\oplus b}$ where $b = 0$ if $y \in X$ and $b = 1$ otherwise, while the sender learns nothing about the receiver's input X. Here, the associated values v_0, v_1 are indistinguishable with their domain $\{0,1\}^\ell$. Since the receiver's obtained value depends on whether $y \in X$, we name our gadget "Membership Oblivious Transfer". We formally describe the mOT ideal functionality in Figure 2.

It is easy to define the correctness of the mOT. For the security guarantee, we say that a mOT is secure if it satisfies two following properties:

- Similar to the traditional one-out-of-two oblivious transfer [38], the receiver $\mathcal{R}$ only learns one of the two associated values of the sender $\mathcal{S}$. In addition, the receiver $\mathcal{R}$ has no information about whether $y \in X$ is from the protocol's output. In fact, the latter is satisfied if the associated values (v_0, v_1) come from the same distribution.
- The sender $\mathcal{S}$ learns nothing about the receiver's input and output.

Our mOT Protocol. Our mOT construction consists of two main phases. The first phase follows the popular steps in the circuit-PSI protocols [35,36], which enables the sender and the receiver to compute a secret share of a membership bit, i.e. the two parties obtain XOR shares of 1 or 0 if the sender's keyword y is or is not in the receiver's set X.

The second phase allows the receiver to obtain the corresponding associated value from the sender, depending on whether the output of the first phase was shares of 0 or 1. Typically, this step can be done using generic two-party secure computation (e.g., garbled circuit) in the literature. However, it is relatively inefficient. Instead, we propose a simple solution that relies on OT. More precisely, the sender randomly chooses a value $r \leftarrow \{0,1\}^\ell$ and masks its associated values by computing $(r \oplus v_0, r \oplus v_1)$. Denote a secret share bit of $\mathcal{S}$ and $\mathcal{R}$ to be $b_{\mathcal{S}}$ and $b_{\mathcal{R}}$ received from the first phase, respectively. Using the choice bit $b_{\mathcal{R}}$, the receiver obliviously obtains $w = r \oplus v_{b_{\mathcal{R}}}$ when interacting with the sender with input $(r \oplus v_0, r \oplus v_1)$ via OT. Next, the sender sends $u = r \oplus b_{\mathcal{S}} \cdot (v_1 \oplus v_0)$ to the receiver $\mathcal{R}$. The value u helps to remove the mask r from the w by computing

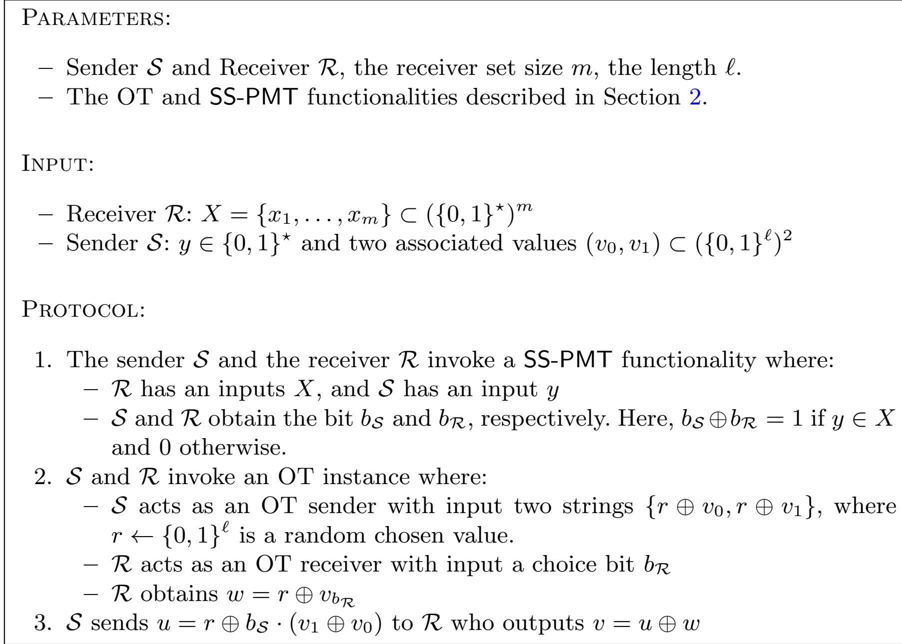
PARAMETERS:

- Sender $\mathcal{S}$ and Receiver $\mathcal{R}$, the receiver set size m, the length ℓ.
- The OT and SS-PMT functionalities described in Section 2.

INPUT:

- Receiver $\mathcal{R}$: $X = \{x_1, \ldots, x_m\} \subset (\{0,1\}^\star)^m$
- Sender $\mathcal{S}$: $y \in \{0,1\}^\star$ and two associated values $(v_0, v_1) \subset (\{0,1\}^\ell)^2$

PROTOCOL:

1. The sender $\mathcal{S}$ and the receiver $\mathcal{R}$ invoke a SS-PMT functionality where:
   - $\mathcal{R}$ has an inputs X, and $\mathcal{S}$ has an input y
   - $\mathcal{S}$ and $\mathcal{R}$ obtain the bit $b_\mathcal{S}$ and $b_\mathcal{R}$, respectively. Here, $b_\mathcal{S} \oplus b_\mathcal{R} = 1$ if $y \in X$ and 0 otherwise.
2. $\mathcal{S}$ and $\mathcal{R}$ invoke an OT instance where:
   - $\mathcal{S}$ acts as an OT sender with input two strings $\{r \oplus v_0, r \oplus v_1\}$, where $r \leftarrow \{0,1\}^\ell$ is a random chosen value.
   - $\mathcal{R}$ acts as an OT receiver with input a choice bit $b_\mathcal{R}$
   - $\mathcal{R}$ obtains $w = r \oplus v_{b_\mathcal{R}}$
3. $\mathcal{S}$ sends $u = r \oplus b_\mathcal{S} \cdot (v_1 \oplus v_0)$ to $\mathcal{R}$ who outputs $v = u \oplus w$

Fig. 3: Membership Oblivious Transfer (mOT) Construction

$v = u \oplus w$, which is the receiver's output. We formally present the construction of our mOT in Figure 3.

For the correctness of the mOT construction, one can rewrite $w = r \oplus b_\mathcal{R} \cdot v_1 \oplus (1 \oplus b_\mathcal{R}) \cdot v_0$. Hence, $v = u \oplus w = (b_\mathcal{R} \oplus b_\mathcal{S}) \cdot v_1 \oplus (1 \oplus b_\mathcal{R} \oplus b_\mathcal{S}) \cdot v_0$ which equals to $v_{b_\mathcal{R} \oplus b_\mathcal{S}}$ as desired (recall that $b_\mathcal{R} \oplus b_\mathcal{S} = 1$ if $y \in X$ and 0 otherwise). We present the security proof of the below theorem in Appendix A.1.

Theorem 1. *The mOT protocol described in Figure 3 securely implements the mOT functionality defined in Figure 2 in the semi-honest setting, given the OT and SS-PMT functionalities described in Section 2.*

3.2 Conditional OPRF (cOPRF)

As described in Section 2.2, an OPRF [15] enables the receiver to learn a PRF value $F(k, q)$ on its input query q without knowing the sender's PRF key k. In this work, we introduce a new notion of a conditional oblivious PRF (cOPRF). Intuitively, the functionality is similar to OPRF, with the additional feature that the sender has a set of elements X, and the receiver only obtains the correct PRF value if its query q is not in the sender's set X. The formal description of a conditional oblivious PRF (cOPRF) functionality is given in Figure 4.

For the correctness, we say that a conditional OPRF satisfies correctness: Given the PRF k key generated from $\mathsf{KeyGen}(1^\kappa)$, (i) if $q \notin X$, then the cOPRF

PARAMETERS: Sender $\mathcal{S}$ and Receiver $\mathcal{R}$, the receiver set size m, and the PRF F.
FUNCTIONALITY:

- Wait for input set $X = \{x_1, ..., x_m\}$ and a PRF key k from $\mathcal{S}$.
- Wait for input a query q from $\mathcal{R}$.
- Give $\mathcal{R}$ the PRF value $F(k, q)$ if $q \notin X$, and random otherwise.

Fig. 4: Conditional OPRF (cOPRF) Ideal Functionality

outputs a valid PRF value $F(k, q)$; (ii) otherwise, the cOPRF outputs a random value. For the security guarantee, we consider the following experiment/game where F_{cOPRF} is the output function of the cOPRF:

$$\begin{aligned} &\underline{\mathsf{Exp}^{\mathcal{A}}(X, q, \kappa):} \\ &\quad k \leftarrow \mathsf{KeyGen}(1^\kappa) \\ &\quad \text{return } \mathcal{A}\big(F_{\mathsf{cOPRF}}(k, X, q)\big) \end{aligned}$$

We say that a cOPRF is m**-secure** if for all $|X_1| = |X_2| = m$ and all polynomial-time $\mathcal{A}$:

$$\Big| \Pr[\mathsf{Exp}^{\mathcal{A}}(X_1, q, \kappa) \Rightarrow 1] - \Pr[\mathsf{Exp}^{\mathcal{A}}(X_2, q, \kappa) \Rightarrow 1] \Big|$$

is negligible in κ

The security definition indicates that the sender's set size m is a public parameter, which allows us to state that the probability of two experiments in the definition is indistinguishable. Thus, given the output of the cOPRF, there is no information about what the sender's set of items was.

