

# Shuffle-based Private Set Union: Faster and More Secure<sup>\*</sup>

Yanxue Jia<sup>1</sup>, Shi-Feng Sun<sup>1(✉)</sup>, Hong-Sheng Zhou<sup>2</sup>, Jiajun Du<sup>1</sup>, and Dawu Gu<sup>1(✉)</sup>

<sup>1</sup> Shanghai Jiao Tong University, China  
{jiayanxue,shifeng.sun,cqdujiajun,dwgu}@sjtu.edu.cn

<sup>2</sup> Virginia Commonwealth University, USA  
hszhou@vcu.edu

**Abstract.** Private Set Union (PSU) allows two players, the sender and the receiver, to compute the *union* of their input datasets without revealing any more information than the result. While it has found numerous applications in practice, not much research has been carried out so far, especially for large datasets.

In this work, we take shuffling technique as a key to design PSU protocols for the first time. By shuffling receiver’s set, we put forward the first protocol, denoted as  $\Pi_{\text{PSU}}^{\text{R}}$ , that eliminates the expensive operations in previous works, such as additive homomorphic encryption and repeated operations on the receiver’s set. It outperforms the state-of-the-art design by Kolesnikov et al. (ASIACRYPT 2019) in both efficiency and security; the unnecessary leakage in Kolesnikov et al.’s design, can be avoided in our design.

We further extend our investigation to the application scenarios in which both players may hold *unbalanced* input datasets. We propose our second protocol  $\Pi_{\text{PSU}}^{\text{S}}$ , by shuffling the sender’s dataset. This design can be viewed as a dual version of our first protocol, and it is suitable in the cases where the sender’s input size is much smaller than the receiver’s.

Finally, we implement our protocols  $\Pi_{\text{PSU}}^{\text{R}}$  and  $\Pi_{\text{PSU}}^{\text{S}}$  in C++ on big datasets, and perform a comprehensive evaluation in terms of both scalability and parallelizability. The results demonstrate that our design can obtain a 4-5 $\times$  improvement over the state-of-the-art by Kolesnikov et al. with a single thread in WAN/LAN settings.

## 1 Introduction

Private set operations allow mutually distrustful parties to perform set operations (like intersection and union) on their datasets, while revealing no more information about their own private input than what can be deduced from the results. Over the past decade, much progress has been made on Private Set Intersection (PSI), which has become considerably efficient and been deployed widely in practice [28,17,31,30,26,27,6,33]. In contrast, little attention has been drawn on Private Set Union (PSU).

Like the well-researched PSI, PSU also has numerous applications [18,12,3,4]. For example, it can be used for *cyber risk assessment and management*. Specifically, “*Individual blacklists today suffer from several drawbacks that limit their accuracy in malicious source identification. ... Aggregating blacklists from different maintainers and across various attack types can improve the accuracy of malicious source identification over any individual blacklist,*” as pointed out by Ramanathan et al. [36]. Therefore, it is significant for organizations (namely, maintainers of IP blacklists) to compute the joint list of individual IP blacklists, which will help minimize vulnerabilities in their infrastructure.

---

<sup>\*</sup> This is a full version of the paper accepted by USENIX Security 2022.

In addition, each individual IP blacklist is generated based on a detection strategy developed by the maintainer, which cannot be leaked; note that certain attacks could be launched by the adversaries via evading the detection strategy.

A straightforward way to obtain the joint list is to let the organizations simply exchange their blacklists. However, this will reveal the intersection of their blacklists. Then a curious organization may deduce the detection strategy of the other organization according to the IP addresses in the intersection<sup>3</sup>. Whereas, he cannot do this via the IP addresses not in his own blacklist<sup>4</sup>. Therefore, to mitigate the privacy concerns, what we essentially need is to generate the joint blacklist with the intersection hidden, which is exactly the functionality of PSU.

According to [36,1], most blacklists in real scenarios contain 1,000 - 500,000 IP addresses, and the update frequency is expected to be 5 - 15 minutes. With our new PSU protocol  $\Pi_{\text{PSU}}^{\text{R}}$  (cf. the high level ideas in Section 1.1), a joint blacklist can be obtained from individual blacklists of size  $2^{20}$  in 67.756s (cf. our experiments in Section 4) with a single thread in the WAN setting, which is efficient enough for this use case.

In addition, as mentioned in [16,18], PSU can be employed to compute the union of cancer patients of different hospitals while hiding the identities of the patients who had cancer treatment at multiple hospitals, which involves patient privacy. Also, it can be used for privacy-preserving aggregation of network traffic statistics [4], merger of two Internet providers without revealing the information of their existing networks [3], and private database supporting full join [18]. Therefore, it is highly desirable to develop efficient, scalable, and secure PSU protocols.

Many interesting variants of PSU (e.g., multi-party, set-size hidden, and shared-output PSU) have been proposed for real-world applications. In this paper, we focus on PSU in the *two-party* setting, where a sender and a receiver hold sets  $X$  and  $Y$  respectively, and aim to compute their union  $X \cup Y$  without revealing any more information than the result (especially what are the items in  $X \cap Y$ ). Particularly, the goal is to enable the receiver to learn no more information than  $X \cup Y$  and the sender to learn nothing (see Section 2.1 for the formal definition).

**PSU for balanced datasets.** The state-of-the-art work is by Kolesnikov et al. [18]; there, they proposed the first *scalable* two-party PSU protocol based on *symmetric-key techniques*. Before that, except for the circuit-based PSU [2], all designs [16,9,7] rely on *public-key operations* like additive homomorphic encryption (AHE), which make the constructed protocols *unscalable* in practice, especially for large datasets. To develop efficient PSU protocols, Kolesnikov et al. introduced the notion of Reverse Private Membership Test (RPMT) functionality as a basic building block.


**A unified view for PSU design with single sender's item.** Interestingly, when focusing on a special case of two-party PSU where the sender's data-size is 1, we observe that, the designs [9,7,18] can be presented in the same framework<sup>5</sup>. Please see Figure 1: the sender has only one item  $x$  and the receiver holds a set  $Y = \{y_1, \dots, y_n\}$ , and we denote this special case of PSU by  $(1, n)$ -PSU. We note that, protocols in this design framework, consist of two phases: **(1)** the two parties execute a RPMT functionality,

<sup>3</sup> The curious organization knows his own detection strategy that is used to identify the IP addresses in the intersection. Thus, it is reasonable for him to deduce that a similar detection strategy is used by the other organization.

<sup>4</sup> Different organizations usually monitor different areas of the Internet, as mentioned in [36]. Therefore, the curious organization knows nothing about the traffic through these IP addresses. Even if the curious organization can monitor the traffic, it is difficult to deduce the detection strategy as e.g., some sensitive strategies are based on features that can be only extracted from encrypted data.

<sup>5</sup> More precisely, [9,7,18] all realize a relaxed RPMT functionality where if  $x \notin Y$ , the information about  $x$  can be leaked to the receiver.

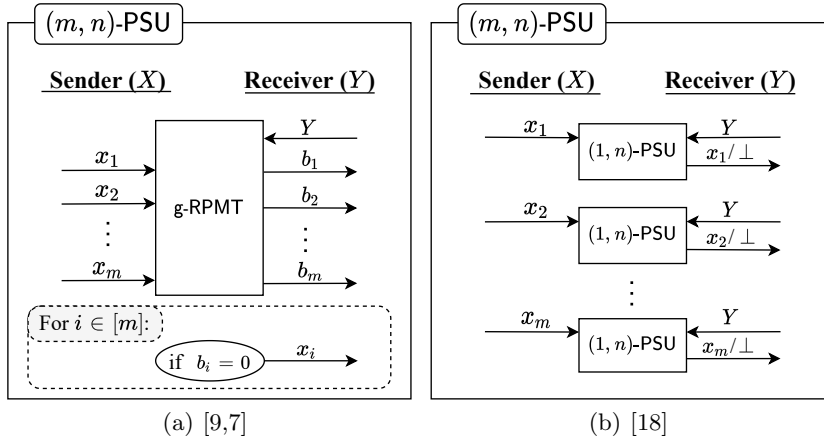
through which the receiver obtains a bit  $b$ , where  $b = 0$  means  $x \notin Y$ , otherwise  $x \in Y$ . The receiver knows no more information about  $x$  beyond whether  $x$  belongs to  $Y$ , and the sender learns nothing about  $Y$ . **(2)** if  $b = 0$  (which means  $x \notin Y$ ), the receiver obtains  $x$  (and thus  $\{x\} \cup Y$ ) while the sender learns nothing.



**Fig. 1.** Design framework for  $(1, n)$ -PSU protocols.

Under this framework, in [9,7] the RPMT functionality has been realized by employing homomorphic encryption, and in this case the receiver is able to obtain  $x$  by a straightforward decryption. As pointed out in [18], however, the public-key operations have become the workhorse of these works. Therefore, Kolesnikov et al. [18] proposed a new way of realizing RPMT in the first phase, based *only on symmetric-key operations*, and then they implemented the second phase of the design by using Oblivious Transfer (OT).

**PSU design: From single to multiple sender's items.** Now let's consider how to extend the special case  $(1, n)$ -PSU into a more general  $(m, n)$ -PSU, where  $m$  can be a very large integer. Two different approaches, as shown in Figure 2, have been proposed in [9,7] and in [18], respectively.




**Fig. 2.** Two design frameworks for  $(m, n)$ -PSU protocols.

More concretely, in [9,7], a generalized version of RPMT, denoted as **g-RPMT**, has been used, to support a large input set  $X = \{x_1, x_2, \dots, x_m\}$  from the sender. Please refer to Figure 2(a). The receiver needs only to encrypt  $Y$  once, and the ciphertext for  $x_i$  can be used to test if  $x_i \in Y$  for all  $i \in [m]$ . However, we note that the schemes in [9,7] suffer from heavy public-key operations.

In contrast, in [18] based on symmetric-key operations, for each  $x_i \in X$  from the sender, a  $(1, n)$ -PSU sub-protocol will be executed; to implement  $(m, n)$ -PSU, thus, in total,  $m$  number of  $(1, n)$ -PSU sub-protocols will be executed. Please refer to Figure 2(b).

Here for each  $(1, n)$ -PSU sub-protocol, the receiver needs to perform a polynomial interpolation with degree  $n$  for  $Y$ , which requires time  $O(n \log^2 n)$ . Thus when  $X$ 's size  $m = n$ , this approach will result in a quadratic computation complexity  $O(n^2 \log^2 n)$ .

**Bucketing techniques: The gain and the loss.** To further improve the performance, Kolesnikov et al. introduced the *bucketing technique* in PSU protocol design for the first time. Please see Figure 3: First, the sets  $X$  and  $Y$  are inserted into two simple hash tables with  $b$  bins, respectively, which means that the set  $X$  (resp.,  $Y$ ) is divided into  $b$  disjoint subsets  $X_1, \dots, X_b$  (resp.,  $Y_1, \dots, Y_b$ ). Then, each bin is padded with dummy items up to the maximum bin size  $\rho$ . The two parties perform a  $(\rho, \rho)$ -PSU sub-protocol on the items of each bin separately. In this way, the complexity of each sub-protocol on each bin is  $O(\rho^2 \log^2 \rho)$ , so the total cost can be reduced to  $O(b\rho^2 \log^2 \rho) = O(n \log n \log^2 \log n)$ .



**Fig. 3.** Bucketing technique in [18].

However, as mentioned by Kolesnikov et al. [18], the bucketing technique will incur certain information leakage about the items in the intersection, during the execution; intuitively, the receiver could learn if there are items in  $X \cap Y$  in certain subsets  $Y_i$ , where  $Y_i \subset Y$ . Note that in the ideal PSU, from the view of receiver, any item in the entire set  $Y$  could be an item in  $X \cap Y$ . But here, the receiver can know that some subsets have items in  $X \cap Y$  and others do not. To address this issue, Kolesnikov et al. proposed to add special items to each bin, with the goal of reducing the leakage. Unfortunately, we find that their way of adding only special items, is *insufficient* to resolve the problem (More details are given in Section 5).

Kolesnikov et al.'s bucketing technique [18] is significant since it improves the performance of the PSU design a lot; however, we must note that, the downside of this bucketing technique is that, the designed protocols will suffer from certain level of information leakage. This leakage may not be necessary! Recall that, in prior works [9,7] under the design framework in Figure 2(a), all items in the receiver's set  $Y$  are processed *at the same time*; although the involved public-key operations make the design much less scalable (than that in [18]) for large datasets, the resulting protocol does *not* suffer from the information leakage issue (as that in [18]).

**Main question.** Based on above discussions, we ask the following natural question:

*Is it possible to design a PSU protocol to achieve the “best of the two worlds,” i.e., (1) fast and scalable, and at the same time, (2) without any unnecessary information leakage?*

We will give an affirmative answer to this question. In particular, we propose a practical, scalable two-party PSU protocol named  $\Pi_{\text{PSU}}^{\text{R}}$  under the design framework in Figure 2(a) by *shuffling* the receiver's set. Our protocol  $\Pi_{\text{PSU}}^{\text{R}}$  relies only on lightweight symmetric-key primitives (along with some OTs; we note the OTs are also needed in the state-of-the-art result [18]). More details will be shown in Section 1.1 and Section 3.

**PSU for unbalanced datasets.** To the best of our knowledge, previous works on PSU mainly focus on designing efficient protocols for balanced datasets. We now investigate

how to design practical PSU protocols in the *unbalanced* application scenarios, in which the receiver's input size is significantly larger than the sender's, or vice versa. In fact, existing protocols (including our  $\Pi_{\text{PSU}}^{\text{R}}$  above) are already very fast when the size of receiver's input is significantly smaller than the sender's, as the relatively heavy operations mainly depend on the size of receiver's set. However, in the case that the size of receiver's input is much larger than the sender's, the performance of these protocols is reduced significantly. Hence, the second question we ask is:


*Is it possible to design a fast and scalable PSU protocol when the sender's input size is much smaller than the receiver's?*

We answer this question affirmatively by presenting a new protocol named  $\Pi_{\text{PSU}}^{\text{S}}$ ; this new protocol can be viewed as the dual version of  $\Pi_{\text{PSU}}^{\text{R}}$ , exactly by shuffling the sender's set. More details about protocol  $\Pi_{\text{PSU}}^{\text{S}}$  are given in Section 1.1 and Section 3.

### 1.1 Technical Overview

**Protocol  $\Pi_{\text{PSU}}^{\text{R}}$ : shuffling receiver's set.** We now present how to construct practical two-party PSU protocols following the design framework in Figure 2(a). A big challenge is that, we need to efficiently realize the functionality *g-RPMT*, which allows the receiver to perform membership tests while not revealing the receiver's set  $Y$  to the sender. In previous works [9,7], functionality *g-RPMT* has been realized but not efficiently due to the heavy public-key operations.