Our cOPRF Protocol. We present the construction of an cOPRF, which is built on our mOT primitive and OPRF. While many OPRF protocols exist such as the BaRK OPRF [26], we use the Diffie-Hellman OPRF protocol [13] in which the PRF value of q has a form $F(k, q) = H(q)^k$ for a hash function H.

The protocol starts with the receiver $\mathcal{R}$ picking a random number $\alpha \leftarrow \{0, 1\}^\kappa$ and computing $v_1 = H(q)^\alpha$. The goal is to allow the receiver $\mathcal{R}$ obliviously send v_1 to the sender $\mathcal{S}$ if the query q is not in the sender's set X. This can be done using the mOT in which the receiver $\mathcal{R}$ acts as the mOT's sender with input (q, v_0, v_1) for a random $v_0 \leftarrow \{0, 1\}^\kappa$ while the sender $\mathcal{S}$ acts as the mOT's receiver with input X. As a result, $\mathcal{R}$ obtains v. Next, the $\mathcal{S}$ raises v to the k power as $w = v^k$ and sends the result w back to the $\mathcal{R}$. Now, the receiver can raise the w to the $1/\alpha$ to obtain the final output y. We formally present our cOPRF construction in Figure 5.

It is not hard to see that the output of the cOPRF protocol satisfies correctness. More precisely, if $q \notin X$, then $v = v_1 = H(q)^\alpha$, thus, the protocol's output $y = w^{1/\alpha} = H(q)^k$ as desired. In case the $q \in X$, the value $y = v_0^{k/\alpha}$ is a random value (i.e. an incorrect PRF value of q). We present the security proof of the below theorem in Appendix A.2.

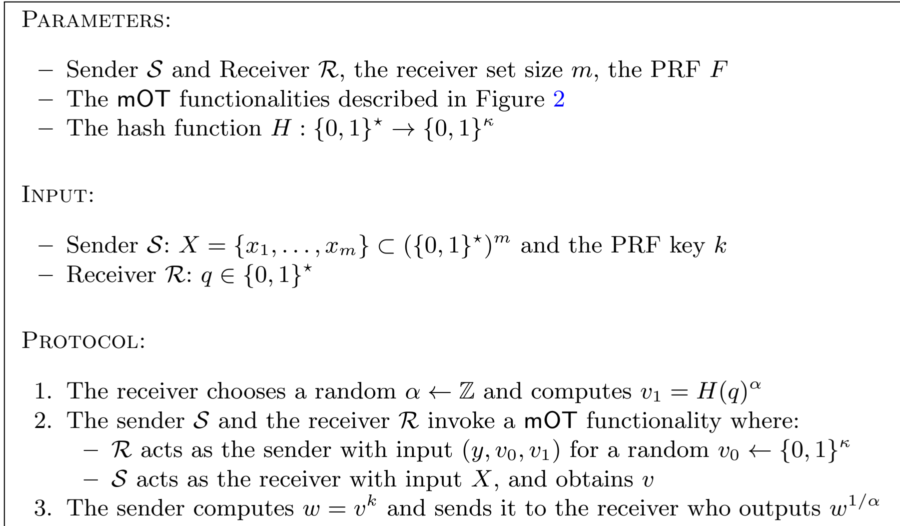

Parameters:

- Sender $\mathcal{S}$ and Receiver $\mathcal{R}$, the receiver set size m, the PRF F
- The mOT functionalities described in Figure 2
- The hash function $H : \{0,1\}^\star \rightarrow \{0,1\}^\kappa$

Input:

- Sender $\mathcal{S}$: $X = \{x_1, \ldots, x_m\} \subset (\{0,1\}^\star)^m$ and the PRF key k
- Receiver $\mathcal{R}$: $q \in \{0,1\}^\star$

Protocol:

1. The receiver chooses a random $\alpha \leftarrow \mathbb{Z}$ and computes $v_1 = H(q)^\alpha$
2. The sender $\mathcal{S}$ and the receiver $\mathcal{R}$ invoke a mOT functionality where:
   - $\mathcal{R}$ acts as the sender with input (y, v_0, v_1) for a random $v_0 \leftarrow \{0,1\}^\kappa$
   - $\mathcal{S}$ acts as the receiver with input X, and obtains v
3. The sender computes $w = v^k$ and sends it to the receiver who outputs $w^{1/\alpha}$

Fig. 5: Conditional OPRF (cOPRF) Construction

Theorem 2. *The cOPRF protocol described in Figure 5 securely implements the cOPRF functionality defined in Figure 4 in the semi-honest setting, given the mOT functionalities described in Section 2.*

3.3 Oblivious Shuffle and Decryption (Shuffle&Decrypt)

Oblivious Shuffle and Decryption (Shuffle&Decrypt) is a n-party protocol, in which each party $P_{i\in[n]}$ holds a permutation $\pi_i : [m] \rightarrow [m]$ and a secret key sk_i of the multi-key cryptosystem as $(\mathsf{pk}, \{\mathsf{sk}_i\}_{i\in[n]}) \leftarrow \mathsf{KeyGen}(1^\kappa, n)$. Given a set of ciphertexts $\{\mathsf{ct}_1, \ldots, \mathsf{ct}_m\}$ where $\mathsf{ct}_i = \mathsf{Enc}(\mathsf{pk}, x_i)$, the Shuffle&Decrypt functionality gives $\{x_{\pi(1)}, \ldots, x_{\pi(m)}\}$ to the party P_1 where $\pi = \pi_n \circ \pi_{n-1} \circ \ldots \circ \pi_1$, and nothing to other parties. The private permutation aims to remove the linkage between the ciphertext ct_i and the plaintext x_i. We formally describe the Shuffle&Decrypt ideal functionality in Figure 6.

Our Shuffle&Decrypt Protocol. The Shuffle&Decrypt construction is simple and directly built from calling algorithms provided in the multi-key cryptosystem. First, the P_1 re-randomizes the ciphertexts and then permutes the result. P_1 then sends the permuted set C_1 to P_2. The re-randomization aims to hide the permutation function from P_2. The P_2 now performs partial decryption using its secret key sk_2. This decryption removes the role of sk_2 from the original ciphertext. P_2 then applies the permutation π_2 on the resulting ciphertexts C_2 and forwards them to P_3. Note that P_2 does not need to rerandomize C_2 as the C_2 is in the random distribution and thus it hides the permutation of P_2. The process repeats sequentially through $P_4, \ldots, P_n$. After the partial decryption

PARAMETERS: n parties, parameter m, and a multi-key encryption scheme defined in Section 2.6
FUNCTIONALITY:

- Wait for input secret key sk_i and a permutation function $\pi_i : [m] \to [m]$ from each party $P_{i \in [n]}$. Here, $(\mathsf{pk}, \{\mathsf{sk}_i\}_{i \in [n]}) \leftarrow \mathsf{KeyGen}(1^\kappa, n)$
- Wait for a combined input a set of ciphertexts $\{\mathsf{ct}_1, \ldots, \mathsf{ct}_m\}$ where $\mathsf{ct}_i = \mathsf{Enc}(\mathsf{pk}, x_i)$ from all parties $\{P_1, \ldots, P_n\}$.
- Give $\{x_{\pi(1)}, \ldots, x_{\pi(m)}\}$ to P_1 where $\pi = \pi_n \circ \pi_{n-1} \circ \ldots \circ \pi_1$.

Fig. 6: Oblivious Shuffle and Decryption (Shuffle&Decrypt) Ideal Functionality

PARAMETERS:

- n parties, the set size m
- A multi-key cryptosystem $(\mathsf{KeyGen}, \mathsf{Enc}, \mathsf{ParDec}, \mathsf{FulDec}, \mathsf{ReRand})$ defined in Section 2.6

INPUT:

- Each party $P_{i \in [n]}$: The secret key sk_i and a permutation function $\pi_i : [m] \to [m]$. Here, $(\mathsf{pk}, \{\mathsf{sk}_i\}_{i \in [n]}) \leftarrow \mathsf{KeyGen}(1^\kappa, n)$
- All parties: $\{\mathsf{ct}_1, \ldots, \mathsf{ct}_m\}$ where $\mathsf{ct}_i = \mathsf{Enc}(\mathsf{pk}, x_i)$

PROTOCOL:

1. P_1 re-randomizes $\mathsf{ct}_j = \mathsf{ReRand}(\mathsf{ct}_j, \mathsf{pk}), \forall j \in [m]$, and sends $C_1 = \{\mathsf{ct}^1_1, \ldots, \mathsf{ct}^1_m\}$ to P_2 where $\mathsf{ct}^1_j = \mathsf{ct}_{\pi_1(j)}, \forall j \in [m]$.
2. For $i = 2$ to n:
   - P_i computes a partial decryption $\bar{\mathsf{ct}}^i_j = \mathsf{PartDec}(\mathsf{sk}_i, \mathsf{ct}^{i-1}_j, A_i), \forall j \in [m]$, where $A_i = \{1, i, i+1, \ldots, n\}$
   - P_i permutes the set $\{\bar{\mathsf{ct}}^i_1, \ldots, \bar{\mathsf{ct}}^i_m\}$ as $\mathsf{ct}^i_j = \bar{\mathsf{ct}}_{\pi_i(j)}, \forall j \in [m]$
   - P_i sends $C_i = \{\mathsf{ct}^i_1, \ldots, \mathsf{ct}^i_m\}$ to $P_{(i+1)\%n}$
3. P_1 outputs $\mathsf{PartDec}(\mathsf{sk}_1, \mathsf{ct}^n_j, \{1\}), \forall j \in [m]$.

Fig. 7: Oblivious Shuffle and Decryption (Shuffle&Decrypt) Construction

was executed by P_n, the ciphertexts require only the secret key sk_1 for the final decryption. P_n now sends these ciphertexts in the permuted order to P_1 which performs the partial decryption and outputs the final result. Figure 7 presents the Shuffle&Decrypt construction. From the high-level description, it is clear that the protocol is correct given the correctness of the underlying multi-key cryptosystem. We present the security proof of the below theorem in Appendix A.3.

Theorem 3. *Given the multi-key cryptosystem defined in Section 2.6, the Shuffle&Decrypt protocol described in Figure 7 securely implements the Shuffle&Decrypt functionality defined in Figure 6 in the semi-honest model, against any number of corrupt, colluding, semi-honest parties.*

	$P_1 : X_1$	$P_2 : \{x_2\}$	$P_3 : \{x_3\}$	$P_4 : \{x_4\}$
Round 1	$\Leftarrow$ OPRF & mOT $\Rightarrow$	$\Leftarrow$ cOPRF	$\Rightarrow$	
			r_3	
S	$F_1(X_1) \cup \{F_1(x_2)\}$	$F_1(x_2)$ $\Leftarrow$	cOPRF	$\Rightarrow$
E	$\{\mathsf{Enc}(x_2)\}$			$F_1(x_4)$
Round 2	$\Leftarrow$	OPRF & mOT	$\Rightarrow$ $\Leftarrow$ cOPRF	$\Rightarrow$
S	$F_2(X_1) \cup \{F_2(x_2)\} \cup \{F_2(r_3)\}$		$F_2(r_3)$	$F_2(x_4)$
E	$\{\mathsf{Enc}(x_2), \mathsf{Enc}(0)\}$			
Round 3	$\Leftarrow$	OPRF & mOT		$\Rightarrow$
S	$F_3(X_1) \cup \{F_3(x_2)\} \cup \{F_3(r_3)\} \cup \{F_(x_4)\}$			$F_3(x_4)$
E	$\{\mathsf{Enc}(x_2), \mathsf{Enc}(0), \mathsf{Enc}(x_4)\}$			
Shuffle	$\{\mathsf{Enc}(0), \mathsf{Enc}(x_4), \mathsf{Enc}(x_2)\}$			
Decrypt	$\{0, x_4, x_2\}$			
Output	$X_1 \cup \{x_4, x_2\}$			

Table 2: Illustration of our mPSU protocol for 4 parties. P_1 has an input set of X_1 while P_i have input set of only one item x_i for $i \in [2, 4]$. In addition, we assume that $x_2 = x_3 \notin X_1$ and $x_4 \in X_1$. The $F_{i-1}(\cdot)$ denotes the multi-key PRF $F((k_{i+1}, \ldots, k_2), \cdot)$, which represents the PRF value received in the i-th round. $\Leftarrow \mathsf{P} \Rightarrow$ denotes the execution of protocol P between two parties and colors are given to show different invocations. We only show the output of each invocation. For example, in round 1, P_1 and P_2 invoke the OPRF and mOT (Step (3,a) and (3,b) in Figure 8). P_1 updates its set by the PRF value and receives the message $(F_1(x_2), \mathsf{Enc}(x_2))$ from P_2. P_2 invokes cOPRF protocol with P_3 and P_4 concurrently. P_3 receives a random value r_3 since $x_3 = x_2$ and P_4 receives the PRF value $F_1(x_4)$ for x_4. The following rounds are similar.

4 Our mPSU Construction

Figure 8 presents our main mPSU protocol, which guarantees security against any number of corrupt, colluding, semi-honest parties. The protocol makes use of our new mOT, cOPRF and Shuffle&Decrypt gadgets.

4.1 Our Protocol

The design of a secure mPSU protocol presents significant challenges, specifically with regard to (1) ensuring that the output does not contain duplicate items, (2) preventing the disclosure of partial union results, and (3) hiding which items from which parties. To illustrate the high-level idea of our protocol, we consider a simple case where the leader party P_1 has a set X_1 of items while the remaining parties $P_2, \ldots, P_n$, each possess a single item $X_i = \{x_i\}$. We assume that the item x_2 of P_2 and x_3 of P_3 are not in the X_1 but x_4 of P_4 is (i.e. $x_2, x_3 \notin X_1$ and $x_4 \in X_1$).

Regarding (1), a potential approach is to enable the leader P_1 to engage with the other parties and obtain an encryption of x_i if $x_i \notin X_1$ and an encryption of the zero otherwise. An encryption of zero indicates the presence of common items between P_1 and P_i, which can be removed after decryption. To this end, the P_1 and P_i invoke the mOT instance in which P_i acts as the sender with input $\big(x_i, \mathsf{Enc}(\mathsf{pk}, 0), \mathsf{Enc}(\mathsf{pk}, x_i)\big)$ and P_1 acts the receiver with input X_1,

thereby obtaining the desired encryption. After executing the mOT instances, the leader party P_1 acquires $E = \{\mathsf{Enc}(\mathsf{pk}, x_2), \mathsf{Enc}(\mathsf{pk}, x_3), \mathsf{Enc}(\mathsf{pk}, 0)\}$ from the party P_2, P_3 and P_4, respectively. The set E allows the leader P_1 to obtain the set union after decryption.

The above protocol description does not entirely address the issue of removing duplicate items since x_2 could be identical to x_3. To overcome the limitation, we leverage the OPRF in the following manner. The leader party P_1 with input X_1 interacts with P_2 holding the PRF key k and receives the PRF values $S = \{y \mid y = F(k, x), x \in X_1\}$. They then execute the mOT where P_2 has keyword $F(k, x_2)$ and messages $(r_2||\mathsf{Enc}(\mathsf{pk}, 0), F(k, x_2)||\mathsf{Enc}(\mathsf{pk}, x_2))$ so that P_1 can obliviously obtain $v_2||e_2$ from P_2 where $v_2||e_2$ equals $F(k, x_2)||\mathsf{Enc}(\mathsf{pk}, x_2)$ if $x_2 \notin X_1$, and $r_2||\mathsf{Enc}(\mathsf{pk}, 0)$ otherwise for a random r_2. The P_1 appends v_2 to S and e_2 to an empty set E. At this point, $S = \{y \mid y = F(k, x), x \in X_1\} \cup \{F(k, x_2)\}$ if $x_2 \notin X_2$, and $S = \{y \mid y = F(k, x), x \in X_1\} \cup \{r_2\}$ otherwise. The updated set S will be used as the input to the next mOT between P_1 and P_3. It helps to remove the duplication items of P_2 and P_3 from the final output. Concretely, P_3 prepare the keyword and messages in the same way for the mOT, the functionality checks the membership condition whether $F(k, x_3)$ is in the updated S and then gives P_1 the corresponding result $v_3||e_3$ which equals to either $r_3||\mathsf{Enc}(\mathsf{pk}, 0)$ for a random r_3 or $F(k, x_3)||\mathsf{Enc}(\mathsf{pk}, x_3)$. If $x_2 = x_3 \notin X_1$, the P_1 obtains the encryption $e_2 = \mathsf{Enc}(\mathsf{pk}, x_2)$ from P_2, but $e_3 = \mathsf{Enc}(\mathsf{pk}, 0)$ from P_3. The r_3 is indistinguishable from the PRF value $F(k, x_3)$, thus r_3 reveals nothing to P_1 about whether $x_2 = x_3$. The P_1 continues to update the PRF set S and the encryption set E by repeating the above process sequentially with other parties $P_4, \ldots, P_n$.

The above description ignores an important security concern – in particular, when P_1 and P_3 collude and the PRF key is belongs to one party (e.g., P_2), the adversary learns P_2's input as $F(k, x_2)$ appears in S, but $F(k, x_3)$ does not. To address this vulnerability, e.g., the challenge (2), we propose to use our new gadget cOPRF as described in Section 4.1.1.

Upon the completion of $(n-1)$ instances of the mOT protocol between the leader party P_1 and other parties P_i, the leader P_1 has acquired an encryption set E, containing encryptions $\mathsf{Enc}(\mathsf{pk}, x)$ for $x \in \bigcup_{i=1}^{n} X_i$ and τ encryptions of zero, where $\tau = \sum_{i=1}^{n} |X_i| - |\bigcup_{i=1}^{n} X_i|$ indicates the number of duplicate items. In order to satisfy requirement (3) of the mPSU protocol, the encryption set E remains secure and cannot be decrypted by any subset of parties. We utilize a multi-key cryptosystem, as detailed in Section 2.6. This cryptosystem requires the participation of all parties in the decryption process. Moreover, to prevent any compromised parties from ascertaining which elements are held by each party, partial decryption results must be permuted by each party prior to transmission to the others. To achieve this, we employ the $\mathsf{Shuffle\&Decrypt}$ functionality, which permits each party to apply its own permutation function to the partial decryption results. We demonstrate our mPSU protocol execution in Table 2 pertaining to the scenario involving four parties.