The high-level idea of our design is shown in Figure 4 and the details are as follows. Initially, each item  $y \in Y$  is split into two shares  $s$  and  $y \oplus s$  by an additive secret sharing, where  $s$  is distributed uniformly at random and perfectly hides  $y$ . The set  $Y = \{y_1, \dots, y_n\}$  can be shared into two sets  $\{s_1, \dots, s_n\}$  and  $\{y_1 \oplus s_1, \dots, y_n \oplus s_n\}$ ; the receiver will keep the former herself, and send the latter to the sender. Now we can see that for each item  $x_i \in X$ , if it belongs to  $Y$  (i.e.,  $\exists j$  s.t.  $x_i = y_j$ ), then  $x_i \oplus y_j \oplus s_j = s_j$ . Thus the sender can compute and send  $\{x_i \oplus y_1 \oplus s_1, \dots, x_i \oplus y_n \oplus s_n\}$  to the receiver, and the receiver can check if the sender's item  $x_i$  belongs to  $Y$  by computing the intersection of this set and  $\{s_1, \dots, s_n\}$ . If empty, it means  $x_i \notin Y$ . Otherwise, the receiver learns that  $x_i \in Y$ .



**Fig. 4.** Core design idea of protocol  $\Pi_{\text{PSU}}^{\text{R}}$  for  $(m, n)$ -PSU.

Now, the receiver can learn if  $x_i \in Y$  without revealing  $Y$  to the sender. However, if  $x_i \in Y$ , the receiver can *additionally* figure out which item of  $Y$  is exactly equal to

$x_i$  according to  $s_j \in \{x_i \oplus y_1 \oplus s_1, \dots, x_i \oplus y_n \oplus s_n\} \cap \{s_1, \dots, s_n\}$  as she knows the mapping of the shares  $\{s_1, \dots, s_n\}$  to the items in  $Y$ . Note that such information is not allowed to be obtained by the receiver in the RPMT. Thus, next we need to find a way to defend against such information leakage.

Recall that, the receiver is able to obtain additional information is due to the fact that the receiver knows which share corresponds to which item in  $Y$ . Our **key design idea** now is to break the mapping by *shuffling* the receiver's set (and shares) with a permutation not known to the receiver. Together with the additive secret sharing explained before, what we essentially need is a **Permute + Share** functionality: taking as input a set  $Y = \{y_1, \dots, y_n\}$  from the receiver and a permutation  $\pi$  (over  $\{1, 2, \dots, n\}$ ) from the sender, the functionality outputs the shares  $\{s_{\pi(1)}, \dots, s_{\pi(n)}\}$  and  $\{y_{\pi(1)} \oplus s_{\pi(1)}, \dots, y_{\pi(n)} \oplus s_{\pi(n)}\}$  to the receiver and the sender, respectively. After executing this functionality, the sender computes  $\{x_i \oplus y_{\pi(1)} \oplus s_{\pi(1)}, \dots, x_i \oplus y_{\pi(n)} \oplus s_{\pi(n)}\}$  and sends it to the receiver. Then the receiver can check if the sender's item belongs to  $Y$  as before. Following this way, the receiver will learn that there is an item  $s_{\pi(j)} \in \{s_{\pi(1)}, \dots, s_{\pi(n)}\}$  equal to  $x_i \oplus y_{\pi(j)} \oplus s_{\pi(j)}$  if  $x_i \in Y$ , but she is unable to find out  $y_{\pi(j)}$  according to  $s_{\pi(j)}$ , as she does not know  $\pi$ .


At this point, it seems that  $x_i$  can be completely hidden from the receiver at the first glance. Unfortunately,  $x_i \oplus y_{\pi(j)} \oplus s_{\pi(j)}$  may still leak partial information about  $x_i$  to the receiver. This is because  $y_{\pi(j)} \oplus s_{\pi(j)}$  is not distributed uniformly and independently from the perspective of the receiver who knows  $s_{\pi(j)}$  and  $y_{\pi(j)}$ . To overcome this obstacle, we further employ multi-point Oblivious Pseudorandom Function (OPRF)<sup>6</sup>  $F(k, \cdot)$  to conceal  $\{x \oplus y_{\pi(1)} \oplus s_{\pi(1)}, \dots, x \oplus y_{\pi(n)} \oplus s_{\pi(n)}\}$ . More concretely, the receiver takes  $\{s_{\pi(1)}, \dots, s_{\pi(n)}\}$  as the input to multi-point OPRF, then the sender receives the PRF key  $k$  and the receiver obtains  $\{F(k, s_{\pi(1)}), \dots, F(k, s_{\pi(n)})\}$ . In this case, the sender with the key  $k$  can compute  $I_i = \{F(k, x_i \oplus y_{\pi(1)} \oplus s_{\pi(1)}), \dots, F(k, x_i \oplus y_{\pi(n)} \oplus s_{\pi(n)})\}$  for each  $x_i$  and sends it to the receiver. Then the receiver proceeds to perform the membership test as before, but learns nothing about  $x_i$  as she does not know the PRF key  $k$ . For the second phase, the receiver can receive  $x_i \notin Y$  through OT as in [18].

*Optimization.* Following our core idea, it can be seen that the protocol executes **Permute+Share** and multi-point OPRF only once for all  $x_i \in X$ , but needs to send  $m$  sets (i.e.,  $I_1, \dots, I_m$ ) to receiver. In addition, the sender and receiver need to execute OT sub-protocol  $m$  times. Fortunately, we find that the functionality **Permute + Share** (resp. multi-point OPRF) can be securely realized by the protocols in [5,21] (resp. [6]) with computation and communication cost  $O(n \log n)$  (resp.  $O(n)$ ). However, the sender needs to compute and send  $I_i$  containing  $n$  PRF values for each  $x_i \in X$ , thus it results in a quadratic computation and communication complexity  $O(mn)$ .

The main reason for this quadratic complexity is that the sender does not know which item in  $Y$  may be equal to his item  $x_i$ , so he has to XOR  $x_i$  with all shares  $\{y_{\pi(1)} \oplus s_{\pi(1)}, \dots, y_{\pi(n)} \oplus s_{\pi(n)}\}$ . To improve the efficiency, our key idea is to reduce the number of items in  $Y$  that could be equal to  $x_i$  by leveraging Cuckoo hashing [25] (defined in Section 2.2). Briefly, we insert  $Y$  into a Cuckoo hash table with  $\gamma$  hash functions,  $\epsilon \cdot n$  bins and stash size  $0^7$ , and then execute  $\Pi_{\text{PSU}}^{\text{R}}$  over the hash table. To make it clear, we take a concrete example to explain the optimization via Cuckoo hashing, as illustrated in Figure 5. In particular, we assume that the sender's item to be checked is  $x_i$ , and that the receiver's set  $Y = \{y_1, \dots, y_6\}$ , can be inserted to the Cuckoo hash table with

<sup>6</sup> Multi-point OPRF is evaluated on different inputs with the same key, while single-point OPRF is evaluated with a different key for each input.

<sup>7</sup> According to the empirical analysis in [33], the stash size can be reduced to 0 by increasing the number of hash functions while achieving a hashing failure probability of  $2^{-40}$ .



**Fig. 5.**  $\Pi_{\text{PSU}}^{\text{R}}$ : Optimization via Cuckoo hashing.

8 bins and 2 hash functions  $h_1(\cdot)$  and  $h_2(\cdot)$ <sup>8</sup>. Then the optimized protocol works as follows: **(1)** The receiver inserts  $Y$  into the Cuckoo hash table and adds a dummy item  $d$  to each empty bin, then obtains the filled table denoted by  $Y^*$ . **(2)** The receiver and the sender execute **Permute + Share** with a randomly chosen permutation  $\pi$  and  $Y^*$  as inputs, and obtain the shuffled secret share sets  $Sh_1$  and  $Sh_2$ , respectively. Here,  $s_d^1$  and  $s_d^2$  in  $Sh_1$  are the shares of dummy item and  $s_d^1 \neq s_d^2$ . The dotted arrows mean that after permutation  $\pi$ , the 4-th (resp. 7-th) item in  $Sh_1$  is the share of the 2-th (resp. 6-th) item in  $Y^*$ , but the receiver does not know these corresponding relations. In addition, the other share is the 4-th (resp. 7-th) item in  $Sh_2$ . **(3)** With  $h_1(\cdot)$  and  $h_2(\cdot)$ , the sender computes  $h_1(x_i)$  and  $h_2(x_i)$ , say  $h_1(x_i) = 2$  and  $h_2(x_i) = 6$ ; we note that according to the principle of filling the Cuckoo hash table, the potential item of  $Y$  that is equal to  $x_i$  must be inserted to the position  $h_1(x_i)$  or  $h_2(x_i)$  of  $Y^*$ . Then he uses  $\pi$  to locate the corresponding shares (namely, 4-th and 7-th items) in  $Sh_2$  and generates  $I_i$  with them. **(4)** The sender sends  $I_i$  to the receiver and the receiver proceeds as before.

Based on the above optimization, we can reduce the number of items in  $I_i$  to a constant  $\gamma$  that is the number of hash functions. Hence, the computation and communication cost incurred by  $\{I_1, \dots, I_m\}$  can be reduced to  $O(\gamma m)$ . In Table 1, we summarize the computation and communication costs of main steps in our  $\Pi_{\text{PSU}}^{\text{R}}$ . And we will give a more detailed complexity analysis by taking account of the error rate and security parameters in Section 3.3.


**Table 1.** The computation and communication costs of  $\Pi_{\text{PSU}}^{\text{R}}$ .

	Permute + Share	multi-point OPRF $\{I_i\}_{i \in [m]}$	OT
Costs	$O(n \log n)$	$O(n)$	$O(m)$

$m$  is the sender's set size;  $n$  is the receiver's set size.

**Protocol  $\Pi_{\text{PSU}}^{\text{S}}$ : shuffling sender's set.** From Table 1 we can see that when  $m \gg n$  the overall cost of  $\Pi_{\text{PSU}}^{\text{R}}$  is dominated by  $\{I_i\}_{i \in [m]}$  and OT that are linear in  $m$ . However, when  $m \ll n$  the cost is dominated by **Permute + Share** that is superlinear in  $n$ . Therefore, when considering unbalanced datasets,  $\Pi_{\text{PSU}}^{\text{R}}$  is more suitable for the case that the sender's set size is much larger than the receiver's (i.e.,  $m \gg n$ ). To develop efficient solutions for the other case where  $m \ll n$ , we propose a second protocol  $\Pi_{\text{PSU}}^{\text{S}}$  by *shuffling the sender's set*. As a whole, it can be regarded as the dual version of  $\Pi_{\text{PSU}}^{\text{R}}$ . The high-level idea is shown in Figure 6.

<sup>8</sup> Note that the parameters used here are to simplify the explanation, please refer to Section 4.1 for the concrete parameter choices.



**Fig. 6.** Core idea of  $\Pi_{\text{PSU}}^{\text{S}}$  for  $(m, n)$ -PSU.

The basic idea of  $\Pi_{\text{PSU}}^{\text{S}}$  is to share the sender's set  $X$  into two share sets obtained by the sender and the receiver, respectively. Then, the sender sends the shares of the items in  $X \setminus Y$  to the receiver such that the receiver can recover the items in  $X \setminus Y$ . While being shared, the sender's set  $X$  needs to be shuffled by a permutation not known by the sender such that the sender cannot know the correspondence between shares and the items in  $X$ . More concretely, through **Permute + Share** with set  $X$  and permutation  $\pi$  as inputs, the sender obtains the set of shuffled shares  $\{s_{\pi(1)}, \dots, s_{\pi(m)}\}$  and the receiver obtains the remaining ones  $\{x_{\pi(1)} \oplus s_{\pi(1)}, \dots, x_{\pi(m)} \oplus s_{\pi(m)}\}$ . Note that for an item  $x_{\pi(i)} \in X$ , if  $x_{\pi(i)} \in Y$ , say  $x_{\pi(i)} = y_j$ , then  $y_j \oplus x_{\pi(i)} \oplus s_{\pi(i)} = s_{\pi(i)}$ , otherwise  $y_j \oplus x_{\pi(i)} \oplus s_{\pi(i)} \neq s_{\pi(i)}$  for any  $y_j \in Y$ . Moreover, it is obvious that the receiver can recover  $x_{\pi(i)}$  by computing  $s_{\pi(i)} \oplus (x_{\pi(i)} \oplus s_{\pi(i)})$  provided that she can also obtain the corresponding share  $s_{\pi(i)}$ . Thus, for each share  $x_{\pi(i)} \oplus s_{\pi(i)}$  the receiver can generate  $I_i = \{y_1 \oplus x_{\pi(i)} \oplus s_{\pi(i)}, \dots, y_n \oplus x_{\pi(i)} \oplus s_{\pi(i)}\}$  and send it to the sender, so that the sender is able to determine whether to send  $s_{\pi(i)}$  by checking if  $s_{\pi(i)} \in I_i$ ; if  $s_{\pi(i)} \notin I_i$  meaning  $x_{\pi(i)} \notin Y$ , the sender sends  $s_{\pi(i)}$  to the receiver, otherwise does nothing. As in  $\Pi_{\text{PSU}}^{\text{R}}$ , however,  $x_{\pi(i)} \oplus s_{\pi(i)}$  is not enough to completely hide  $Y$  as the sender knows both  $x_{\pi(i)}$  and  $s_{\pi(i)}$ . To avoid this leakage, we adopt the same approach as  $\Pi_{\text{PSU}}^{\text{R}}$ . That is, we compute  $I_i = \{F(k, y_1 \oplus x_{\pi(i)} \oplus s_{\pi(i)}), \dots, F(k, y_n \oplus x_{\pi(i)} \oplus s_{\pi(i)})\}$  through an mpOPRF  $F$  with the key not known to the sender.

We remark that the sender does not know the correspondence between shares and his original items, so he cannot learn which item is in  $X \cap Y$ . Nevertheless, the sender can know  $|X \cap Y|$ , which is not revealed to the sender in  $\Pi_{\text{PSU}}^{\text{R}}$ .

*Optimization.* Similar to  $\Pi_{\text{PSU}}^{\text{R}}$  without optimization,  $\{I_1, \dots, I_m\}$  here also results in a high computation and communication complexity  $O(mn)$ . To further reduce the cost, we adopt the similar idea as  $\Pi_{\text{PSU}}^{\text{R}}$  to reduce the size of  $I_i$ . Besides the Cuckoo hashing, here we additionally need a simple hash table to record  $Y$ , which is crucial for locating the candidate items of  $Y$  equal to the item of  $X$ . For ease of exposition, we take a concrete example to explain the main idea, also illustrated in Figure 7. Particularly, we assume that  $s_4$  is the share to be checked by using  $I_1$ . Then the optimized protocol works as follows: (1) The sender and receiver inserts their sets  $X$  and  $Y$  into a Cuckoo hash table and a simple hash table respectively, using hash functions  $h_1(\cdot)$  and  $h_2(\cdot)$ . We denote the Cuckoo hash table containing  $X$  as  $X^*$ . (2) After **Permute + Share** on  $X^*$ , the sender and receiver can obtain share sets  $Sh_1$  and  $Sh_2$ , respectively. (3) The counterpart of  $s_4$  is the first item in  $Sh_2$ , say  $x_4 \oplus s_4$ . Since the receiver knows the permutation  $\pi$ , she can learn that  $x_4 \oplus s_4$  is the share of 3-th item (namely,  $x_4$ ) in  $X^*$ . According to the principle of filling the Cuckoo hash table and simple hash table, the receiver knows that only the items ( $y_9$  and  $y_{16}$ ) mapped to the 3-th position of the simple hash table may



equal the 3-th item in  $X^*$ . Then she computes  $F(k, y_9 \oplus x_4 \oplus s_4)$  and  $F(k, y_{16} \oplus x_4 \oplus s_4)$  and adds them into  $I_1$ . To hide the number of the receiver's items inserted into the bin, the receiver further pads  $I_1$  up to the maximum size of bin with random elements  $r$  chosen from the range of  $F(k, \cdot)$ . (4) Once receiving  $I_1$ , the sender determines whether to send  $s_4$  as before.