At this stage, we are currently focusing on a simple scenario where each $P_{i\in[2,n]}$ possesses only one item. In order to generalize our method to a set X_i, we apply a popular technique known as bin-and-ball technique. At the high level, the party $P_{i\in[2,n]}$ places its input values into β bins through the use of Cuckoo hashing, where each bin is allowed to contain at most one item. The leader P_1 utilizes the same set of Cuckoo hash functions to map the input values in S into β bins using Simple hashing. The mapping allows the parties to execute the simple case above bin-by-bin efficiently. As a result, for each bin, the P_1 obtains encryptions of the partial union set which are subsequently combined into a big encryption set E before being subjected to decryption.

4.1.1 The Usage of cOPRF As previously mentioned, when the set S of the leader party P_1 contains the PRF values of the union items but the PRF is under a single key k generated by a party, the adversary can learn partial information regarding the P_i's input set X_i if it colludes P_1 and any P_t. To overcome this limitation, the values added to set S after each mOT instance with P_i must not reveal any information about the set X_i. To accomplish this, we propose using the multi-key PRF $(k_i, \ldots, k_2)$ where each key k_t is chosen by $P_{t\in[2,i]}$. Assume that P_i can compute $y_i = F((k_i, \ldots, k_2), x_i)$ for each item $x_i \in X_i$ (we describe how to compute the PRF value y_i in detail in the following paragraph). In that case, by invoking mOT, P_i can securely send y_i to P_1 if the item x_i is not in the union of previous sets $\{X_1, \ldots, X_{i-1}\}$ and random otherwise. The key k_i is kept secret and only known by P_i, which protects colluding parties from learning the input set of P_i.

There is a simple way to enable the computation of the PRF values by party P_i. This involves a sequential interaction between P_i and other parties, $P_{t\in[2,i-1]}$, via an OPRF. During this interaction, P_t acts as the sender with its PRF key k_t, and P_i acts as the receiver with $F((k_{t-1}, \ldots, k_2), x_i), \forall x_i \in X_i$ or x_i if $t = 2$. The receiver P_i then receives $F(k_t, F((k_{t-1}, \ldots, k_2), x_i))$ which we denote as $F((k_t, \ldots, k_2), x_i)$, while the sender P_t learns nothing. However, the use of multi-key PRF values could reveal information about X_t when P_1 and $P_{t+1}, \ldots, P_i$ were to collude. For simplicity, we assume that P_1 and $P_{i>t}$ engage in collusion. If the item x_i of P_i is identical to the item $x_t \in X_t$, but does not appear in any set $X_{j\in[1,t-1]}$, the corrupt leader party P_1 would obtain $y_t = F((k_t, \ldots, k_2), x_t)$ from the honest party P_t through the mOT instance. Additionally, the corrupt P_i would also possess $y_i = F((k_t, \ldots, k_2), x_i)$ after engaging in interactions with each $P_2, \ldots, P_t$ via OPRF. By comparing the two PRF values y_t and y_i, the adversary (i.e., the corrupt P_1 and P_i) would be able to obtain partial information about X_t. To address the leakage, we ensure that the corrupt party P_i obtains a "fake" PRF value of their input item x_i if the x_i appears in any input set of previous parties $P_{t\in[2,i-1]}$. This objective can be achieved through the use of our cOPRF primitive.

Recall that our cOPRF is a single-query PRF, thus, we use the bin-and-ball technique to enhance the performance of our protocol. We now demonstrate how to execute the cOPRF using the recursive method. We assume that at the $i-1^{th}$

PARAMETERS:

- n parties $P_{i\in[n]}$ for $n > 1$.
- The mOT, OPRF, Shuffle&Decrypt functionalities described in Figures 2&9&6, respectively.
- A multi-key cryptosystem (KeyGen, Enc, ParDec, FulDec, ReRand) defined in Section 2.6.
- Hashing parameters: a number of bins μ, maximum bin sizes $\beta : \mathbb{Z} \to \mathbb{Z}$ for simple-hashing bins, the h hash functions $H_{j\in[h]} : \{0,1\}^\star \to [\mu]$.

INPUT:
- Party $P_{i\in[n]}$ has $X_i = \{x_{i,1}, \dots, x_{i,m}\}$.

PROTOCOL:

1. All n parties call the key generation algorithm $\mathsf{KeyGen}(1^\lambda, 1^\kappa)$. Each P_i receives a private key sk_i and a joint public key pk.
2. Local Execution:
   (a) $P_{i\in[2,n]}$ hashes items X_i into μ bins using the Cuckoo hashing. Let $C^i_{b,1}$ denote the items in the P_i's bth bin. P_i computes the encryption $e^i_{b,1} = \mathsf{Enc}(pk, C^i_{b,1})$, for $b \in [\mu]$.
   (b) $P_{i\in[n]}$ hashes X_i into μ bins under k hash functions. Let $S^i_{b,1}$ denote the set of items in the P_i's bth bin. P_i pads $S^i_{b,1}$ with dummy values to the maximum bin size $\beta(m)$.
   (c) For bin $b \in [\mu]$, the P_1 initials an empty set E_b.
3. P_1 sequentially interacts with P_i for $i \in [2, n]$ as follow.
   (a) For each bin $b \in [\mu]$, the P_1 and P_i invoke the functionality of OPRF where:
      - P_1 acts as the receiver with input the set $S^1_{b,i-1}$.
      - P_i acts as the sender with input the random-chosen PRF key k_i.
      - P_1 obtains a set $S^1_{b,i}$ of the PRF values $F(k_i, y)$ for $y \in S^1_{b,i-1}$.
   (b) For each bin $b \in [\mu]$, P_1 and P_i invoke a mOT instance where:
      - P_1 acts as the receiver with input $S^1_{b,i}$.
      - P_i acts as the sender with input $(y_{b,i}, r_{b,i}||\mathsf{Enc}(\mathsf{pk}, 0), y_{b,i}||\bar{e}_{b,i})$. Here, $r_{b,i}$ is a random value; $y_{b,i} = F(k_i, C^i_{b,i-1})$ if $C^i_{b,i-1} \neq \varnothing$ and random otherwise; $\bar{e}_{b,i} = \mathsf{Enc}(\mathsf{pk}, 0)$ if $C^i_{b,i-1} = \varnothing$, otherwise, $\bar{e}_{b,i} = e^i_{b,i-1}$.
      - P_1 obtains $v_b||e_b$.

      P_1 appends e_b to E_b. P_1 hashes $\bigcup_{b=1}^{\mu} \left(S^1_{b,i} \cup v_b\right)$ into μ bins under h hash functions. The P_1 redefines $S^1_{b,i}$ to be the set of items in its bth bin, and then pads $S^1_{b,i}$ with dummy values to the maximum bin size $\beta(im)$.
   (c) For each bin $b \in [\mu]$, P_i and $P_{t\in[i+1,n]}$ invoke the cOPRF where:
      - P_t acts as the receiver with input $C^t_{b,i-1}$ or a dummy if $C^t_{b,i-1} = \varnothing$.
      - P_i acts as the sender with input the PRF key k_i and the set $S^i_{b,i-1}$.
      - P_t obtains w_b, and sets $w_b = \varnothing$ if if $C^t_{b,i-1} = \varnothing$.

      P_t hashes $W = \{w_b \mid b \in [\mu]\ \&\ w_b \neq \varnothing\}$ into μ bins using the Cuckoo and Simple hashing. Let $S^t_{b,i}$ and $C^t_{b,i}$ denote the items in the Simple and Cuckoo b-th bin, respectively. P_t pads $S^t_{b,i}$ with dummy to maximum bin size $\beta(m)$.
   (d) The P_i and $P_{t\in[i+1,n]}$ invoke a mOT instance where:
      - P_i acts as the receiver with input $\{F(k_i, y) \mid y \in S^i_{b,i-1}\}$
      - P_t acts as the sender with input $(C^t_{b,i}, \mathsf{Enc}(\mathsf{pk}, 0), e^t_{b,i-1})$.
      - P_i obtains c and sends $c' = \mathsf{ReRand}(v, \mathsf{pk})$ to P_t

      P_t computes $e^t_{b,i} = \mathsf{ReRand}(c', \mathsf{pk})$
4. All the parties invoke the Shuffle&Decrypt functionality where:
   - P_1 inputs $E = \bigcup_{b=2}^{\mu} E_b$, the sk_1 and a random permutation $\pi_1 : [m] \to [m]$.
   - P_i inputs the private key sk_i and a random permutation $\pi_i : [m] \to [m]$.
   - P_1 obtains a set U.
5. P_1 removes all zero from U, and outputs $U \cup X_1$.

Fig. 8: Our mPSU Protocol

round of the protocol execution, the party $P_{t\in[i+1,n]}$ possesses a set $S^t_{b,i-1}$ for the b-th bin. This set contains $F((k_{i-1}\ldots,k_2),x_{t,j})$ if $x_{t,j}$ is not present in any set $X_{t\in[2,i-1]}$, and random values otherwise. For the baseline with $i=2$, $S^t_{b,1}$ corresponds to the input set of P_t, which was mapped to the b-th Simple-hashing bin. To compute the desired PRF values which contain the key k_i, the P_i and P_t invoke a cOPRF in which P_i acts as the sender with input $S^i_{b,i-1}$ and the PRF key k_i while P_t acts as the receiver and queries on the input $c\in S^t_{b,i-1}$. The output of cOPRF provides P_t the correct PRF value $F(k_i,c)$. This value equals $F((k_i,k_{i-1},\ldots,k_2),x_{t,j})$ for an input $x_{t,j}$ if $x_{t,j}\notin X_i$ and c has a form $F((k_{i-1},\ldots,k_2),x_{t,j})$ – in other words, $x_{t,j}\notin X_2\cup\ldots\cup X_{i-1}$. The cOPRF gives P_t a random (i.e., a "fake" PRF value) if $x_{t,j}\in X_i$ or c is random. We present the cOPRF execution in Step (3,c), Figure 8.

When using the hash-to-bin scheme, parties are required to apply the mapping to their multi-key PRF values. Additionally, each individual $P_{i\in[2,n]}$ must also map its PRF values using Simple hashing, and the resulting Simple-hashing bin serves as the input of the sender P_i in the cOPRF process. Since the P_1 and P_i invoke mOT in Step (3,b), thus these parties only need to execute the OPRF computation (instead of cOPRF). Note that P_1 does not need to query the PRF values for dummy items, but instead aggregates all non-dummy PRF values in the set $S^1_{b,i-1}$ and employs them as input to the OPRF execution.

4.1.2 The Usage of mOT

As per the overview description, the parties P_1 and P_i engage in the mOT to incrementally acquire the PRF values and encrypted union items. As the PRF values are generated under distinct keys in this protocol, the input for the mOT must differ from what we described before. We introduce two modifications to the input as follows.

The first modification is that the P_1 has to compute multi-key PRF values for its set $S^1_{b,i}$ for $i\in[2,n]$. It can be done easily by executing an OPRF with P_i which holds the PRF key k_i. Concretely, assuming that after the mOT with P_{i-1}, the set $S^1_{b,i-1}$ contains either random values r or $F((k_{i-1},\ldots,k_2),x)$ for $x\in\bigcup_{t=1}^{i-1}X_i$. In the case where $i=2$, the $S^1_{b,i-1}$ is indeed the P_1's input set. The P_1 submits OPRF queries on S^1_{i-1} to P_i and obtains $S^1_{b,i}=\{F(k_i,s)\mid s\in S^1_{i-1}\}$ whiles P_i learns nothing. Clearly, the set $S^1_{b,i}$ consists of $F((k_i,\ldots,k_2),x)$ for $x\in\bigcup_{t=1}^{i-1}X_i$ and $F(k_i,r)$ for random $r\in S^1_{i-1}$. We present the OPRF in Step (3,a), Figure 8.

The mOT execution between P_1 and P_i should allow P_1 to add the PRF values $F((k_i,\ldots,k_2),x)$ of the P_i's item $x\in X_i$ to $S^1_{b,i}$ if x is not in the union of $\{X_1,\ldots,X_{i-1}\}$. To achieve this, P_i should use input $(y_{b,i},r_{b,i}||\mathsf{Enc}(\mathsf{pk},0),y_{b,i}||\bar{e}_{b,i})$ to the mOT execution. The $y_{b,i}$ is $F(k_i,c)$ where c is either random or has a form of the multi-key PRF $F((k_{i-1},\ldots,k_2),x))$ obtained from cOPRF executions with previous parties $P_{i\in[2,i-1]}$. The $r_{b,i}$ is randomly chosen. The main tricky part is how to define $\bar{e}_{b,i}$. If $\bar{e}_{b,i}$ is an encryption $\mathsf{Enc}(\mathsf{pk},x)$ of the P_i's input item x, and if the the $y_{b,i}$ is a "fake" PRF value, then $\bar{e}_{b,i}$ is added to E_b as $y_{b,i}$ never appears in the P_1's set $S^1_{b,i}$. Recall that the "fake" PRF $y_{b,i}$ indicates that the correspond-

ing x is in the set of previous parties $P_{t\in[2,i-1]}$. Hence, if we set $\bar{e}_{b,i} = \mathsf{Enc}(\mathsf{pk}, x)$, the E_b contains two ciphertexts that are encryptions of the same x. This will reveal to P_1 the multiplicity of each element in the union after decrypting E_b. To avoid this issue, we propose the following modification on how to compute $\bar{e}_{b,i}$.

Our objective is to ensure that $\bar{e}_{b,i}$ is the encryption of the P_i's input item x if x does not appear in any set $X_2, \ldots, X_{i-1}$; otherwise it should be the encryption of zero. We achieve this by having P_i participate with each party $P_2, \ldots, P_{i-1}$ using mOT. For clarity, we describe the scenario for $i = 3$ as follows. Initially, P_3 computes $e^3_{b,1} = \mathsf{Enc}(\mathsf{pk}, x)$ for its item x. Next, the party P_2 and P_3 invoke a mOT instance where P_2 acts as the receiver with input $S^2_{b,2} = \{F(k_2, y) \mid y \in S^2_{b,1}\}$ and P_3 acts as the sender with input $(C^3_{b,2}, \mathsf{Enc}(\mathsf{pk}, 0), e^3_{b,1})$. Here, $C^3_{b,2} = F(k_2, x)$ if $x \notin X_2$ or random otherwise (the value $C^3_{b,2}$ is the output of the cOPRF execution between the sender P_2 and the receiver P_3). Therefore, the mOT functionality returns the vaue c to the P_2, which is the $\mathsf{Enc}(\mathsf{pk}, 0)$ if $x \in X_2$ and the $\mathsf{Enc}(\mathsf{pk}, x)$ if $x \notin X_2$. The P_2 then rerandomizes c by computing $c' \leftarrow \mathsf{ReRand}(c, \mathsf{pk})$ and returns the result to P_3. The rerandomization aims to prevent the P_3 from determining the output of P_2, which would reveal whether $x \in X_2$. The P_3 then computes $e^3_{b,2} = \mathsf{ReRand}(c', \mathsf{pk})$, and we define $\bar{e}_{b,3}$ as $e^3_{b,2}$, which is the input to mOT. The OPRF is presented in Step (3,d), Figure 8.

4.2 Correctness and Security

Correctness. We consider three following cases depending on whether a specific item $x_{i,k} \in X_i$ of the smallest-index party P_i is in P_1 or other parties P_t for $n \geq t > i > 1$. Since P_i is the smallest index that has $x_{i,k}$, no previous parties have $x_{i,k}$. Thus, P_i obtains $C^i_{b,i-1} = F(k_{i-1}, \ldots, k_2, x_{i,k})$ after interacting with $P_{t\in[2,i-1]}$ via the cOPRF.

- Case 1 ($x_{i,k} \in X_1$) – the P_1 has $x_{i,k}$: As $x_{i,k} \in X_1$, the OPRF with $P_{t\in[2,i]}$ in Step (3,a) gives P_1 the multi-key PRF value $F((k_i, \ldots, k_2), x_{i,k})$. In the mOT execution between P_1 and P_i, the input keyword of P_i as $y_{b,i} = F(k_i, C^i_{b,i-1})$ is in $S^1_{b,i}$. Thus, P_1 receives the encryption of zero $\mathsf{Enc}(\mathsf{pk}, 0)$ from the mOT functionality. As a result, $x_{i,k}$ does not appear in the final result from the $\mathsf{Shuffle\&Decrypt}$ execution.
- Case 2 ($x_{i,k} \notin X_1$ and $x_{i,k} \in X_t$) – the P_1 does not have $x_{i,k}$, but another party P_t has $x_{t,j} = x_{i,k}$ for $t > i$: The mOT execution between P_1 and P_i on input related to $y_{b,i} = F(k_i, C^i_{b,i-1})$ gives P_1 the $y_{b,i}$ and $e_{b,i} = \mathsf{Enc}(\mathsf{pk}, x_{i,k})$. The PRF value $y_{b,i}$ is added to the set $S^1_{b,i}$, and the $x_{i,k}$ will appear in the final union output.

  Assume that $x_{t,j} = x_{i,k} \in X_j$ was mapped into the b-th Cuckoo bin. Since $x_{t,j} \in X_i$, the cOPRF and mOT with P_i gives P_j a random "fake" PRF value $C^t_{b,i}$ and $e_{t,b} = \mathsf{Enc}(\mathsf{pk}, 0)$. Thus, when executing mOT with P_t, the P_1 obtains $\mathsf{Enc}(\mathsf{pk}, 0)$. Hence, the $x_{t,j}$ does not appear in the final result.

– Case 3 ($x_{i,k} \notin \bigcup_{j=1,j\neq i}^{n} X_j$) – no party has $x_{i,k}$: The mOT execution between P_1 and P_i gives P_1 an $\mathsf{Enc}(\mathsf{pk}, x_{i,k})$. Thus, $x_{i,k}$ appears in the final result from the Shuffle&Decrypt functionality.

Security. The security of our mPSU protocol is given as below.