Unfortunately, we can see from above that if  $x_4 \notin Y$  and the sender sends  $s_4$  immediately to the receiver, then the receiver can recover  $x_4$  and learn its position in the Cuckoo hash table. Since an item's position in the Cuckoo hash table is also impacted by the other items, the position will leak information about the entire set  $X$ . In addition, if  $x_4 \in Y$  and the sender does not send  $s_4$  to the receiver, then the receiver knows that an item belonging to  $X \cap Y$  is in  $\{y_9, y_{16}\}$ , which is not allowed in PSU. To overcome the above problems, our basic idea is to let the sender record the shares that should be sent first and then send them together in a random order. We implement this by using Permute + Share again. For more details, please refer to Section 3.2.




Fig. 7.  $\Pi_{\text{PSU}}^{\text{S}}$ : Optimization via Cuckoo hashing.

In general, when  $n$  items are inserted into a simple hash table with  $m$  bins using  $\gamma$  hash functions, the maximum bin size is  $O(\gamma n/m)$ <sup>9</sup> when  $n > m \log m$  according to [34]. Therefore, the computation and communication costs of  $\{I_1, \dots, I_m\}$  can be reduced to  $m \cdot O(\gamma n/m) = O(\gamma n)$ .

## 1.2 Our Contributions

We explore new techniques of designing two-party PSU protocols for both balanced and unbalanced datasets, and propose two efficient and secure PSU protocols  $\Pi_{\text{PSU}}^{\text{R}}$  and  $\Pi_{\text{PSU}}^{\text{S}}$  supporting big datasets in Section 3. More specifically, our main contributions are summarized as below.

**New protocols.** We for the first time, give a scalable and secure construction, named  $\Pi_{\text{PSU}}^{\text{R}}$ , for realizing two-party PSU. Note that the state-of-the-art design by Kolesnikov et al. [18] faces the issue of partial information leakage of items in intersection. While this protocol is efficient for balanced datasets, we further extend our study to the unbalanced case in the sense that the receiver's input size is significantly larger than the sender's, or vice versa. Then we propose a second efficient and secure protocol, dubbed  $\Pi_{\text{PSU}}^{\text{S}}$ . This protocol is suitable in the applications where the sender's input size is much smaller than the receiver's; this can be viewed as a dual version of our first protocol which is more suitable for the opposing case.

<sup>9</sup> To be precise, the maximum bin size should be  $\Theta(\gamma n/m)$ .

**New design techniques.** To avoid the leakage incurred by the leverage of bucketing technique on the receiver’s set, our key point is to process the receiver’s set *at the same time*. Then we design  $\Pi_{\text{PSU}}^{\text{R}}$  under the framework in Figure 2(a) by shuffling the receiver’s set. Regarding designing PSU protocols for unbalanced datasets, our observation is to perform heavy operations on the smaller dataset. Thus we design  $\Pi_{\text{PSU}}^{\text{S}}$  by shuffling the sender’s set for the case that the sender’s set size is much smaller than the receiver’s. With the key technique *shuffling*, our design avoids expensive computations like public-key operations and repeated operations on sender/receiver’s set. Furthermore, we reduce the communication and computation overhead by employing the Cuckoo hashing, which is for the first time used in PSU.

**Implementation & evaluation.** We implement our protocols in C++ and perform a comprehensive evaluation in Section 4. The results demonstrate that  $\Pi_{\text{PSU}}^{\text{R}}$  is  $4\text{-}5\times$  faster than the state-of-the-art PSU protocol [18] with a single thread in WAN/LAN settings. Moreover, we show that our protocols support parallelization;  $\Pi_{\text{PSU}}^{\text{R}}$  and  $\Pi_{\text{PSU}}^{\text{S}}$  can achieve a speedup of  $2.89\times$  and  $3.49\times$  respectively at 8 threads in LAN setting. Beyond, the results indicate that our protocols are also efficient and scalable for unbalanced datasets.

**New leakage analysis.** In Section 5 we show the bucketing technique adopted in the state-of-the-art design [18] will leak the information of intersection. Specifically, we demonstrate that after knowing all the sender’s items in a bin belong to the intersection, the receiver can learn that her corresponding bin has items in  $X \cap Y$  with an extremely large probability.

### 1.3 Related Work

As a special case of secure two-party computation, privacy-preserving set operation also includes generic and custom constructions. The generic PSI protocol (also called circuit-based PSI) was firstly proposed by Huang et al. [13], and the first generic PSU protocol was proposed in [2]. In general, the generic protocols are less efficient than the custom ones but more flexible to support different functionalities. In this work, we are mainly interested in the custom constructions.

Over the past decades, a large amount of work has been done on specific PSI (e.g., [20,28,17,31,30,26,27,6,33]). The recent works are mainly based on oblivious transfer extension [28,32,38,24] and various OPRF constructions [17,26,27,6,39]. The state-of-the-art protocols have become considerably efficient for practical applications.

Although PSI and PSU are similar and they share some building blocks (e.g., OPRF [17,26,18]), PSU cannot be obtained by directly employing existing PSI techniques, and little process has been made towards practical PSU so far.

The first PSU protocol was proposed by Kissner and Song in [16], and realized by using threshold additively homomorphic encryption (AHE) and polynomial representation. Later, Frikken [9] proposed a new PSU protocol with intersection hidden by leveraging similar techniques. Roughly, their protocol works as follows: the receiver holding the secret key of AHE sends to the sender the encrypted polynomial representation  $\text{Enc}(f)$  of her own set  $Y$ , then the sender with set  $X$  calculates the tuples  $(\text{Enc}(xf(x)), \text{Enc}(f(x)))$  for each  $x \in X$ , and sends them to the receiver. If  $x \in Y$ , the receiver can only recover  $(0, 0)$  from the tuple without learning any information about  $x$ . Otherwise, the receiver can recover  $(xf(x), f(x))$  and then obtain  $x$ . Following the similar idea in [9], Davidson and Cid [7] presented a new protocol by replacing polynomial representation with inverted Bloom Filter<sup>10</sup>.

<sup>10</sup> If there is an item mapped to an entry of Bloom Filter, the entry will be filled with a bit 0, otherwise, with a bit 1.

All the above protocols encrypt the (polynomial or Bloom Filter) representation of the receiver’s set using AHE and perform a large number of operations in an encrypted manner. As pointed out by Kolesnikov et al. in [18], the public-key operations have become the workhorse of these works. Then they proposed the first scalable PSU protocol using only symmetric-key techniques. In their work, a polynomial is also used to represent the receiver’s set, but the receiver is required to *re-generate* her polynomial representation for testing each item of the sender. By this way, the design in [18] avoids the usage of the expensive additive homomorphic encryption, but still suffers from the repeated high-degree polynomial interpolation. To further reduce this cost, Kolesnikov et al. proposed an efficient optimization by using the bucketing technique.

Next, we summarize the asymptotic complexities of the above PSU protocols [16,9,7,18] and ours in Table 2. In terms of asymptotic complexity, the scheme in [7] is the most efficient. However, according to the experimental comparison shown in [18], the protocol in [7] is 7607× slower than [18] due to heavy public-key operations. Note that our protocol  $\Pi_{\text{PSU}}^{\text{R}}$  is 4-5× faster than [18]. On the other hand, the PSU protocols in [18] and our work are only based on symmetric-key operations, but the complexities are super-linear. Therefore, designing a PSU protocol with linear complexity by using symmetric-key operations is still left open.

Protocol	Comm. (bits)	Comp. (#Ops)	
		pub-key	symm-key
[16]	$O(n^2)$	$O(n^2)$	-
[9]	$O(n)$	$O(n \log \log(n))$	-
[7]	$O(n)$	$O(n)$	-
[18]	$O(n \log(n))$	-	$O(n \log(n))$
$\Pi_{\text{PSU}}^{\text{R/S}}$	$O(n \log(n))$	-	$O(n \log(n))$

**Table 2.** Comparisons of asymptotic communication (bits) and computation (#operations) costs of two-party PSU protocols in the semi-honest setting. pub/symm-key: public/symmetric-key operations. Here,  $n$  is the size of the parties’ input sets. For [18] and ours, we ignore the pub-key cost of  $\kappa$  base OTs where  $\kappa$  is computational security parameter.

We now provide a concrete comparison for the state-of-the-art results on PSU and on PSI for the large datasets of size  $2^{20}$  in LAN setting. We can see that the performance of PSI is far better than that of PSU. Our  $\Pi_{\text{PSU}}^{\text{R}}$  outperforms Kolesnikov et al. [18] by a factor of 5.4, but is still 20× lower than the PSI protocol in [17]. More research on PSU design should be encouraged to further improve the performance.

	PSU		PSI	
	[18]	$\Pi_{\text{PSU}}^{\text{R}}$	[17]	[39]
Time (s)	263.476	48.703	2.441	5.396
Comm.(MB)	2470.11	1338.79	128.5	53.55

**Table 3.** Comparisons of total runtime (in seconds) and communication (in MB) between the state-of-the-art works on PSU and PSI for set size  $2^{20}$  in LAN setting.

Finally, we remark that in a concurrent and independent work, Garimella et al. [10] improve the results in [18] by a factor of 2 - 2.5× also by leveraging shuffling technique. Similar to that in [18], polynomial representation is also used for receiver’s set in [10]; the difference is that, repeated polynomial interpolations can be avoided in [10], leading to better performance than that in [18]. However, we note that in our work the polynomial interpolations can be *entirely* avoided; thus we can obtain even better performance than that in [10].

## 2 Preliminaries

**Notation.** We denote by  $\kappa$  and  $\lambda$  the computational and statistical security parameters, respectively. We use  $[m]$  to denote the set  $\{1, 2, \dots, m\}$ , and  $X = \{x_1, \dots, x_n\}$  to denote a set with size  $|X| = n$ . Given a permutation  $\pi$  on  $n$  items, we use  $\pi(X)$  to denote the set  $\{x_{\pi(1)}, x_{\pi(2)}, \dots, x_{\pi(n)}\}$ .

### 2.1 Security Model

Our PSU protocol involves two parties, and we follow the static semi-honest security definition in [19] for secure two-party computation in this work.


**Static Semi-Honest Security.** There are two parties denoted by  $P_0$  and  $P_1$ . Let  $f_i(X, Y)$  be the output for  $P_i$  in the ideal functionality  $\mathcal{F}$  and  $f(X, Y) = (f_0(X, Y), f_1(X, Y))$  be the joint output. Let the view of  $P_i$  during an execution of  $\Pi$  on inputs  $(X, Y)$  be  $\text{view}_i^\Pi(X, Y)$  that consists of the input  $X$  or  $Y$ , the contents of  $P_i$ 's internal random tape and the messages received during the execution. Similarly,  $\text{output}_i^\Pi(X, Y)$  is the output of  $P_i$  during an execution of  $\Pi$  on inputs  $(X, Y)$  and can be computed from the  $\Pi$ 's view. And the joint output of both parties is  $\text{output}^\Pi(X, Y) = (\text{output}_0^\Pi(X, Y), \text{output}_1^\Pi(X, Y))$ .

**Definition 1.** A protocol  $\Pi$  securely computes  $\mathcal{F}$  against static semi-honest adversaries if there exist probabilistic polynomial-time (PPT) algorithms  $\text{Sim}_0$  and  $\text{Sim}_1$  such that

$$\begin{aligned} (\text{Sim}_0(X, f_0(X, Y)), f(X, Y)) &\stackrel{c}{=} (\text{view}_0^\Pi(X, Y), \text{output}^\Pi(X, Y)), \\ (\text{Sim}_1(Y, f_1(X, Y)), f(X, Y)) &\stackrel{c}{=} (\text{view}_1^\Pi(X, Y), \text{output}^\Pi(X, Y)). \end{aligned}$$

**PSU Functionality.** The ideal functionality for PSU, denoted as  $\mathcal{F}_{\text{PSU}}^{n_1, n_2}$ , is shown in Figure 8 (except for the text marked in blue). This functionality allows two players, the sender and the receiver, who hold private datasets with size  $n_1$  and  $n_2$ , respectively, to compute the union of the both input datasets. Note that our formulation of PSU functionality is identical to that in [18], and we allow only the receiver, not the sender, to obtain the union of the two input sets.

We remark that, in this formulation, based on the obtained output, the receiver can easily calculate the size of the intersection of the two input sets. However, the receiver is not allowed to learn any additional information about the data items in the intersection. On the other hand, the sender is not allowed to learn any information about the union or the intersection of the two private input sets.




**Fig. 8.** Ideal Functionalities for PSU (The difference is marked in blue).

With the goal of investigating PSU design comprehensive, we further consider a natural relaxation of the ideal PSU functionality, by allowing the sender to learn the size of the intersection; as mentioned above, the receiver by default, is allowed to obtain the intersection size. We denote the relaxed ideal functionality as  $\mathcal{F}_{\text{PSU}^*}^{n_1, n_2}$ ; we also show it in Figure 8 while the difference from  $\mathcal{F}_{\text{PSU}}^{n_1, n_2}$  is marked in blue.

## 2.2 Building Blocks

We briefly recollect the main cryptographic tools, including Permute + Share, multi-point Oblivious PRF, 1-out-of-2 Oblivious Transfer, simple hashing and Cuckoo hashing.

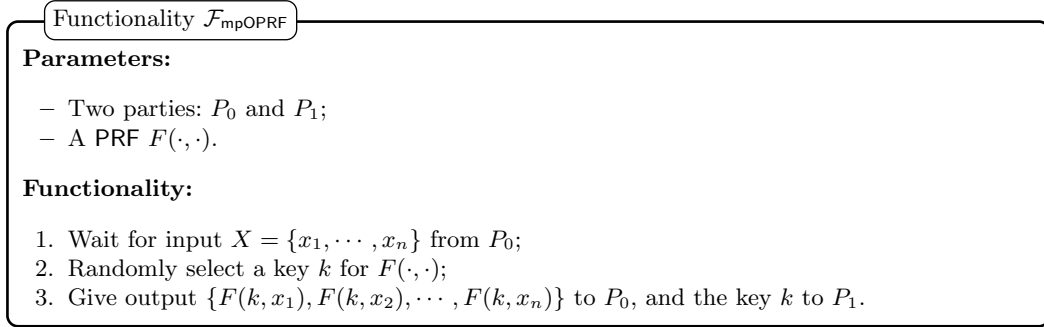
**Permute + Share.** The Permute + Share functionality  $\mathcal{F}_{\text{PS}}$  shown in Figure 9 is defined by Chase et al. in [5]. There are two parties  $P_0$  and  $P_1$  in this functionality, where  $P_0$  possesses a set  $X = \{x_1, \dots, x_n\}$  of size  $n$  and  $P_1$  picks a permutation  $\pi$  on  $n$  elements. The goal of  $\mathcal{F}_{\text{PS}}$  is to let  $P_0$  learn the shares  $\{s_{\pi(1)}, s_{\pi(2)}, \dots, s_{\pi(n)}\}$  and  $P_1$  learn nothing but the other shares  $\{x_{\pi(1)} \oplus s_{\pi(1)}, x_{\pi(2)} \oplus s_{\pi(2)}, \dots, x_{\pi(n)} \oplus s_{\pi(n)}\}$ . As mentioned in [5], some earlier works [13,21] can also be used for securely realizing  $\mathcal{F}_{\text{PS}}$ . These solutions all have computation/communication complexity  $O(n \log n)$ .




**Fig. 9.** Permute + Share functionality.

**Multi-Point Oblivious PRF.** Oblivious PRF (OPRF) is a protocol involving two parties  $P_0$  and  $P_1$ , where  $P_1$  obtains the key of the PRF  $F(\cdot, \cdot)$  and  $P_0$  takes as input  $x$  and obtains  $F(k, x)$ . OPRF has been widely used in PSI protocols, and extensive efforts have been made to develop efficient *single-point* OPRF protocols [8,11,29,17]. Most recently, Pinkas et al. [26] proposed for the first time to use *multi-point* OPRF to realize more efficient PSI protocols. Particularly, in a multi-point OPRF,  $P_0$  takes as input  $\{x_1, \dots, x_n\}_{n \geq 1}$  and obtains  $\{F(k, x_1), \dots, F(k, x_n)\}_{n \geq 1}$  while  $P_1$  obtains the PRF key  $k$ . Later, Chase et al. [6] designed a more efficient multi-point OPRF with computation complexity  $O(n)$  while the computation cost of [26] is  $O(n \log^2 n)$ . Moreover, the scheme in [6] only involves efficient OT extension and AES operations, rather than the high-degree polynomial interpolation/evaluation over a large field as in [26]. The functionality  $\mathcal{F}_{\text{mpOPRF}}$  is shown in Figure 10.

**1-out-of-2 Oblivious Transfer.** 1-out-of-2 oblivious transfer (OT) is a two-party protocol, where party  $P_0$  takes as input two strings  $\{x_0, x_1\}$ , and the other party  $P_1$  chooses a random bit  $b$  and obtains nothing other than  $x_b$  while  $P_0$  learns nothing about  $b$ . The first OT protocol was proposed by Rabin in [35]. And due to the lower bound in [14], all

**Fig. 10.** Multi-Point OPRF functionality.

the OT protocols require expensive public-key operations. To improve the performance, Ishai et al. [15] introduced the concept of OT extension that enables us to carry out many OTs based on a small number of basic OTs. The functionality  $\mathcal{F}_{\text{OT}}$  is shown in Figure 11.

**Fig. 11.** 1-out-of-2 Oblivious Transfer functionality.

**Simple Hashing.** In the simple hashing scheme, there are  $\gamma$  hash functions  $h_1, \dots, h_\gamma : \{0, 1\}^* \rightarrow [b]$  used to map  $n$  items into  $b$  bins  $B_1, \dots, B_b$ . An item  $x$  will be added into  $B_{h_1(x)}, B_{h_2(x)}, \dots, B_{h_\gamma(x)}$ , regardless of whether these bins are empty. According to the following inequality [22], the maximum bin size  $\rho$  can be set to ensure that no bin will contain more than  $\rho$  items except with probability  $2^{-\lambda}$  when hashing  $n$  items into  $b$  bins.

$$\Pr[\exists \text{ bin with } \geq \rho \text{ items}] \leq b \left[ \sum_{i=\rho}^n \binom{n}{i} \cdot \left(\frac{1}{b}\right)^i \cdot \left(1 - \frac{1}{b}\right)^{n-i} \right]$$

**Cuckoo Hashing.** Cuckoo hashing was introduced by Pagh and Rodler in [25]. In this hashing scheme, there are  $\gamma$  hash functions  $h_1, \dots, h_\gamma$  used to map  $n$  items into  $b = \epsilon n$  bins and a stash, and we denote the  $i$ -th bin as  $B_i$ . Unlike the simple hashing, the Cuckoo hashing can guarantee that there is only one item in each bin, and the approach to avoid collisions is as follows: For an item  $x$ , it can be inserted into any empty bin of  $B_{h_1(x)}, B_{h_2(x)}, \dots, B_{h_\gamma(x)}$ . If there are no empty bins in the  $k$  bins, randomly select a bin  $B_{h_r(x)}$  in these  $\gamma$  bins, and evict the prior item  $y$  in  $B_{h_r(x)}$  where  $h_r(x) = h_r(y)$  to a new bin  $B_{h_i(y)}$  where  $i \neq r$ . The above procedure is repeated until no more evictions are necessary, or until the number of evictions has reached a threshold. In the latter case, the last item will be put in the stash. According to the empirical analysis in [33], we

can adjust the values of  $\gamma$  and  $\epsilon$  to reduce the stash size to 0 while achieving a hashing failure probability of  $2^{-40}$ .

### 3 Private Set Union via Shuffling

In this section, we propose two scalable PSU protocols  $\Pi_{\text{PSU}}^{\text{R}}$  and  $\Pi_{\text{PSU}}^{\text{S}}$  by leveraging shuffling and Cuckoo hashing techniques. The first protocol  $\Pi_{\text{PSU}}^{\text{R}}$  realizes  $\mathcal{F}_{\text{PSU}}^{n_1, n_2}$ , which is obtained by *shuffling the receiver's set*. In contrast, the second  $\Pi_{\text{PSU}}^{\text{S}}$  realizes  $\mathcal{F}_{\text{PSU}^*}^{n_1, n_2}$ , which is obtained by *shuffling the sender's set*. To ease the understanding of our main idea, we also present the simplified versions of  $\Pi_{\text{PSU}}^{\text{R}}$  and  $\Pi_{\text{PSU}}^{\text{S}}$  (without using the Cuckoo hashing) denoted by  $\hat{\Pi}_{\text{PSU}}^{\text{R}}$  and  $\hat{\Pi}_{\text{PSU}}^{\text{S}}$ , respectively, in Appendix A.

#### 3.1 Protocol $\Pi_{\text{PSU}}^{\text{R}}$ : Shuffling Receiver's Set

The first protocol  $\Pi_{\text{PSU}}^{\text{R}}$  is designed under the framework in Figure 2(a). Our basic idea is to realize the functionality  $g$ -RPMT by shuffling the secret shares of the receiver's set. In the following, we first give a brief description of this protocol and then present the details in Figure 12.

We assume that the sender's set is  $X = \{x_1, \dots, x_{n_1}\}$  and the receiver's set is  $Y = \{y_1, \dots, y_{n_2}\}$ . Then the protocol proceeds as follows. Firstly, the receiver chooses the parameters of Cuckoo hash table without a stash, including the number of bins  $b = \epsilon \cdot n_2$  and  $\gamma$  hash functions  $h_1, \dots, h_\gamma$ . Then she inserts  $Y$  into this table and pads each empty bin with a dummy item  $d$ . Please refer to Section 2.2 for the details of Cuckoo hashing. After successfully inserting  $Y$  into the Cuckoo hash table, the receiver sends the parameters to the sender. Hereafter, we denote by  $Y_C$  the Cuckoo hash table filled with  $Y$  and  $Y_C[i]$  the item in the  $i$ -th bin of the table.

Secondly, the two parties invoke  $\mathcal{F}_{\text{PS}}$  with inputs  $Y_C$  and  $\pi$  randomly chosen by the receiver. After this, the sender and receiver obtain the shares  $\{a'_1, a'_2, \dots, a'_b\}$  and  $\{a_1, a_2, \dots, a_b\}$  respectively, where  $a'_i \oplus a_i = Y_C[\pi(i)]$ . Further through  $\mathcal{F}_{\text{mpOPRF}}$  with  $\{a_i\}_{i \in [b]}$  as the input, the sender receives PRF key  $k$  and the receiver obtains  $\{F(k, a_i)\}_{i \in [b]}$  where  $F(k, a_i) \in \{0, 1\}^{\ell_2}$ .

Next, for each  $x_i \in X$ , the sender generates a set  $I_i$  so that the receiver can test if  $x_i \in Y$  via  $I_i$ . If not, the receiver can obtain the item  $x_i$ . However, we observe that if the sender picks items from  $X$  in a special order, then when the receiver obtains a certain item  $x_i$  she can obtain extra information about  $X$  according to the order<sup>11</sup>. To avoid this leakage, the sender permutes his set  $X$  to  $\pi'(X) = \{x'_1, \dots, x'_{n_1}\}$  by a randomly chosen permutation  $\pi'$ , and then generates  $I_i$  for each item  $x'_i$  of  $\pi'(X)$  in turn, the details of which are shown below.

Note that for each item  $x'_i \in \pi'(X)$ , if there is an item  $y \in Y$  equal to  $x'_i$ , then  $y$  must be inserted into one of the positions of  $Y_C$  indexed by  $\{h_j(x'_i)\}_{j \in [\gamma]}$ , according to the property of the Cuckoo hashing. Hence, to test if  $x'_i \in Y$ , we need only to check if  $x'_i \in \{Y_C[h_1(x'_i)], \dots, Y_C[h_\gamma(x'_i)]\}$ . To do so, the sender first uses the permutation  $\pi$  of  $\mathcal{F}_{\text{PS}}$  to identify the shares of  $Y_C[h_1(x'_i)], \dots, Y_C[h_\gamma(x'_i)]$  from  $\{a'_1, a'_2, \dots, a'_b\}$ , say  $\{a'_{q_1}, a'_{q_2}, \dots, a'_{q_\gamma}\}$ , where  $q_j = \pi^{-1}(h_j(x'_i))$  for  $j \in [\gamma]$ . Then he computes  $I_i$  as  $I_i = \{F(k, x'_i \oplus a'_{q_1}), \dots, F(k, x'_i \oplus a'_{q_\gamma})\}$ . However, we notice that if there are distinct hash functions, say  $h_{j_s}$  and  $h_{j_t}$  s.t.  $h_{j_s}(x'_i) = h_{j_t}(x'_i)$ , then we have  $I_i[j_s] = I_i[j_t]$ , from which the receiver may learn partial information about  $x'_i$ . To overcome this shortcoming,  $I_i$  is generated in a slightly different way:  $I_i[j_s]$  is computed and recorded in  $I_i$  as before,

<sup>11</sup> For example, assuming that  $X$  consists of the ages of a group people, if the sender picks the items of  $X$  in an ascending order and the receiver obtains an item  $x = 16$  through the third OT, then the receiver can learn that two people in  $X$  are under the age of 16.

Protocol  $\Pi_{\text{PSU}}^{\text{R}}$  using Cuckoo hashing**Parameters:**

- Hash functions  $h_1, \dots, h_\gamma : \{0, 1\}^{\ell_1} \rightarrow [b]$ ;
- A Cuckoo hash table based on  $h_1, \dots, h_\gamma$  with  $b = \epsilon \cdot n_2$  bins, stash size  $s = 0$ ;
- Ideal functionalities  $\mathcal{F}_{\text{PS}}$ ,  $\mathcal{F}_{\text{OT}}$  and  $\mathcal{F}_{\text{mpOPRF}}$  (the underlying PRF is  $F(k, \cdot) : \{0, 1\}^{\ell_1} \rightarrow \{0, 1\}^{\ell_2}$ );

**Inputs:**

- Sender  $\mathcal{S}$ : set  $X = \{x_1, \dots, x_{n_1}\}, x_i \in \{0, 1\}^{\ell_1}$ ;
- Receiver  $\mathcal{R}$ : set  $Y = \{y_1, \dots, y_{n_2}\}, y_i \in \{0, 1\}^{\ell_1}$ ;

**Protocol:**

1.  $\mathcal{R}$  inserts set  $Y$  into the Cuckoo hash table based on  $h_1, \dots, h_\gamma$  as shown in Section 2.2, and adds a dummy item  $d$  in each empty bin, then denotes the filled Cuckoo hash table as  $Y_C$  and the item in  $i$ -th bin as  $Y_C[i]$ ;
2.  $\mathcal{S}$  and  $\mathcal{R}$  invoke the ideal functionality  $\mathcal{F}_{\text{PS}}$ :
  - $\mathcal{R}$  acts as  $P_0$  with input set  $Y_C$ , and  $\mathcal{S}$  acts as  $P_1$  with a permutation  $\pi$ ;
  - $\mathcal{R}$  obtains the shuffled shares  $\{a_1, a_2, \dots, a_b\}$ , and  $\mathcal{S}$  obtains the other shuffled shares  $\{a'_1, a'_2, \dots, a'_b\}$  where  $Y_C[\pi(i)] = a'_i \oplus a_i$ ;
3.  $\mathcal{S}$  and  $\mathcal{R}$  invoke the ideal functionality  $\mathcal{F}_{\text{mpOPRF}}$ :
  - $\mathcal{R}$  acts as  $P_0$  with her shuffled shares  $\{a_i\}_{i \in [b]}$ , and obtains the outputs  $\{F(k, a_i)\}_{i \in [b]}$ ;
  - $\mathcal{S}$  obtains the key  $k$ ;
4.  $\mathcal{R}$  initializes set  $Z = \emptyset$ ,  $\mathcal{S}$  randomly selects a permutation  $\pi'$ , and obtains  $\pi'(X) = \{x'_1, x'_2, \dots, x'_{n_1}\}$ ;
5. For  $i \in [n_1]$ :
  - $\mathcal{S}$  initializes sets  $Q_i = \emptyset$  and  $I_i = \emptyset$ ;
  - For  $j \in [\gamma]$ :
    - $\mathcal{S}$  computes  $q_j = \pi^{-1}(h_j(x'_i))$ ;
    - if  $q_j \notin Q_i$ ,  
 $Q_i = Q_i \cup \{q_j\}$ ,  $I_i = I_i \cup \{F(k, x'_i \oplus a'_{q_j})\}$ , else,  
 $r \xleftarrow{\$} \{0, 1\}^{\ell_2}$  and  $I_i = I_i \cup \{r\}$ ;
  - $\mathcal{S}$  sends  $I_i$  to  $\mathcal{R}$ ;
  - $\mathcal{R}$  checks if  $\{F(k, a_j)\}_{j \in [b]} \cap I_i \neq \emptyset$ . If so,  $\mathcal{R}$  sets  $b_i = 1$ , otherwise, sets  $b_i = 0$ ;
  - $\mathcal{S}$  and  $\mathcal{R}$  invoke the ideal functionality  $\mathcal{F}_{\text{OT}}$ :
    - $\mathcal{S}$  acts as  $P_0$  with input  $\{x'_i, \perp\}$ ;
    - $\mathcal{R}$  acts as  $P_1$  with input  $b$ ;
    - if  $b_i = 0$ ,  $\mathcal{R}$  obtains  $x'_i$ , otherwise, obtains  $\perp$ ;
  - Once receiving  $x'_i$ ,  $\mathcal{R}$  sets  $Z = Z \cup \{x'_i\}$ ;
6.  $\mathcal{R}$  outputs  $Y \cup Z$ ;

**Fig. 12.** Protocol  $\Pi_{\text{PSU}}^{\text{R}}$  using Cuckoo Hashing.



but  $I_i[j_i]$  is replaced with a random value  $r \xleftarrow{\$} \{0, 1\}^{\ell_2}$ . In Particular, for each  $x'_i$ , the sender initializes a set  $I_i = \emptyset$  and a set  $Q_i = \emptyset$ , where  $Q_i$  is used to record the indices of the shares (in  $\{a'_1, \dots, a'_b\}$ ) that are XORed with  $x'_i$ . Then for each  $j \in [\gamma]$ , the sender computes  $q_j = \pi^{-1}(h_j(x'_i))$ . If it does not appear before (i.e.,  $q_j \notin Q_i$ ), the sender adds it into  $Q_i$  and records  $F(k, x'_i \oplus a'_{q_j})$  into  $I_i$ . Otherwise, the sender randomly chooses  $r \in \{0, 1\}^{\ell_2}$  and records it into  $I_i$ . At the end, the sender sends  $I_i$  to the receiver. Recall that if  $x'_i \in Y$ , then there is an item in  $I_i$  that belongs to  $\{F(k, a_i)\}_{i \in [b]}$ .