Theorem 4. *Given the multi-key cryptosystem, mOT, OPRF, cOPRF and Shuffle&Decrypt functionalities described in Section 2.6, and Figures 2&9&6, respectively, the mPSU protocol described in Figure 8 securely implements the mPSU functionality defined in Figure 1 in the semi-honest model, against any number of corrupt, colluding, semi-honest parties.*

Proof. Let C and H be a coalition of corrupt and honest parties, respectively. We must show how to simulate C's view in the ideal model. We consider three following cases based on whether C has an item x:

1. C does not have x, but H has x: If H contains only one honest party P_i, then P_i has x. The corrupted parties C can deduce that the honest party P_i has x from the output of the set union. Hence, there is nothing to hide about whether P_i has x in this case. If H has more than one honest party (say P_i and $P_{j>i}$). We consider two following cases:
   – Only P_i has x: we must show that the protocol must hide the identity of P_i. If $P_1 \in H$, only the honest party P_1 learns the union $\bigcup_{i=1}^{n} X_i$ in Step 5. In addition, the cOPRF and mOT between P_i and previous corrupt parties $P_{t<i} \in C$ reveals nothing to C. Thus, the simulation is simple.
   If $P_1 \in C$, the corrupt P_1 obtains $\mathsf{Enc}(\mathsf{pk}, x)$ and the $F((k_i, \dots, k_2), x)$ from P_i. Since the encryption is protected under the Shuffle&Decrypt functionality until the P_1 learns the union sets which was permuted by the honest party P_i, the encryption reveals nothing to C. Similarly, the PRF value consists of the P_i's key k_i, which hides x from C.
   – Both P_i and P_j have x: If $i = 1$, then the honest leader P_1 receives encryptions of zeros $\mathsf{Enc}(\mathsf{pk}, 0)$ and "fake" PRF values when executing mOT with P_j. Thus, the C learns nothing about which parties in H have x. If $P_1 \in C$, the corrupt P_1 receives the encryption $\mathsf{Enc}(\mathsf{pk}, x)$ from P_i and $\mathsf{Enc}(\mathsf{pk}, 0)$ from P_j. Thanks to the CCA property of the encryption scheme and the permutation in Shuffle&Decrypt, C cannot distinguish the two encryptions. Thus, the protocol hides the identity of which honest party has x.
2. C have x, but H does not have x: We must show that the protocol must hide the information that H does not have x. Consider the cOPRF (or OPRF) and mOT executions where a party in H acts as the sender and a party in C acts as the receiver, the corrupt set C receives nothing related to x. In the final step, the encryption set E contains $\mathsf{Enc}(\mathsf{pk}, x)$, which was permuted by the honest parties H. Hence, all honest parties have an indistinguishable effect on the Shuffle&Decrypt step.

Our Protocol		$m=2^8$			$m=2^{12}$			$m=2^{16}$			$m=2^{20}$		
		$t=1$	$t=4$	$t=16$	$t=1$	$t=4$	$t=16$	$t=1$	$t=4$	$t=16$	$t=1$	$t=4$	$t=16$
LAN (s)	$n=3$	2.34	0.68	0.36	36.19	9.30	2.91	601.18	153.76	41.99	5175.56	1395.41	378.74
	$n=4$	4.41	1.23	0.60	69.64	17.50	4.98	1147.32	291.31	78.60	11468.34	3046.06	826.02
	$n=6$	10.62	2.89	1.18	167.80	42.44	11.75	2786.67	724.16	189.26	30266.82	8351.98	2175.02
	$n=8$	19.68	5.24	1.96	312.52	78.98	21.49	5176.67	1331.20	349.87	57776.16	15680.08	4130.61
WAN (s)	$n=3$	12.03	10.58	10.32	52.38	26.56	20.48	655.97	208.69	96.97	9920.07	2961.36	1219.62
	$n=4$	13.28	11.35	10.98	94.21	43.12	30.91	1214.65	370.88	161.04	18937.01	5444.82	2088.87
	$n=6$	36.72	29.04	27.36	210.05	85.91	55.57	2909.79	859.65	327.68	45732.48	12943.89	4434.85
	$n=8$	56.70	42.36	39.13	371.11	139.78	82.74	5356.24	1523.28	544.93	84535.25	23227.10	7576.88
Comm. Cost (MB)	$n=3$			1.52			21.11			344.83			5517.28
	$n=4$			2.29			31.88			505.67			8090.72
	$n=6$			3.83			53.47			848.28			13572.48
	$n=8$			5.38			75.12			1191.67			19066.72

Table 3: The running time and communication cost of our mPSU protocol: the number of parties $n \in \{3, 4, 6, 8\}$, set size $m \in \{2^8, 2^{12}, 2^{16}, 2^{20}\}$, and numbers of thread $t = \{1, 4, 16\}$. The reported running time represents the time taken for the entire protocol to complete. Communication cost is computed as the average cost across all parties.

3. Both C and H have x. When C acts as the receiver invokes the cOPRF (or OPRF) with an honest sender which does not have x, the C obtains the correct PRF values. When C interacts with an honest party that has x, the C obtains the PRF values (if $P_1 \in C$) and the "fake" PRF values. Since the PRF values contain the PRF key of the honest set H and their distribution is random. Thus, C learns nothing from cOPRF or OPRF executions. Similarly, the corrupt coalition's view is simulated from Step (3b, 3d) based on the functionality of mOT and encryption scheme. Moreover, the $\mathsf{Enc}(\mathsf{pk}, x)$ appears only once in the encryption set E, thus, C learns nothing about whether H has x.

5 Implementation and Performance

We implement our protocol and evaluate it with various number of parties, set sizes, and number of threads. All evaluations were performed with an item input length 128 bits, a statistical security parameter $\lambda = 40$, and a computational security parameter $\kappa = 128$. We do a number of experiments on a single server that has AMD EPYC 74F3 processors and 256GB of RAM. We run all parties in the same network, but simulate a network connection using the Linux tc command: a LAN setting with 0.02ms round-trip latency, 10 Gbps network bandwidth; a WAN setting with a 80ms round-trip latency, 400 Mbps network bandwidth.

Our mPSU protocol is built on Elgamal encryption scheme (multi-key cryptosystem), Diffie-Hellman OPRF (cOPRF), SS-PMT, and OT (mOT and cOPRF). We implement the exponentiation for OPRF and Elgamal encryption using the elliptic curve code (Curve25519) from Relic [40]. For the SS-PMT implementation which requires garbled circuit for two strings comparison, we use the EMP-toolkit library [45]. Finally, we use the OT-extension [23] provided in [33] to implement mOT. Our complete implementation will be available on GitHub.

	m	Our	[19]	[5]	[16]
Running Time LAN (second)	2^8	4.41	155.63	**2.70**	6009.00
	2^{12}	69.64	2490.04	**57.725**	-
	2^{16}	**1147.32**	39840.65	1158.53	-
	2^{20}	**11468.34**	637450.40	25279.39	-
Running Time WAN (second)	2^8	**13.28**	-	128.05	-
	2^{12}	**94.21**	-	2387.70	-
	2^{16}	**1214.65**	-	45939.46	-
Comm. Cost (MB)	2^8	**2.29**	15.72	740.67	2.50
	2^{12}	**31.88**	251.65	15058.25	40.00
	2^{16}	**505.67**	4026.53	308523.50	640.00
	2^{20}	**8090.72**	64424.48	6279871.50	10240.00

Table 4: Performance comparison of different mPSU protocols with $n = 4$ parties, each having $m \in \{2^8, 2^{12}, 2^{16}, 2^{20}\}$.

Our protocol scales well using multi-threading between the parties. In each round, the party $P_{i\in[2,n-1]}$ can use $n-i+1$ threads so that each party operates OPRF with mOT and cOPRF building blocks with other parties P_1 and $P_{j\in[i+1,n]}$ at the same time. In addition, each pair of parties can use multiple threads to execute these building blocks bin-by-bin in parallel. We evaluate it on number of threads $t \in \{1, 4, 16\}$ to show the performance of our protocols running with multi-threading. Table 3 presents the overall running time and communication overhead of our mPSU protocol.

Comparison with Previous Work To demonstrate the performance of our mPSU protocols with a comparison, we have implemented the semi-honest protocols proposed in [16,5] and estimate the performance for protocol proposed in [19]. Table 4 presents the running time and communication cost of various mPSU protocols [16,5,19] which are secure in the dishonest majority [1] and semi-honest setting. We do not incorporate the results from [44] into our comparison, as their protocol only work for a small universe. Even in their largest setting, with a universe size of 2^{32}, it is considerably smaller than the general scenario involving 128-bit elements. According to [44, Figure 7], in a scenario involving 5 parties, each with only 32 elements of 32-bit length, their protocol takes around 10 seconds. Interestingly, this is comparable to the runtime of our protocol involving 6 parties, each with 256 elements in a 128-bit universe.

[1] The recent mPSU protocol [29] provides a weak security guarantee wherein the leader does not collude with any parties.

In [5], each input set X_i is initially shared among n parties using a secret-sharing scheme. Subsequently, these parties employ a generic secure computation technique to compute the union on the shares. Our implementation of the [5]'s method, however, is limited to the two-party scenario where each X_i is secret-shared between only two parties (which is in favor of [5]). Consequently, the secure computation takes place exclusively between these two parties. We implement [5] using EMP-toolkit library [45] which provides the most of the state-of-the-art techniques for two-party secure computation in the semi-honest setting. As shown in Table 4, for $n = 4$, our protocol is 2.20 times faster for the large set size of 2^{20} in the LAN setting and $9.64 - 37.82\times$ faster than [5] in the WAN. Additionally, our protocol incurs an average communication cost of 505.67 MB per party, whereas the cost for [5] is significantly $(323 - 776\times)$ higher.