Finally, upon receiving  $I_i$ , the receiver checks if the intersection of  $I_i$  and  $\{F(k, a_j)\}_{j \in [b]}$  is non-empty. If not, the receiver sets  $b_i = 0$  and obtains  $x_i$  through  $\mathcal{F}_{\text{OT}}$  and adds it to an initially empty set  $Z$ , otherwise sets  $b_i = 1$  and obtains nothing. At last, the receiver outputs  $Y \cup Z$ .

Next we first argue that the protocol  $\Pi_{\text{PSU}}^{\text{R}}$  in Figure 12 realizes the functionality  $\mathcal{F}_{\text{PSU}}^{n_1, n_2}$  correctly, and then show it satisfies the security properties.

**Correctness.** The receiver obtains the Cuckoo hash table  $Y_C$  filled with the set  $Y$  and dummy items  $d$ . Then through  $\mathcal{F}_{\text{PS}}$ , the receiver obtains the shuffled secret share set  $\{a_1, \dots, a_b\}$  of  $Y_C$  and the sender receives the other shares  $\{a'_1, \dots, a'_b\}$ , where  $a'_i \oplus a_i = Y_C[\pi(i)]$ . For an item  $x^* \in X$ , if  $x^* \in Y$ , say  $x^* = y_i$ , then  $y_i$  must be inserted into one of the bins of  $Y_C$  indexed by  $\{h_j(x^*)\}_{j \in [\gamma]}$ , and so the share of  $y_i$  held by the sender must belong to  $\{a'_{q_1}, \dots, a'_{q_\gamma}\}$  where  $q_j = \pi^{-1}(h_j(x^*))$  for all  $j \in [\gamma]$ . Without loss of generality, we assume that the share of  $y_i$  is  $a'_{q_w}$ , then  $x^* \oplus a'_{q_w} = y_i \oplus a'_{q_w} = a_{q_w}$ , and thus  $F(k, x^* \oplus a'_{q_w}) = F(k, a_{q_w})$ . So in this case the intersection of  $I^* = \{F(k, x^* \oplus a'_{q_j})\}_{j \in [\gamma]}$  and  $\{F(k, a_i)\}_{i \in [b]}$  is non-empty, the receiver sets  $b^* = 1$  and receives nothing from  $\mathcal{F}_{\text{OT}}$ . Otherwise (i.e.,  $x^* \notin Y$ ), we have  $x^* \oplus a'_{q_j} \neq a_{q_j}$  for all  $j \in [\gamma]$ . Moreover, for any  $a_t$  where  $t \in [b] \setminus \{q_j\}_{j \in [\gamma]}$ ,  $x^* \oplus a'_{q_j} \neq a_t$  with an overwhelming probability, as long as the length  $\ell_1$  of the share is sufficiently large. Thus the intersection of  $\{F(k, x^* \oplus a'_{q_j})\}_{j \in [\gamma]}$  and  $\{F(k, a_i)\}_{i \in [b]}$  is empty except for a negligible probability, then the receiver will set  $b^* = 0$  and receive  $x^*$  through  $\mathcal{F}_{\text{OT}}$ .

We remark that the correctness error comes from the following two types of collisions. Specifically, the first type is incurred by the secret shares, that is, for  $x^* \notin Y$  and some  $j \in [\gamma]$ ,  $x^* \oplus a'_{q_j} \in \{a_1, \dots, a_b\}$  holds. The other case is incurred by PRF, that is  $F(k, x^* \oplus a'_{q_j}) = F(k, a_t)$  for some  $x^* \oplus a'_{q_j} \neq a_t$ . To make the correctness hold with an overwhelming probability, we need to ensure the probability of collisions happening is less than  $2^{-\lambda}$ . To this end, we set both the share length  $\ell_1$  and the PRF output length  $\ell_2$  to be at least  $\lambda + \log(\epsilon n_2) + \log(\gamma n_1)$ .

**Security.** Now we proceed to show the semi-honest security of  $\Pi_{\text{PSU}}^{\text{R}}$  in the  $\{\mathcal{F}_{\text{PS}}, \mathcal{F}_{\text{mpOPRF}}, \mathcal{F}_{\text{OT}}\}$ -hybrid model.

**Theorem 1.** *The protocol  $\Pi_{\text{PSU}}^{\text{R}}$  presented in Figure 12 securely realizes  $\mathcal{F}_{\text{PSU}}^{n_1, n_2}$  in the  $\{\mathcal{F}_{\text{PS}}, \mathcal{F}_{\text{mpOPRF}}, \mathcal{F}_{\text{OT}}\}$ -hybrid model, in the presence of semi-honest adversaries.*

*Proof.* We construct  $\text{Sim}_{\mathcal{S}}$  and  $\text{Sim}_{\mathcal{R}}$  to simulate the views of corrupted sender  $\mathcal{S}$  and corrupted receiver  $\mathcal{R}$  respectively.

**Corrupt sender.** Simulator  $\text{Sim}_{\mathcal{S}}$  simulates a real execution in which the sender  $\mathcal{S}$  is corrupted. Since  $\mathcal{A}$  is semi-honest,  $\text{Sim}_{\mathcal{S}}$  can obtain the input  $X = \{x_1, \dots, x_{n_1}\}$  where  $x_i \in \{0, 1\}^{\ell_1}$  of  $\mathcal{S}$  directly, and externally send the set  $X$  to  $\mathcal{F}_{\text{PSU}}^{n_1, n_2}$ . When receiving a permutation  $\pi$  from  $\mathcal{A}$ ,  $\text{Sim}_{\mathcal{S}}$  checks if it is a permutation of  $b$  items. If so,  $\text{Sim}_{\mathcal{S}}$  randomly selects  $\{a'_1, \dots, a'_b\}$  where  $a'_i \in \{0, 1\}^{\ell_1}$  as shuffled shares, and simulates  $\mathcal{F}_{\text{PS}}$  sending them to  $\mathcal{A}$ . Then,  $\text{Sim}_{\mathcal{S}}$  randomly selects a key  $k$  of PRF and sends it to  $\mathcal{A}$  to simulate  $\mathcal{F}_{\text{mpOPRF}}$ , and waits for the set  $I$  from  $\mathcal{A}$ . As for the  $\mathcal{F}_{\text{OT}}$ ,  $\mathcal{A}$  does not need to obtain outputs from it, and thus  $\text{Sim}_{\mathcal{S}}$  does nothing.

We argue that the outputs of  $\text{Sim}_{\mathcal{S}}$  are indistinguishable from the real view of  $\mathcal{S}$  by the following hybrids:

Hyb<sub>0</sub>:  $\mathcal{S}$ 's view in the real protocol.

Hyb<sub>1</sub>: Same as Hyb<sub>0</sub> except that the output of  $\mathcal{F}_{\text{PS}}$  is replaced by  $\{a'_1, \dots, a'_b\}$  chosen by  $\text{Sim}_{\mathcal{S}}$ , and  $\text{Sim}_{\mathcal{S}}$  runs the  $\mathcal{F}_{\text{PS}}$  simulator to produce the simulated view for  $\mathcal{S}$ . The security of protocol  $\Pi_{\text{PS}}$  guarantees the view in simulation is computationally indistinguishable from the view in the real protocol.

Hyb<sub>2</sub>: Same as Hyb<sub>1</sub> except that the output key of  $\mathcal{F}_{\text{mpOPRF}}$  is replaced by the  $k$  chosen by  $\text{Sim}_{\mathcal{S}}$ , and  $\text{Sim}_{\mathcal{S}}$  runs the  $\mathcal{F}_{\text{mpOPRF}}$  simulator to produce the simulated view for  $\mathcal{S}$ . The security of protocol  $\Pi_{\text{mpOPRF}}$  guarantees the view in simulation is computationally indistinguishable from the view in the real protocol.

Hyb<sub>3</sub>: Same as Hyb<sub>2</sub> except that  $\text{Sim}_{\mathcal{S}}$  runs the  $\mathcal{F}_{\text{OT}}$  simulator to produce the simulated view for  $\mathcal{S}$ . This hybrid is computationally indistinguishable from Hyb<sub>2</sub> by the security of the protocol  $\Pi_{\text{OT}}$ . The hybrid is the view output by  $\text{Sim}_{\mathcal{S}}$ .

**Corrupt receiver.** Simulator  $\text{Sim}_{\mathcal{R}}$  simulates a real execution in which the receiver  $\mathcal{R}$  is corrupted.  $\text{Sim}_{\mathcal{R}}$  begins by invoking  $\mathcal{A}$  who corrupts  $\mathcal{R}$ , and receives the set  $Y = \{y_1, \dots, y_{n_2}\}$  (s.t.  $y_i \in \{0, 1\}^{\ell_1}$ ) that  $\mathcal{A}$  would send to  $\mathcal{F}_{\text{PS}}$  in a real execution. Then,  $\text{Sim}_{\mathcal{R}}$  forwards the set  $Y$  to the ideal functionality  $\mathcal{F}_{\text{PSU}}^{n_1, n_2}$ , and waits for the output from  $\mathcal{F}_{\text{PSU}}^{n_1, n_2}$ . To simulate  $\mathcal{F}_{\text{PS}}$  outputting the shuffled shares of  $Y_C$ ,  $\text{Sim}_{\mathcal{R}}$  randomly chooses  $\{a_1, \dots, a_b\}$  where  $a_i \in \{0, 1\}^{\ell_1}$ , and sends them to  $\mathcal{A}$ . Once receiving the input  $\{a_1, \dots, a_b\}$  of  $\mathcal{F}_{\text{mpOPRF}}$  from  $\mathcal{A}$ ,  $\text{Sim}_{\mathcal{R}}$  randomly chooses a key  $k$  of the PRF  $F(\cdot, \cdot)$ , and simulate  $\mathcal{F}_{\text{mpOPRF}}$  sending  $\{F(k, a_1), \dots, F(k, a_b)\}$  to  $\mathcal{A}$ .

After obtaining the output  $|X \cap Y|$  and  $X \cup Y$  from  $\mathcal{F}_{\text{PSU}}^{n_1, n_2}$ ,  $\text{Sim}_{\mathcal{R}}$  needs to simulate generating a set  $I$  for each item in  $\pi'(X) = \{x'_1, \dots, x'_{n_1}\}$ . Note that  $\text{Sim}_{\mathcal{R}}$  does not know  $\pi'(X) = \{x'_1, \dots, x'_{n_1}\}$ .  $\text{Sim}_{\mathcal{R}}$  knows the number  $|X \cap Y|$ , so he can randomly select  $|X \cap Y|$  items from  $[n_1]$  to form a set  $L$  in ascending order. This process refers to  $\text{Sim}_{\mathcal{R}}$  randomly selecting the indexes of items in  $\pi'(X)$  that belong to the intersection and recording these indexes in the set  $L$ . In addition,  $\text{Sim}_{\mathcal{R}}$  randomly chooses  $|X \cap Y|$  items from  $\{F(k, a_1), \dots, F(k, a_b)\}$  to form a set  $S$  in random order, which is used to generate the set  $I$  for each item in  $X \cap Y$ . For each  $x'_i$ , if  $i \in L$  and  $i = L[\hat{i}]$ ,  $\text{Sim}_{\mathcal{R}}$  inserts  $S[\hat{i}]$  and other  $\gamma - 1$  randomnesses that are not equal to any  $\{F(k, a_1), \dots, F(k, a_b)\}$  into set  $I$  and sends  $I$  to  $\mathcal{A}$ . If  $i \notin L$ ,  $\text{Sim}_{\mathcal{R}}$  randomly selects  $\gamma$  randomnesses that are not equal to any  $\{F(k, a_1), \dots, F(k, a_b)\}$  to insert into the corresponding set  $I$  and sends it to  $\mathcal{A}$ . When receiving the input  $b_i$  of  $\mathcal{F}_{\text{OT}}$  from  $\mathcal{A}$ , if  $b_i = 0$ ,  $\text{Sim}_{\mathcal{R}}$  takes one item from  $\{X \cup Y\} - Y$  in random order and sends it to  $\mathcal{A}$ , otherwise,  $\text{Sim}_{\mathcal{R}}$  sends  $\perp$ .

We argue that the outputs of  $\text{Sim}_{\mathcal{R}}$  are indistinguishable from the real execution through the following hybrids:

Hyb<sub>0</sub>:  $\mathcal{R}$ 's view in the real protocol.

Hyb<sub>1</sub>: Same as Hyb<sub>0</sub> but the protocol aborts if there exists  $a_t = a'_i \oplus x'_j$  where  $a_t$  is not  $a_i$  or  $F(k, a_t) = F(k, a'_i \oplus x'_j)$  for  $a_t \neq a'_i \oplus x'_j$ . The aborting probability can be  $2^{-\lambda}$  by setting the share length  $\ell_1$  and  $F(k, \cdot)$  output length  $\ell_2$ .

Hyb<sub>2</sub>: Same as Hyb<sub>1</sub> except that the output of  $\mathcal{F}_{\text{PS}}$  is replaced by  $\{a_1, \dots, a_b\}$  chosen by  $\text{Sim}_{\mathcal{R}}$ , and  $\text{Sim}_{\mathcal{R}}$  runs the  $\mathcal{F}_{\text{PS}}$  simulator to produce the simulated view for  $\mathcal{R}$ . The security of protocol  $\Pi_{\text{PS}}$  and the pseudorandomness of  $F(k, \cdot)$  guarantee the view in simulation is computationally indistinguishable from the view in the real protocol.

Hyb<sub>3</sub>: Same as Hyb<sub>2</sub> except that the output of  $\mathcal{F}_{\text{mpOPRF}}$  and corresponding items in sets  $I$  are replaced by  $\{F(k, a_1), \dots, F(k, a_b)\}$  computed by  $\text{Sim}_{\mathcal{R}}$ , and  $\text{Sim}_{\mathcal{R}}$  runs the  $\mathcal{F}_{\text{mpOPRF}}$  simulator to produce the simulated view for  $\mathcal{R}$ . More concretely, in Hyb<sub>2</sub>, the output of  $\mathcal{F}_{\text{mpOPRF}}$  to  $\mathcal{R}$  is still  $\{F(k^*, a_1^*), \dots, F(k^*, a_b^*)\}$ , which are the values in the real protocol. And the set  $I$  of an item in  $X \cap Y$  will contain an item in  $\{F(k^*, a_1^*), \dots, F(k^*, a_b^*)\}$ . In this hybrid, each  $F(k^*, a_i^*)$  is replaced by  $F(k, a_i)$ . The security of protocol  $\Pi_{\text{mpOPRF}}$  guarantees the view in simulation is computationally indistinguishable from the view in the real protocol. Note that at this point, the items in

$\{F(k, a_1), \dots, F(k, a_b)\}$  that also exist in set  $I$  are not the items selected into set  $S$  by  $\text{Sim}_{\mathcal{R}}$ . The next hybrid is to change these items.

**Hyb<sub>4</sub>**: Same as **Hyb<sub>3</sub>** except that the items in  $\{F(k, a_1), \dots, F(k, a_b)\}$  that belong to some sets  $I$  are changed to the items in  $S$  and the corresponding items in set  $I$  are also changed. This hybrid is perfectly indistinguishable from **Hyb<sub>3</sub>** as the permutation  $\pi$  is random.

**Hyb<sub>5</sub>**: Same as **Hyb<sub>4</sub>** except that the items in sets  $I$  that are not equal to the PRF values of the receiver are changed to the randomnesses chosen by  $\text{Sim}_{\mathcal{R}}$ . **Hyb<sub>5</sub>** is computationally indistinguishable from **Hyb<sub>4</sub>** because of the pseudorandom property of PRF.

**Hyb<sub>6</sub>**: Same as **Hyb<sub>5</sub>** except that the sets  $I$  are sent in the order chosen by  $\text{Sim}_{\mathcal{R}}$ . This hybrid is perfectly indistinguishable from **Hyb<sub>5</sub>** as the permutation  $\pi'$  is random.

**Hyb<sub>7</sub>**: Same as **Hyb<sub>6</sub>** except that  $\text{Sim}_{\mathcal{R}}$  runs the  $\mathcal{F}_{\text{OT}}$  simulator to produce the simulated view for  $\mathcal{R}$ . This hybrid is computationally indistinguishable from **Hyb<sub>5</sub>** by the security of the protocol  $\Pi_{\text{OT}}$ . The hybrid is the view output by  $\text{Sim}_{\mathcal{R}}$ .  $\square$

The protocol  $\Pi_{\text{PSU}}^{\text{R}}$  is very scalable for balanced datasets, as demonstrated in Section 4. When considering unbalanced datasets, we observe that it is already considerably efficient for the case that the sender's set size is much larger than the receiver's, but not so friendly for the opposite case. To deal with this case, we propose a second protocol  $\Pi_{\text{PSU}}^{\text{S}}$  as below.

### 3.2 Protocol $\Pi_{\text{PSU}}^{\text{S}}$ : Shuffling Sender's Set

In contrast to  $\Pi_{\text{PSU}}^{\text{R}}$ , the core idea of designing  $\Pi_{\text{PSU}}^{\text{S}}$  is to shuffle the sender's set, rather than the receiver's. A brief description is given below and the details are shown in Figure 13.

Similarly, the sender's set and the receiver's set are assumed to be  $X = \{x_1, \dots, x_{n_1}\}$  and  $Y = \{y_1, \dots, y_{n_2}\}$ , respectively. Then the protocol works as follows. Firstly, the sender inserts  $X$  into the Cuckoo hash table with  $b$  bins by using  $\gamma$  hash functions  $\{h_1, \dots, h_\gamma\}$  and fills each empty bin with a dummy item  $d$ . After that, the sender sends to the receiver the parameters of the Cuckoo hashing. Then the receiver inserts  $Y$  into a simple hash table with  $b$  bins by using the same hash functions. In general, when  $n$  items are inserted into a simple hash table with  $m$  bins using  $\gamma$  hash functions, the maximum bin size is  $O(\gamma n/m)$ <sup>12</sup> when  $n > m \log m$  according to [34]. Therefore, the maximum bin size  $\rho$  of simple hash table is  $O(\gamma n_2/b)$ . For simplicity, we denote by  $X_C$  and  $Y_B$  the filled Cuckoo hash table and simple hash table, respectively.

Secondly, the sender and the receiver invoke  $\mathcal{F}_{\text{PS}}$  with input  $X_C$  and a random permutation  $\pi$ , respectively. Then the sender obtains the shuffled secret shares  $\{a_1, \dots, a_b\}$  of  $X$  while the receiver obtains other shares  $\{a'_1, \dots, a'_b\}$ . Further, by running  $\mathcal{F}_{\text{mpOPRF}}$  with the input  $\{a_i\}_{i \in [b]}$  the sender obtains  $\{F(k, a_i)\}_{i \in [b]}$  and the receiver obtains the PRF key  $k$ .

Thirdly, for  $i \in [b]$  the receiver generates a set  $I_i$  with her  $i$ -th share  $a'_i$  and sends it to the sender, so that the sender can test if the item of  $X$  associated with  $a'_i$  belongs to  $Y$ . Note that as the receiver selects the permutation  $\pi$ , she knows that  $a'_i$  is the share of the  $\pi(i)$ -th item of  $X_C$ , whereas the sender does not know which item of  $X_C$  is being tested. Moreover, if  $X_C[\pi(i)] \in Y$ , then the item in  $Y$  equal to  $X_C[\pi(i)]$  must be contained in the  $\pi(i)$ -th bin of  $Y_B$ . So to check if  $X_C[\pi(i)] \in Y$ , we need only to check if  $X_C[\pi(i)] \in Y_B[\pi(i)]$ . To do so, the receiver computes the XOR of  $a'_i$  and each item in  $Y_B[\pi(i)]$ , then evaluates PRFs over them, namely  $\{F(k, y_j \oplus a'_i)\}_{y_j \in Y_B[\pi(i)]}$ , and adds them into set  $I_i$ . To further hide from the sender the actual number of items in  $Y$  mapped to the bin, the receiver pads  $I_i$  with  $\rho - |Y_B[\pi(i)]|$  random values from  $\{0, 1\}^{\ell_2}$  if

<sup>12</sup> To be precise, the maximum bin size should be  $\Theta(\gamma n/m)$ .

$\Pi_{\text{PSU}}^{\text{S}}$  Protocol using Cuckoo hashing**Parameters:**

- Hash functions  $h_1, \dots, h_\gamma : \{0, 1\}^{\ell_1} \rightarrow [b]$ ;
- A Cuckoo hash table based on  $h_1, \dots, h_\gamma$ , with  $b = \epsilon \cdot n_1$  bins, stash size  $s = 0$ ;
- A simple hash table based on  $h_1, \dots, h_\gamma$ , with  $b = \epsilon \cdot n_1$  bins and bin size  $\rho$ , where  $\rho = O(\gamma n_2 / b)$ ;
- Ideal functionalities  $\mathcal{F}_{\text{PS}}$  and  $\mathcal{F}_{\text{mpOPRF}}$  (the underlying PRF is  $F(k, \cdot) : \{0, 1\}^{\ell_1} \rightarrow \{0, 1\}^{\ell_2}$ );

**Inputs:**

- Sender  $\mathcal{S}$ : set  $X = \{x_1, \dots, x_{n_1}\}, x_i \in \{0, 1\}^{\ell_1}$ ;
- Receiver  $\mathcal{R}$ : set  $Y = \{y_1, \dots, y_{n_2}\}, y_i \in \{0, 1\}^{\ell_1}$ ;

**Protocol:**

1.  $\mathcal{S}$  inserts set  $X$  into the Cuckoo hash table, and fills empty bins with the dummy item  $d$ , then denotes the filled Cuckoo hash table as  $X_C$  and the item in  $i$ -th bin as  $X_C[i]$ ;  $\mathcal{R}$  inserts set  $Y$  into the simple hash table, and deletes the duplicates in each bin, then denotes the set of items in the  $i$ -th bin as  $Y_B[i]$ ;
2.  $\mathcal{S}$  and  $\mathcal{R}$  invoke the ideal functionality  $\mathcal{F}_{\text{PS}}$ :
  - $\mathcal{S}$  acts as  $P_0$  with input set  $X_C$ , and  $\mathcal{R}$  acts as  $P_1$  with a permutation  $\pi$ ;
  - $\mathcal{S}$  obtains the shuffled shares  $\{a_1, a_2, \dots, a_b\}$ ,  $\mathcal{R}$  obtains the another shuffled shares  $\{a'_1, a'_2, \dots, a'_b\}$  where  $X_C[\pi(i)] = a_i \oplus a_i$ ;
3.  $\mathcal{S}$  and  $\mathcal{R}$  invoke the ideal functionality  $\mathcal{F}_{\text{mpOPRF}}$ :
  - $\mathcal{S}$  acts as  $P_0$  with his shuffled shares  $\{a_i\}_{i \in [b]}$ , and obtains the outputs  $\{F(k, a_i)\}_{i \in [b]}$ ;
  - $\mathcal{R}$  obtains the key  $k$ ;
4.  $\mathcal{R}$  initializes a set  $Z = \emptyset$ ,  $\mathcal{S}$  initializes a string  $U = 0^b$ ;
5. For  $i \in [b]$ :
  - $\mathcal{R}$  initializes a set  $I_i = \emptyset$ ;
  - For each  $y_j \in Y_B[\pi(i)]$ ,  $\mathcal{R}$  adds  $F(k, y_j \oplus a'_i)$  to  $I_i$ ;
  - $\mathcal{R}$  pads  $I_i$  up to bin size  $\rho$  by different  $r \xleftarrow{\$} \{0, 1\}^{\ell_2}$ , and sends  $I_i$  to  $\mathcal{S}$ ;
  - $\mathcal{S}$  checks if  $F(k, a_i)$  is in  $I_i$ , if not,  $\mathcal{S}$  sets  $U[i] = 1$ , otherwise, sets  $U[i] = 0$ ;
6.  $\mathcal{S}$  and  $\mathcal{R}$  invoke the ideal functionality  $\mathcal{F}_{\text{PS}}$ :
  - $\mathcal{R}$  acts as  $P_0$  with input set  $\{a'_i\}_{i \in [b]}$ , and  $\mathcal{S}$  acts as  $P_1$  with a random permutation  $\pi'$ ;
  - $\mathcal{S}$  and  $\mathcal{R}$  obtains the shuffled share sets  $\{s_1^1, s_2^1, \dots, s_b^1\}$  and  $\{s_1^2, s_2^2, \dots, s_b^2\}$  respectively, where  $s_i^1 \oplus s_i^2 = a'_{\pi'(i)}$ ;
7. For  $i \in [b]$ :
  - If  $U[\pi'(i)] = 1$ ,  $\mathcal{S}$  sets  $z_i = a_{\pi'(i)} \oplus s_i^1$ , otherwise, sets  $z_i = \perp$ , then sends  $z_i$  to  $\mathcal{R}$ ;
  - If  $z_i \neq \perp$  and  $z_i \oplus s_i^2 \neq d$ ,  $\mathcal{R}$  sets  $Z = Z \cup \{z_i \oplus s_i^2\}$ ;
8.  $\mathcal{R}$  outputs  $Y \cup Z$ ;

**Fig. 13.** Protocol  $\Pi_{\text{PSU}}^{\text{S}}$  using Cuckoo Hashing.

the number (i.e.,  $|Y_B[\pi(i)]|$ ) of items in  $I_i$  is less than the maximum bin size  $\rho$ . Note that if there is a  $y_j$  that is mapped to the bin multiple times using different hash functions,  $\Pi_{\text{PSU}}^{\text{S}}$  only puts it to the bin once, and thus there are no duplicates in  $I_i$ .

Finally, after receiving  $I_i \supseteq \{F(k, y_j \oplus a'_i)\}_{y_j \in Y_B[\pi(i)]}$ , the sender checks if  $F(k, a_i) \in I_i$ . If not, the sender sends  $a_i$  to the receiver in order for the receiver to obtain  $X_C[\pi(i)] = a_i \oplus a'_i$ , otherwise sends  $\perp$ . By this way, however, the receiver will learn that an item belonging to  $X \cap Y$  is in the bin  $Y_B[\pi(i)]$  if she receives  $\perp$ , which will leak the information about intersection to the receiver as mentioned in [18]. On the other hand, if the receiver obtains the sender's item, she may learn partial information about the whole sender's set as she knows the position in Cuckoo hash table to which the item is mapped.

To solve the above problems, we opt to postpone sending the sender's shares. More specifically, instead of sending the share directly, the sender first records which shares should be sent using a bit-string  $U$  and then sends them in a new order. Since the receiver needs to match the shares from the sender with her shares, the receiver also needs to permute her shares in the same order. To this end, we leverage  $\mathcal{F}_{\text{PS}}$  again. Roughly speaking, the receiver takes her shares  $\{a'_1, a'_2, \dots, a'_b\}$  as the input and the sender randomly chooses a permutation  $\pi'$  as his input. After  $\mathcal{F}_{\text{PS}}$ , the sender receives the shuffled shares  $\{s_1^1, s_2^1, \dots, s_b^1\}$ , and the receiver obtains the other shares  $\{s_1^2, s_2^2, \dots, s_b^2\}$  where  $s_i^1 \oplus s_i^2 = a'_{\pi'(i)}$ . As the sender knows the permutation  $\pi'$ , he can check if  $a_{\pi'(i)}$  should be sent to the receiver according to  $U[\pi'(i)]$ . If so, the sender sends  $z_i = a_{\pi'(i)} \oplus s_i^1$  to the receiver, otherwise sends  $z_i = \perp$ . Once receiving  $z_i$  from the sender, if it is not  $\perp$ , the receiver can calculate  $z_i \oplus s_i^2 = a_{\pi'(i)} \oplus s_i^1 \oplus s_i^2 = a_{\pi'(i)} \oplus a'_{\pi'(i)}$ . If the recovered item is not the dummy item  $d$ , the receiver will add it into  $Z$ . Finally, the receiver outputs  $Y \cup Z$ .

In what follows, we first show the correctness of the protocol  $\Pi_{\text{PSU}}^{\text{S}}$  in Figure 13 and then argue that it securely realizes the functionality  $\mathcal{F}_{\text{PSU}^*}^{n_1, n_2}$ .

**Correctness.** The sender gets the Cuckoo hashing table  $X_C$  filled with the set  $X$  and dummy items  $d$ , and the receiver gets the simple hashing table  $Y_B$  filled with the set  $Y$ . Then through  $\mathcal{F}_{\text{PS}}$ , the receiver gets the shuffled secret shares  $\{a'_1, \dots, a'_b\}$  of  $X_C$  and the sender gets the other shares  $\{a_1, \dots, a_b\}$ , where  $a'_i \oplus a_i = X_C[\pi(i)]$ .

For an item  $x^* \in X$ , if  $x^* \in Y$ , say  $x^* = y_j$  and  $a_q, a'_q$  are the shares of  $x^*$ , then  $y_j$  must be in the  $\pi(q)$ -th bin of the simple hashing table (i.e.,  $y_j \in Y_B[\pi(q)]$ ) and  $y_j \oplus a'_q = x^* \oplus a'_q = a_q$ . Therefore, once the sender gets  $I_q$  containing  $F(k, y_j \oplus a'_q)$ , he will find that  $F(k, a_q) \in I_q$  and thus sets  $U[q] = 0$ . Otherwise (i.e.,  $x^* \notin Y$ ),  $y_j \oplus a'_q \neq a_q$  for all  $y_j \in Y_B[\pi(i)]$ , thus  $F(k, a_q) \notin I_q$  and  $U[q]$  is set to be 1 with an overwhelming probability. Then through  $\mathcal{F}_{\text{PS}}$  with  $\{a'_1, \dots, a'_b\}$  and  $\pi'$  as the inputs, the sender gets the shuffled secret shares  $\{s_1^1, \dots, s_b^1\}$  and the receiver gets the other secret shares  $\{s_1^2, \dots, s_b^2\}$  where  $s_i^1 \oplus s_i^2 = a'_{\pi'(i)}$ . If  $U[q] = 1$  meaning  $x^* \notin Y$ , the receiver will get  $a_q \oplus s_i^1$  where  $i = \pi'^{-1}(q)$  and thus recover  $x^* = a_q \oplus s_i^1 \oplus s_i^2 = a_q \oplus a'_q$ , otherwise gets nothing about  $x^*$ .