We report the partial running time and communication cost of the mPSU protocol proposed by [19]. The first step of their protocol is for each party to locally compute an encryption of a local Bloom filter. To achieve a false positive rate of 2^{-40}, the table size should be at least $60nm$. We estimate the time and communication cost for this single step of each party based on the performance shown in [30] (as well as our [16]'s implementation), where each Paillier encryption takes about 2.5 ms with a key length of 2048 bits, and report the numbers in Table 4.

Our mPSU protocol outperforms previous works in the LAN setting. Despite the low communication cost due to the usage of homomorphic encryption, the running time of [16,19] is not practical even for small set sizes. Thus, we skip the evaluation of the [19,16] in the WAN setting.

References

1. Shuffle-based private set union: Faster and more secure. In *31st USENIX Security Symposium (USENIX Security 22)*, Boston, MA, August 2022. USENIX Association.
2. Gilad Asharov, Abhishek Jain, Adriana López-Alt, Eran Tromer, Vinod Vaikuntanathan, and Daniel Wichs. Multiparty computation with low communication, computation and interaction via threshold FHE. In David Pointcheval and Thomas Johansson, editors, *EUROCRYPT 2012*, volume 7237 of *LNCS*, pages 483–501. Springer, Heidelberg, April 2012.
3. Alex Berke, Michiel Bakker, Praneeth Vepakomma, Kent Larson, and Alex 'Sandy' Pentland. Assessing disease exposure risk with location data: A proposal for cryptographic preservation of privacy, 2020.
4. Alexander Bienstock, Sarvar Patel, Joon Young Seo, and Kevin Yeo. Near-Optimal oblivious Key-Value stores for efficient PSI, PSU and Volume-Hiding Multi-Maps. In *32nd USENIX Security Symposium (USENIX Security 23)*, pages 301–318, Anaheim, CA, August 2023. USENIX Association.
5. Marina Blanton and Everaldo Aguiar. Private and oblivious set and multiset operations. In Heung Youl Youm and Yoojae Won, editors, *ASIACCS 12*, pages 40–41. ACM Press, May 2012.

6. Zvika Brakerski. Fully homomorphic encryption without modulus switching from classical GapSVP. In Reihaneh Safavi-Naini and Ran Canetti, editors, *CRYPTO 2012*, volume 7417 of *LNCS*, pages 868–886. Springer, Heidelberg, August 2012.
7. Justin Brickell and Vitaly Shmatikov. Privacy-preserving graph algorithms in the semi-honest model. In Bimal K. Roy, editor, *ASIACRYPT 2005*, volume 3788 of *LNCS*, pages 236–252. Springer, Heidelberg, December 2005.
8. Prasad Buddhavarapu, Andrew Knox, Payman Mohassel, Shubho Sengupta, Erik Taubeneck, and Vlad Vlaskin. Private matching for compute. Cryptology ePrint Archive, Paper 2020/599, 2020. https://eprint.iacr.org/2020/599.
9. Dung Bui and Geoffroy Couteau. Improved private set intersection for sets with small entries. PKC, 2023. https://eprint.iacr.org/2022/334.
10. Martin Burkhart, Mario Strasser, Dilip Many, and Xenofontas Dimitropoulos. SEPIA: Privacy-Preserving aggregation of Multi-Domain network events and statistics. In *19th USENIX Security Symposium (USENIX Security 10)*, Washington, DC, August 2010. USENIX Association.
11. Nishanth Chandran, Nishka Dasgupta, Divya Gupta, Sai Lakshmi Bhavana Obbattu, Sruthi Sekar, and Akash Shah. Efficient linear multiparty PSI and extensions to circuit/quorum PSI. In Giovanni Vigna and Elaine Shi, editors, *ACM CCS 2021*, pages 1182–1204. ACM Press, November 2021.
12. Daniel Demmler, Peter Rindal, Mike Rosulek, and Ni Trieu. Pir-psi: Scaling private contact discovery. Cryptology ePrint Archive, Paper 2018/579, 2018. https://eprint.iacr.org/2018/579.
13. W. Diffie and M. Hellman. New directions in cryptography. *IEEE Transactions on Information Theory*, 22(6):644–654, 1976.
14. Taher ElGamal. A public key cryptosystem and a signature scheme based on discrete logarithms. In G. R. Blakley and David Chaum, editors, *CRYPTO'84*, volume 196 of *LNCS*, pages 10–18. Springer, Heidelberg, August 1984.
15. Michael J. Freedman, Yuval Ishai, Benny Pinkas, and Omer Reingold. Keyword search and oblivious pseudorandom functions. In Joe Kilian, editor, *Theory of Cryptography*, pages 303–324, Berlin, Heidelberg, 2005. Springer Berlin Heidelberg.
16. Keith B. Frikken. Privacy-preserving set union. In Jonathan Katz and Moti Yung, editors, *ACNS 07*, volume 4521 of *LNCS*, pages 237–252. Springer, Heidelberg, June 2007.
17. Gayathri Garimella, Payman Mohassel, Mike Rosulek, Saeed Sadeghian, and Jaspal Singh. Private set operations from oblivious switching. In Juan A. Garay, editor, *Public-Key Cryptography – PKC 2021*, pages 591–617, Cham, 2021. Springer International Publishing.
18. Craig Gentry. Fully homomorphic encryption using ideal lattices. In *Proceedings of the Forty-First Annual ACM Symposium on Theory of Computing*, STOC '09, page 169–178, New York, NY, USA, 2009. Association for Computing Machinery.
19. Xuhui Gong, Qiang-Sheng Hua, and Hai Jin. Nearly optimal protocols for computing multi-party private set union. In *2022 IEEE/ACM 30th International Symposium on Quality of Service (IWQoS)*, pages 1–10, 2022.
20. Christoph Hagen, Christian Weinert, Christoph Sendner, Alexandra Dmitrienko, and Thomas Schneider. Contact discovery in mobile messengers: Low-cost attacks, quantitative analyses, and efficient mitigations. *ACM Trans. Priv. Secur.*, 26(1), nov 2022.
21. Kyle Hogan, Noah Luther, Nabil Schear, Emily Shen, David Stott, Sophia Yakoubov, and Arkady Yerukhimovich. Secure multiparty computation for cooperative

cyber risk assessment. In *2016 IEEE Cybersecurity Development (SecDev)*, pages 75–76, 2016.

22. Mihaela Ion, Ben Kreuter, Ahmet Erhan Nergiz, Sarvar Patel, Shobhit Saxena, Karn Seth, Mariana Raykova, David Shanahan, and Moti Yung. On deploying secure computing: Private intersection-sum-with-cardinality. In *IEEE European Symposium on Security and Privacy, EuroS&P 2020, Genoa, Italy, September 7-11, 2020*, pages 370–389. IEEE, 2020.
23. Yuval Ishai, Joe Kilian, Kobbi Nissim, and Erez Petrank. Extending oblivious transfers efficiently. In Dan Boneh, editor, *CRYPTO 2003*, volume 2729 of *LNCS*, pages 145–161. Springer, Heidelberg, August 2003.
24. M. Kantarcioglu and C. Clifton. Privacy-preserving distributed mining of association rules on horizontally partitioned data. *IEEE Transactions on Knowledge and Data Engineering*, 16(9):1026–1037, 2004.
25. Lea Kissner and Dawn Xiaodong Song. Privacy-preserving set operations. In Victor Shoup, editor, *CRYPTO 2005*, volume 3621 of *LNCS*, pages 241–257. Springer, Heidelberg, August 2005.
26. Vladimir Kolesnikov, Ranjit Kumaresan, Mike Rosulek, and Ni Trieu. Efficient batched oblivious PRF with applications to private set intersection. In Edgar R. Weippl, Stefan Katzenbeisser, Christopher Kruegel, Andrew C. Myers, and Shai Halevi, editors, *ACM CCS 2016*, pages 818–829. ACM Press, October 2016.
27. Vladimir Kolesnikov, Mike Rosulek, Ni Trieu, and Xiao Wang. Scalable private set union from symmetric-key techniques. In Steven D. Galbraith and Shiho Moriai, editors, *ASIACRYPT 2019, Part II*, volume 11922 of *LNCS*, pages 636–666. Springer, Heidelberg, December 2019.
28. Tancrède Lepoint, Sarvar Patel, Mariana Raykova, Karn Seth, and Ni Trieu. Private join and compute from PIR with default. In Mehdi Tibouchi and Huaxiong Wang, editors, *ASIACRYPT 2021, Part II*, volume 13091 of *LNCS*, pages 605–634. Springer, Heidelberg, December 2021.
29. Xiang Liu and Ying Gao. Scalable multi-party private set union from multi-query secret-shared private membership test. Cryptology ePrint Archive, Paper 2023/1413, 2023. https://eprint.iacr.org/2023/1413.
30. Huanyu Ma, Shuai Han, and Hao Lei. Optimized paillier's cryptosystem with fast encryption and decryption. In *Annual Computer Security Applications Conference*, ACSAC '21, page 106–118, New York, NY, USA, 2021. Association for Computing Machinery.
31. Ofri Nevo, Ni Trieu, and Avishay Yanai. Simple, fast malicious multiparty private set intersection. In Yongdae Kim, Jong Kim, Giovanni Vigna, and Elaine Shi, editors, *CCS '21: 2021 ACM SIGSAC Conference on Computer and Communications Security, Virtual Event, Republic of Korea, November 15 - 19, 2021*, pages 1151–1165. ACM, 2021.
32. Duong Tung Nguyen and Ni Trieu. Mpccache: Privacy-preserving multi-party cooperative cache sharing at the edge. Cryptology ePrint Archive, Report 2021/317, 2021. https://eprint.iacr.org/2021/317.
33. Lance Roy Peter Rindal. libOTe: an efficient, portable, and easy to use Oblivious Transfer Library. https://github.com/osu-crypto/libOTe.
34. Benny Pinkas, Thomas Schneider, Gil Segev, and Michael Zohner. Phasing: Private set intersection using permutation-based hashing. In Jaeyeon Jung and Thorsten Holz, editors, *USENIX Security 2015*, pages 515–530. USENIX Association, August 2015.