As noted for  $\Pi_{\text{PSU}}^{\text{R}}$ , the correctness here is guaranteed unless collisions occur. In  $\Pi_{\text{PSU}}^{\text{S}}$ ,  $a_i$  is not equal to  $y_k \oplus a'_i$  for any  $y_k \in Y_B[\pi(i)]$ , if  $y_k \neq X_C[\pi(i)]$ . Therefore, the collisions can only come from the PRF, i.e.,  $F(k, y_j \oplus a'_i) = F(k, a_i)$  for  $y_j \oplus a'_i \neq a_i$ . Then by setting the output length  $\ell_2$  of  $F(k, \cdot)$  as  $\lambda + \log(\epsilon n_1) + \log(\gamma n_2)$ , we can bound the probability of collision happening to  $2^{-\lambda}$ .

**Security.** Now we show that  $\Pi_{\text{PSU}}^{\text{S}}$  securely realizes  $\mathcal{F}_{\text{PSU}^*}^{n_1, n_2}$ .

**Theorem 2.** *The protocol  $\Pi_{\text{PSU}}^{\text{S}}$  presented in Figure 13 securely realizes  $\mathcal{F}_{\text{PSU}^*}^{n_1, n_2}$  in the  $\{\mathcal{F}_{\text{PS}}, \mathcal{F}_{\text{mpOPRF}}\}$ -hybrid model, in the presence of semi-honest adversaries.*

*Proof.* We construct  $\text{Sim}_{\mathcal{S}}$  and  $\text{Sim}_{\mathcal{R}}$  to simulate the views of corrupted sender  $\mathcal{S}$  and corrupted receiver  $\mathcal{R}$  respectively.

**Corrupt sender.** Simulator  $\text{Sim}_{\mathcal{S}}$  simulates a real execution where the sender  $\mathcal{S}$  is corrupted.  $\text{Sim}_{\mathcal{S}}$  begins by invoking adversary  $\mathcal{A}$  who corrupts  $\mathcal{S}$ , and receives the

set  $X = \{x_1, \dots, x_{n_1}\}, x_i \in \{0, 1\}^{\ell_1}$  that  $\mathcal{A}$  would send to  $\mathcal{F}_{\text{PS}}$  in a real execution. Then,  $\text{Sim}_{\mathcal{S}}$  forwards the set  $X$  to the ideal functionality  $\mathcal{F}_{\text{PSU}^*}^{n_1, n_2}$ , and waits for the output from  $\mathcal{F}_{\text{PSU}^*}^{n_1, n_2}$ . To simulate  $\mathcal{F}_{\text{PS}}$  outputting the shuffled shares of  $X$ ,  $\text{Sim}_{\mathcal{S}}$  randomly chooses  $\{a_1, \dots, a_b\}$  where  $a_i \in \{0, 1\}^{\ell_1}$ , and sends them to  $\mathcal{A}$ . Once receiving the input  $\{a_1, \dots, a_b\}$  of  $\mathcal{F}_{\text{mpOPRF}}$  from  $\mathcal{A}$ ,  $\text{Sim}_{\mathcal{S}}$  will randomly choose a key  $k$  for PRF  $F(\cdot, \cdot)$ , and simulate  $\mathcal{F}_{\text{mpOPRF}}$  sending  $\{F(k, a_1), \dots, F(k, a_b)\}$  to  $\mathcal{A}$ .

After obtaining the output  $|X \cap Y|$  from  $\mathcal{F}_{\text{PSU}^*}^{n_1, n_2}$ ,  $\text{Sim}_{\mathcal{S}}$  begins to simulate generating a set  $I$  for each item in  $\{F(k, a_1), \dots, F(k, a_b)\}$ .  $\text{Sim}_{\mathcal{S}}$  randomly selects  $|X \cap Y|$  items from  $[b]$  to form set  $L$  in ascending order. The above process refers to  $\text{Sim}_{\mathcal{S}}$  randomly selecting items belonging to the intersection and recording their index in the set  $L$ . For each  $F(k, a_i)$ , if  $i \in L$ ,  $\text{Sim}_{\mathcal{S}}$  inserts  $F(k, a_i)$  and other  $\rho - 1$  randomnesses that are not equal to any  $\{F(k, a_1), \dots, F(k, a_b)\}$  into set  $I$  and sends  $I$  to  $\mathcal{A}$ . If  $i \notin L$ ,  $\text{Sim}_{\mathcal{S}}$  randomly selects  $\rho$  randomnesses that are not equal to any  $\{F(k, a_1), \dots, F(k, a_b)\}$  to insert into the corresponding set  $I$  and sends it to  $\mathcal{A}$ . When  $\mathcal{A}$  invokes  $\mathcal{F}_{\text{PS}}$  again with random permutation  $\pi'$  as input,  $\text{Sim}_{\mathcal{S}}$  randomly chooses  $\{s_1^1, \dots, s_b^1\}$  and sends the share set to  $\mathcal{A}$ .

We argue that the outputs of  $\text{Sim}_{\mathcal{S}}$  are indistinguishable from the real execution through the following hybrids:

**Hyb<sub>0</sub>**:  $\mathcal{S}$ 's view in the real protocol.

**Hyb<sub>1</sub>**: Same as **Hyb<sub>0</sub>** but the protocol aborts if there exists  $a_i \neq a'_i \oplus y_j$  such that  $F(k, a_i) = F(k, a'_i \oplus y_j)$ . The aborting probability can be  $2^{-\lambda}$  by setting the  $F(k, \cdot)$  output length  $\ell_2$ .

**Hyb<sub>2</sub>**: Same as **Hyb<sub>1</sub>** except that the output of the first  $\mathcal{F}_{\text{PS}}$  is replaced by  $\{a_1, \dots, a_b\}$  chosen by  $\text{Sim}_{\mathcal{S}}$ , and  $\text{Sim}_{\mathcal{S}}$  runs the  $\mathcal{F}_{\text{PS}}$  simulator to produce the simulated view for  $\mathcal{S}$ . The security of protocol  $\Pi_{\text{PS}}$  and pseudorandomness of  $F(k, \cdot)$  guarantee the view in simulation is computationally indistinguishable from the view in the real protocol.

**Hyb<sub>3</sub>**: Same as **Hyb<sub>2</sub>** except that the output of  $\mathcal{F}_{\text{mpOPRF}}$  and corresponding items in sets  $I$  are replaced by  $\{F(k, a_1), \dots, F(k, a_b)\}$  computed by  $\text{Sim}_{\mathcal{S}}$ , and  $\text{Sim}_{\mathcal{S}}$  runs the  $\mathcal{F}_{\text{mpOPRF}}$  simulator to produce the simulated view for  $\mathcal{S}$ . The security of protocol  $\Pi_{\text{mpOPRF}}$  guarantees the view in simulation is computationally indistinguishable from the view in the real protocol.

**Hyb<sub>4</sub>**: Same as **Hyb<sub>3</sub>** except that the items in  $\{F(k, a_1), \dots, F(k, a_b)\}$  that belong to some sets  $I$  are changed to the items indexed by the set  $L$  and the corresponding items in set  $I$  are also changed. This hybrid is perfectly indistinguishable from **Hyb<sub>3</sub>** as the permutation  $\pi$  is random.

**Hyb<sub>5</sub>**: Same as **Hyb<sub>4</sub>** except that the items in sets  $I$  that are not equal to the PRF values of the receiver are changed to the randomnesses chosen by  $\text{Sim}_{\mathcal{S}}$ . **Hyb<sub>5</sub>** is computationally indistinguishable from **Hyb<sub>4</sub>** because of the pseudorandom property of PRF.

**Hyb<sub>6</sub>**: Same as **Hyb<sub>5</sub>** except that the output of the second  $\mathcal{F}_{\text{PS}}$  is replaced by  $\{s_1^1, \dots, s_b^1\}$  chosen by  $\text{Sim}_{\mathcal{S}}$ , and  $\text{Sim}_{\mathcal{S}}$  runs the  $\mathcal{F}_{\text{PS}}$  simulator to produce the simulated view for  $\mathcal{S}$ . The security of protocol  $\Pi_{\text{PS}}$  guarantees the view in simulation is computationally indistinguishable from the view in the real protocol. The hybrid is the view output by  $\text{Sim}_{\mathcal{S}}$ .

**Corrupt receiver.** Simulator  $\text{Sim}_{\mathcal{R}}$  simulates a real execution in which the receiver  $\mathcal{R}$  is corrupted. Since  $\mathcal{A}$  is semi-honest,  $\text{Sim}_{\mathcal{R}}$  can obtain the input  $Y = \{y_1, \dots, y_{n_2}\}$  (s.t.  $y_i \in \{0, 1\}^{\ell_1}$ ) of  $\mathcal{R}$  directly, and externally send  $Y$  to  $\mathcal{F}_{\text{PSU}^*}^S$ . When receiving a permutation  $\pi$  from  $\mathcal{A}$ ,  $\text{Sim}_{\mathcal{R}}$  checks if it is a permutation of  $b$  elements. If so,  $\text{Sim}_{\mathcal{R}}$  randomly selects  $\{a'_1, \dots, a'_b\}$  where  $a'_i \in \{0, 1\}^{\ell_1}$  as shuffled shares, and simulates  $\mathcal{F}_{\text{PS}}$  sending them to  $\mathcal{A}$ . Then,  $\text{Sim}_{\mathcal{R}}$  randomly selects a key  $k$  of PRF and sends it to  $\mathcal{A}$  to simulate  $\mathcal{F}_{\text{mpOPRF}}$ , and waits for the set  $I$  from  $\mathcal{A}$ . When  $\mathcal{A}$  invokes  $\mathcal{F}_{\text{PS}}$  with  $\{a'_1, \dots, a'_b\}$  as input,  $\text{Sim}_{\mathcal{R}}$  randomly chooses  $\{s_1^2, \dots, s_b^2\}$  and sends them to  $\mathcal{A}$ .

Once receiving the output  $X \cup Y$  from  $\mathcal{F}_{\text{PSU}^*}^{n_1, n_2}$ ,  $\text{Sim}_{\mathcal{R}}$  randomly selects  $|\{X \cup Y\} - Y|$  items from  $[b]$  to form set  $L$  in ascending order. In addition, to simulate the dummy item,  $\text{Sim}_{\mathcal{R}}$  also needs to randomly selects  $b - n_1$  items from  $[b]$  to form set  $D$  in ascending order while guarantee that  $L \cap D = \emptyset$ . And we denote the set  $\{X \cup Y\} - Y$  as  $\hat{X}$ . For  $i \in [b]$ , if  $i \in L$  and  $i = L[i]$ ,  $\text{Sim}_{\mathcal{R}}$  sends  $z_i = \hat{X}[i] \oplus s_i^2$  to  $\mathcal{A}$ . If  $i \in D$ ,  $\text{Sim}_{\mathcal{R}}$  sends  $z_i = d \oplus s_i^2$  to  $\mathcal{A}$ . If  $i \notin L \cup D$ ,  $\text{Sim}_{\mathcal{R}}$  sends  $z_i = \perp$  to  $\mathcal{A}$ .

We argue that the outputs of  $\text{Sim}_{\mathcal{R}}$  are indistinguishable from the real view of  $\mathcal{R}$  by the following hybrids:

**Hyb<sub>0</sub>**:  $\mathcal{R}$ 's view in the real protocol.

**Hyb<sub>1</sub>**: Same as **Hyb<sub>0</sub>** except that the output of the first  $\mathcal{F}_{\text{PS}}$  is replaced by  $\{a'_1, \dots, a'_b\}$  chosen by  $\text{Sim}_{\mathcal{R}}$ , and  $\text{Sim}_{\mathcal{R}}$  runs the  $\mathcal{F}_{\text{PS}}$  simulator to produce the simulated view for  $\mathcal{R}$ . The security of protocol  $\Pi_{\text{PS}}$  guarantees the view in simulation is computationally indistinguishable from the view in the real protocol.

**Hyb<sub>2</sub>**: Same as **Hyb<sub>1</sub>** except that the output key of  $\mathcal{F}_{\text{mpOPRF}}$  is replaced by the  $k$  chosen by  $\text{Sim}_{\mathcal{R}}$ , and  $\text{Sim}_{\mathcal{R}}$  runs the  $\mathcal{F}_{\text{mpOPRF}}$  simulator to produce the simulated view for  $\mathcal{R}$ . The security of protocol  $\Pi_{\text{mpOPRF}}$  guarantees the view in simulation is computationally indistinguishable from the view in the real protocol.

**Hyb<sub>3</sub>**: Same as **Hyb<sub>2</sub>** except that  $z_i$  is sent in the order chosen by  $\text{Sim}_{\mathcal{R}}$  and the items in share set generated by the second  $\mathcal{F}_{\text{PS}}$  are also permuted according to the order chosen by  $\text{Sim}_{\mathcal{R}}$ . This hybrid is perfectly indistinguishable from **Hyb<sub>2</sub>** as the permutation  $\pi'$  is random.

**Hyb<sub>4</sub>**: Same as **Hyb<sub>3</sub>** except that the output of the second  $\mathcal{F}_{\text{PS}}$  is replaced by  $\{s'_1, \dots, s'_b\}$  chosen by  $\text{Sim}_{\mathcal{R}}$  and the corresponding  $z_i$  is changed according to  $s'_i$ , and  $\text{Sim}_{\mathcal{R}}$  runs the  $\mathcal{F}_{\text{PS}}$  simulator to produce the simulated view for  $\mathcal{R}$ . The security of protocol  $\Pi_{\text{PS}}$  guarantees the view in simulation is computationally indistinguishable from the view in the real protocol. The hybrid is the view output by  $\text{Sim}_{\mathcal{R}}$ .  $\square$

### 3.3 Cost Analysis

Given the statistical security parameter  $\lambda$ , according to [33], we choose the parameters of Cuckoo hashing without stash, exactly including  $b = \epsilon \cdot n$  bins (where  $n = n_2$  for  $\Pi_{\text{PSU}}^{\text{R}}$  and  $n = n_1$  for  $\Pi_{\text{PSU}}^{\text{S}}$ ) and  $\gamma$  hash functions, to ensure that the hashing failure probability is less than  $2^{-\lambda}$ . Besides, to guarantee the error rate incurred by collisions is less than  $2^{-\lambda}$  in  $\Pi_{\text{PSU}}^{\text{R}}$ , we set the share/item length  $\ell_1$  and the output length  $\ell_2$  of  $F(k, \cdot)$  to be at least  $\lambda + \log(\gamma n_1) + \log(\epsilon n_2)$ . Likewise, the output length  $\ell_2$  of  $F(k, \cdot)$  in  $\Pi_{\text{PSU}}^{\text{S}}$  is at least  $\lambda + \log(\epsilon n_1) + \log(\gamma n_2)$ . In addition, the costs of  $\Pi_{\text{PSU}}^{\text{R}}$  and  $\Pi_{\text{PSU}}^{\text{S}}$  also rely on the sub-protocols used to realize the building blocks. Particularly, we realize  $\mathcal{F}_{\text{PS}}$ ,  $\mathcal{F}_{\text{mpOPRF}}$  and  $\mathcal{F}_{\text{OT}}$  with the protocols in [21], [6] and [15], respectively.