35. Benny Pinkas, Thomas Schneider, Oleksandr Tkachenko, and Avishay Yanai. Efficient circuit-based PSI with linear communication. In Yuval Ishai and Vincent Rijmen, editors, *EUROCRYPT 2019, Part III*, volume 11478 of *LNCS*, pages 122–153. Springer, Heidelberg, May 2019.
36. Benny Pinkas, Thomas Schneider, Christian Weinert, and Udi Wieder. Efficient circuit-based PSI via cuckoo hashing. In Jesper Buus Nielsen and Vincent Rijmen, editors, *EUROCRYPT 2018, Part III*, volume 10822 of *LNCS*, pages 125–157. Springer, Heidelberg, April / May 2018.
37. Benny Pinkas, Thomas Schneider, and Michael Zohner. Scalable private set intersection based on ot extension. *ACM Trans. Priv. Secur.*, 21(2), jan 2018.
38. Tal Rabin. A simplified approach to threshold and proactive RSA. In Hugo Krawczyk, editor, *CRYPTO'98*, volume 1462 of *LNCS*, pages 89–104. Springer, Heidelberg, August 1998.
39. Srinivasan Raghuraman and Peter Rindal. Blazing fast PSI from improved OKVS and subfield VOLE. In Heng Yin, Angelos Stavrou, Cas Cremers, and Elaine Shi, editors, *ACM CCS 2022*, pages 2505–2517. ACM Press, November 2022.
40. RELIC. A modern research-oriented cryptographic meta-toolkit with emphasis on efficiency and flexibility. https://github.com/relic-toolkit.
41. Jae Hong Seo, Jung Cheon, and Jonathan Katz. Constant-round multi-party private set union using reversed laurent series. volume 7293, pages 398–412, 05 2012.
42. Katsunari Shishido and Atsuko Miyaji. Efficient and quasi-accurate multiparty private set union. In *2018 IEEE International Conference on Smart Computing (SMARTCOMP)*, pages 309–314, 2018.
43. Ni Trieu, Kareem Shehata, Prateek Saxena, Reza Shokri, and Dawn Song. Epione: Lightweight contact tracing with strong privacy. *CoRR*, abs/2004.13293, 2020.
44. Jelle Vos, Mauro Conti, and Zekeriya Erkin. Fast multi-party private set operations in the star topology from secure ands and ors. Cryptology ePrint Archive, Paper 2022/721, 2022. https://eprint.iacr.org/2022/721.
45. Xiao Wang, Alex J Malozemoff, and Jonathan Katz. Emp-toolkit: Efficient multiparty computation toolkit, 2016.
46. Cong Zhang, Yu Chen, Weiran Liu, Min Zhang, and Dongdai Lin. Optimal private set union from multi-query reverse private membership test. Cryptology ePrint Archive, Paper 2022/358, 2022. https://eprint.iacr.org/2022/358.

A Security Proof

A.1 Security Proof of Theorem 1

Proof. We construct simulators $\mathsf{Sim}_{\mathcal{S}}$ and $\mathsf{Sim}_{\mathcal{R}}$ to simulate the view of corrupted sender $\mathcal{S}$ and corrupted receiver $\mathcal{R}$, respectively. We argue the indistinguishability of the simulator and the real execution.

Simulating $\mathcal{S}$: The simulator $\mathsf{Sim}_{\mathcal{S}}$ has input (y, v_0, v_1) and receives output from the $\mathsf{SS\text{-}PMT}$ ideal functionality, consisting of a secret-shared membership bit $b_{\mathcal{S}}$. For the OT execution, the simulator $\mathsf{Sim}_{\mathcal{S}}$ obtains nothing, except the random OT transcript which is random. Since the output of $\mathsf{SS\text{-}PMT}$ is secret-shared amongst the corrupt sender and honest receiver, one can replace the bit $b_{\mathcal{S}}$ with a random. It is straightforward to check that the simulation is perfect.

Parameters: A PRF F, and a bound m on the number of queries.
Functionality:

- Wait for input $(q_1, \dots, q_m)$ from the receiver where $q_i \in \{0,1\}^\kappa$.
- Sample a random PRF key k and give it to the sender.
- Give $\{F(k, q_1), \dots, F(k, q_m)\}$ to the receiver.

Fig. 9: OPRF Ideal Functionality

Parameters: Two parties: Sender and Receiver
Functionality:

- Wait for input strings $(x_0, x_1) \subset (\{0,1\}^*)^2$ from the sender.
- Wait for input choice bit $b \in \{0,1\}$ from the receiver.
- Give x_b to the receiver.

Fig. 10: Oblivious Transfer (OT) Ideal Functionality.

Parameters: Two parties: P_0 and P_1, and the set size n.
Functionality:

- Wait for input a set of items $X = \{x_1, \dots, x_n\} \subset (\{0,1\}^*)^n$ from the P_0.
- Wait for input item $y \in \{0,1\}^\star$ from the P_1.
- Give b_i to the $P_{i \in \{0,1\}}$ where $b_0 \oplus b_1 = 1$ if $y \in X$ and 0 otherwise.

Fig. 11: Secret-shared Private Membership Test (SS-PMT) Ideal Functionality.

Simulating $\mathcal{R}$: $\mathsf{Sim}_\mathcal{R}$ with input X receives nothing from the SS-PMT ideal functionality, expect a secret-shared membership bit $b_\mathcal{R}$. $\mathsf{Sim}_\mathcal{R}$ obtains w from the OT and u from the sender in the last step. We show that the output of the simulator $\mathsf{Sim}_\mathcal{R}$ is indistinguishable from the real execution. For this, we formally show the simulation by proceeding with the sequence of hybrid transcripts T_0, T_1, T_2 where T_0 is real view of the receiver, and T_2 is the output of $\mathsf{Sim}_\mathcal{R}$.

- Let T_1 be the same as T_0, except the SS-PMT output which can be replaced with random as the honest sender holds a secret-shared of the output. Thus, T_0 and T_1 are indistinguishable.
- Let T_2 be the same as T_1, except the OT execution and obtaining u. Due to the underlying security property of OT, the receiver only learns one of the two strings related to v_0 or v_1. In addition, the sender's associated values were masked with a random value r before the OT execution. Thus, w reveals nothing about $v_{i \in \{0,1\}}$. When having $u = r \oplus b_\mathcal{S} \cdot (v_1 \oplus v_0)$, the corrupt receiver might try to unmask r by computing $u \oplus w$. However, the resulting value is indeed the protocol's output which can be simulated. Therefore, we can replace both w and u with random (the receiver sees a system of two equations that contains three unknown variables). In summary, T_2 and T_1 are indistinguishable.

A.2 Security Proof of Theorem 2

Proof. The security follows from the security of the mOT functionality and the fact the value $v_1 = H(q)^\alpha$ and $y = w^{1/\alpha}$ is distributed uniformly.

More precisely, the corrupt sender $\mathcal{S}$ learns nothing from the mOT execution as v_0 and v_1 are in the same distribution. The value v_1 reveals nothing about the receiver's input q due to the secret α under the Diffie–Hellman assumption.

The corrupt receiver obtains $w = v^k$ from the honest sender. Due to the secret PRF key k, the receiver learns nothing from v. Thus, simulation is trivial, as the parties' views in the protocol are exactly the cOPRF output.

A.3 Security Proof of Theorem 3

Proof. Let A be a coalition of corrupt parties. The view of A is a set of ciphertexts $\{C_i \mid P_i \in A\}$, and the output of the Shuffle&Decrypt which is $\{x_{\pi(1)}, \ldots, x_{\pi(m)}\}$ if the leader $P_1 \in A$.

Thanks to the property of the multi-key cryptosystem, $C_{i\in[n]}$ reveals nothing about the underlying plaintexts. If P_1 is honest, the randomization hides the party's permutation function. Moreover, when assuming $\{P_i, P_j\} \in A$ but $\{P_{i+1}, \ldots, P_{j-1}\} \notin A$, one might think that A might learn the permutation functions of honest parties $\{P_{i+1}, \ldots, P_{j-1}\}$. However, the output of the partial decryption gives ciphertexts in the random distribution. Thus, the resulting view is random to A (i.e., the corrupt coalition's view is simulated).