**Table 4.** The costs of  $\Pi_{\text{PSU}}^{\text{R}}$  and  $\Pi_{\text{PSU}}^{\text{S}}$

	Part-1	Part-2	Part-3	Part-4
Comp.	$O(b \log b)$	$O(b)$	$O(\gamma n)$	$O(n_1)$
Comm.	$O(\ell_1 b \log b)$	$O(b)$	$O(\ell_2 \gamma n)$	$O(\ell_1 n_1)$

$\Pi_{\text{PSU}}^{\text{R}}$ :  $b = \epsilon n_2, n = n_1$ ;  $\Pi_{\text{PSU}}^{\text{S}}$ :  $b = \epsilon n_1, n = n_2$ ;  
Part-1 is executed twice.

Since  $\Pi_{\text{PSU}}^{\text{R}}$  and  $\Pi_{\text{PSU}}^{\text{S}}$  mainly consist of 4 parts: (1) Permute + Share, (2) multi-point OPRF, (3) computing and sending  $\{I_i\}$ , and (4) obtaining items in  $X \setminus Y$ , we for clarity summarize their complexities in Table 4 according to each part. In  $\Pi_{\text{PSU}}^{\text{R}}$  (resp.  $\Pi_{\text{PSU}}^{\text{S}}$ ), Part-1 and Part-2 are performed on the receiver's set  $Y$  (resp. the sender's set  $X$ ) while Part-3 is performed on the sender's set  $X$  (resp. the receiver's set  $Y$ ), and thus  $b = \epsilon n_2, n = n_1$  (resp.  $b = \epsilon n_1, n = n_2$ ).

		Protocol	set size $n$							
			$2^8$	$2^{10}$	$2^{12}$	$2^{14}$	$2^{16}$	$2^{18}$	$2^{20}$	$2^{22}$
Time (s)	WAN	[18]	1.064	1.379	2.164	5.326	17.541	86.358	333.073	1459.539
		$\Pi_{\text{PSU}}^{\text{R}}$	<b>0.671</b>	<b>0.892</b>	<b>1.132</b>	<b>1.778</b>	<b>4.412</b>	<b>16.104</b>	<b>67.756</b>	<b>341.758</b>
		$\Pi_{\text{PSU}}^{\text{S}}$	0.712	0.993	1.238	2.214	6.233	22.78	102.039	458.731
	LAN	[18]	0.578	0.69	1.278	3.551	13.285	69.19	263.476	1191.703
		$\Pi_{\text{PSU}}^{\text{R}}$	<b>0.265</b>	<b>0.308</b>	<b>0.412</b>	<b>0.87</b>	<b>2.702</b>	<b>10.751</b>	<b>48.703</b>	<b>251.091</b>
		$\Pi_{\text{PSU}}^{\text{S}}$	0.274	0.32	0.434	1.051	3.452	13.382	60.16	279.97
Comm.(MB)	[18]	0.41	1.86	7.72	31.8	131.17	600.62	2470.11	10233.28	
	$\Pi_{\text{PSU}}^{\text{R}}$	<b>0.22</b>	<b>0.814</b>	<b>3.576</b>	<b>15.848</b>	<b>70.198</b>	<b>307.192</b>	<b>1338.79</b>	<b>5779.599</b>	
	$\Pi_{\text{PSU}}^{\text{S}}$	0.376	1.554	7.019	31.381	140.604	617.654	2725.932	11746.69	

**Table 5.** Comparisons of total runtime (in seconds) and communication (in MB) between  $\Pi_{\text{PSU}}^{\text{R}}$ ,  $\Pi_{\text{PSU}}^{\text{S}}$  and [18] with a single thread in WAN/LAN settings where  $n_1 = n_2 = n$ . Best results are marked in bold.

## 4 Performance Evaluation

In this section, we experimentally evaluate our PSU protocols  $\Pi_{\text{PSU}}^{\text{R}}$  and  $\Pi_{\text{PSU}}^{\text{S}}$ . In Section 4.1, we first give our benchmarking environment. In section 4.2, we compare our protocols with the state-of-the-art work [18] in terms of communication cost and single-threaded runtime on different networks, and the results are reported in Table 5. To demonstrate the scalability and parallelizability of our protocols, we evaluate our two protocols on small and large sets with different threads in Section 4.3, and show the results in Table 6, Table 7 and Table 8. Besides the equal set sizes, we also consider the unbalanced sets in Section 4.4. We perform  $\Pi_{\text{PSU}}^{\text{R}}$  in the cases where the sender’s set is larger than the receiver’s set, and  $\Pi_{\text{PSU}}^{\text{S}}$  in the opposite cases, and show the results in Table 9. Our complete implementation is available on GitHub: <https://github.com/dujiajun/PSU>.

### 4.1 Benchmarking Environment

We implement  $\Pi_{\text{PSU}}^{\text{R}}$  and  $\Pi_{\text{PSU}}^{\text{S}}$  in C++, and run our experiments on a single Intel Xeon with 2.39GHz and 128GB RAM. We evaluate our protocols in two networks settings, LAN network with 10Gbps bandwidth and 0.02 ms RTT and WAN network with 400Mbps and 80ms RTT, which are emulated using Linux tc command. We set the computational security parameter  $\kappa = 128$  and statistical security parameter  $\lambda = 40$ , and the item length in bits  $\ell_1$  is 128.

Our protocols are built on Permute + Share, multi-point OPRF, and OT extension. We implement Permute + Share with the design in [21] and OT extension [15] using libOTe library [37] with Naor-Pinkas Base OT [23]. For multi-point OPRF, we use the source code from [6]. The details of Permute + Share sub-protocol and multi-point OPRF sub-protocol are shown in Appendix B.

**Parameters about Cuckoo Hashing.** For the equal set sizes, Permute + Share sub-protocol costs most of the runtime, and thus we need to minimize the number of items to be shuffled as far as possible. Moreover, items in the stash of Cuckoo hashing need to be compared with each item of the other party rather than certain items picked out by the hash functions. Hence, we also need to limit the stash size to be 0. The empirical analysis in [33] shows that increasing the number of hash functions can drastically reduce the number of bins and the required stash size. According to the results reported in [33], we decide to use 4 hash functions to implement Cuckoo hashing with  $1.09 \cdot n$  bins and 0 stash. However, for the unbalanced set sizes, calculating PRFs will dominate, to reduce the number of PRFs calls and keep the stash size 0, we choose to use 3 hash functions and  $1.27 \cdot n$  bins.



## 4.2 Performance Comparisons

In this section, we compare  $\Pi_{\text{PSU}}^{\text{R}}$ ,  $\Pi_{\text{PSU}}^{\text{S}}$  and [18] in terms of runtime and communication, and the results are reported in Table 5. More concretely, compared to [18], our  $\Pi_{\text{PSU}}^{\text{R}}$  can obtain a 4-5 $\times$  improvement in runtime for large datasets ( $n_1 = n_2 \geq 2^{14}$ ) in both WAN and LAN settings. And the communication is about 50% communication of [18].

Although we also consider  $\Pi_{\text{PSU}}^{\text{S}}$  in this comparison, it is worth noting that  $\Pi_{\text{PSU}}^{\text{S}}$  realises a different functionality than the other two protocols, because the sender can obtain the intersection size in advance. Moreover, to avoid the leakage of intersection information,  $\Pi_{\text{PSU}}^{\text{S}}$  has to execute the Permute + Share sub-protocol twice (cf. Section 3.2 for more details). Therefore, the runtime of  $\Pi_{\text{PSU}}^{\text{S}}$  is longer than that of  $\Pi_{\text{PSU}}^{\text{R}}$ . Nevertheless, compared to [18], our  $\Pi_{\text{PSU}}^{\text{S}}$  can still obtain a 3-4 $\times$  improvement in runtime for large datasets ( $n_1 = n_2 \geq 2^{14}$ ) in both WAN and LAN settings. But for communication, the cost of  $\Pi_{\text{PSU}}^{\text{S}}$  is almost equal to that of [18].

## 4.3 Scalability and Parallelizability

In this section, we show that our two protocols can be efficiently executed on small sets without Cuckoo hashing, which can simplify development in practice. And using Cuckoo hashing, our protocols can be scaled to large sets. Moreover, we show that our protocols can be executed in parallel. Specifically, for the set of size  $2^{22}$ ,  $\Pi_{\text{PSU}}^{\text{R}}$  and  $\Pi_{\text{PSU}}^{\text{S}}$  can achieve a speedup of 3.49 $\times$  and 2.89 $\times$  using 8 threads, respectively.

**The cases with small sets.** In table 6, we show that  $\hat{\Pi}_{\text{PSU}}^{\text{R}}$  and  $\hat{\Pi}_{\text{PSU}}^{\text{S}}$  can be executed on small sets directly without Cuckoo hashing (please see Appendix A for more details of  $\hat{\Pi}_{\text{PSU}}^{\text{R}}$  and  $\hat{\Pi}_{\text{PSU}}^{\text{S}}$ ). In LAN setting, they cost about 10 seconds with 8 threads on the set of size  $2^{12}$ . However, we can see that as the set size increases, the runtime and communication increase dramatically. For set size of  $2^{14}$ , the two protocols both need more than 3 minutes with 8 threads, and the communication is so much that it will affect the protocol executed in the WAN setting. Therefore, for large sets, we test the runtime and communication of the protocols with Cuckoo hashing.

		$\hat{\Pi}_{\text{PSU}}^{\text{R}}$				$\hat{\Pi}_{\text{PSU}}^{\text{S}}$			
		$2^8$	$2^{10}$	$2^{12}$	$2^{14}$	$2^8$	$2^{10}$	$2^{12}$	$2^{14}$
Time (s)	T=1	0.620	5.582	89.862	1423.955	0.526	5.827	86.037	1425.376
	T=2	0.432	3.108	45.295	722.29	0.358	2.862	44.967	719.325
	T=4	0.356	1.722	23.094	363.270	0.295	1.861	22.231	360.131
	T=8	0.349	1.067	11.713	183.181	0.261	0.986	11.640	183.838
Comm. (MB)		0.665	9.624	162.759	2828.476	0.641	9.588	162.63	2827.972

**Table 6.** Runtime (in seconds) and communication (in MB) of  $\hat{\Pi}_{\text{PSU}}^{\text{R}}$  and  $\hat{\Pi}_{\text{PSU}}^{\text{S}}$  for small set ( $n_1 = n_2 \in \{2^8, 2^{10}, 2^{12}, 2^{14}\}$ ) and threads  $T \in \{1, 2, 4, 8\}$  threads in LAN setting.

**The cases with large sets.** In both  $\Pi_{\text{PSU}}^{\text{R}}$  and  $\Pi_{\text{PSU}}^{\text{S}}$ , PRF values in set  $I$  are independent of each other, and thus can be calculated in parallel. In addition, the Permute + Share sub-protocol in [21] and the multi-point OPRF sub-protocol in [6] can be partially parallelized. We demonstrate the scalability and parallelizability of  $\Pi_{\text{PSU}}^{\text{R}}$  and  $\Pi_{\text{PSU}}^{\text{S}}$  by evaluating it on the large set sizes  $n_1 = n_2 = n \in \{2^{12}, 2^{14}, 2^{16}, 2^{18}, 2^{20}, 2^{22}\}$  in parallel with  $T \in \{1, 2, 4, 8\}$  threads. Table 7 shows the experimental results of  $\Pi_{\text{PSU}}^{\text{R}}$  in both WAN/LAN settings, and the last row presents the ratio between the runtime of the single thread and 8 threads. We can see that, the speedup becomes better as the set size increases. Specifically, when the set size is  $2^{22}$ , we can obtain a speedup of 2.08 $\times$  in WAN setting and 3.49 $\times$  in LAN setting. Similarly, we report the results of  $\Pi_{\text{PSU}}^{\text{S}}$  in

Table 8. On the whole, as the set sizes and number of threads increase, the runtime of  $\Pi_{\text{PSU}}^{\text{S}}$  changes in the same way as that of  $\Pi_{\text{PSU}}^{\text{R}}$ . However, the speedup of  $\Pi_{\text{PSU}}^{\text{S}}$  is less than that of  $\Pi_{\text{PSU}}^{\text{R}}$ , since  $\Pi_{\text{PSU}}^{\text{S}}$  performs Permute + Share sub-protocol twice, which is not completely parallelized.

		set size $n$					
		$2^{12}$	$2^{14}$	$2^{16}$	$2^{18}$	$2^{20}$	$2^{22}$
WAN	T=1	1.132	1.778	4.412	16.104	67.756	341.758
	T=2	1.127	1.658	3.315	11.025	48.321	230.218
	T=4	1.117	1.553	2.965	8.852	37.847	181.657
	T=8	0.957	1.512	2.626	7.666	34.701	163.82
LAN	T=1	0.412	0.87	2.702	10.751	48.703	251.091
	T=2	0.367	0.615	1.721	6.221	29.812	148.538
	T=4	0.351	0.489	1.256	3.96	21.272	107.298
	T=8	0.325	0.477	1.093	3.582	14.304	71.782
<b>Speedup</b>		<b>1.31-1.26×</b>	<b>1.18-1.95×</b>	<b>1.68-2.47×</b>	<b>2.10-3.00×</b>	<b>1.95-3.40×</b>	<b>2.08-3.49×</b>

**Table 7.** Scaling of  $\Pi_{\text{PSU}}^{\text{R}}$  with set size ( $n_1 = n_2 = n$ ) and number of threads ( $T \in \{1, 2, 4, 8\}$ ) in WAN/LAN settings.

		set size $n$					
		$2^{12}$	$2^{14}$	$2^{16}$	$2^{18}$	$2^{20}$	$2^{22}$
WAN	T=1	1.238	2.214	6.233	22.78	102.039	458.731
	T=2	1.368	1.984	4.731	16.346	77.137	347.897
	T=4	1.388	1.79	3.909	13.856	65.319	292.226
	T=8	1.196	1.711	3.504	12.041	59.736	258.244
LAN	T=1	0.434	1.051	3.452	13.382	60.16	279.97
	T=2	0.378	0.764	2.322	7.863	38.434	175.485
	T=4	0.356	0.614	1.685	5.632	25.842	116.678
	T=8	0.408	0.606	1.397	5.204	20.992	96.723
<b>Speedup</b>		<b>1.04-1.06×</b>	<b>1.29-1.73×</b>	<b>1.78-2.47×</b>	<b>1.89-2.57×</b>	<b>1.70-2.87×</b>	<b>1.78-2.89×</b>

**Table 8.** Scaling of  $\Pi_{\text{PSU}}^{\text{S}}$  with set size ( $n_1 = n_2 = n$ ) and number of threads ( $T \in \{1, 2, 4, 8\}$ ) in WAN/LAN settings.

#### 4.4 Design for Unbalanced Datasets

In this section, we show that  $\Pi_{\text{PSU}}^{\text{R}}$  and  $\Pi_{\text{PSU}}^{\text{S}}$  can be chosen according to the sizes of the two sets. Considering that the set to be shuffled is small ( $2^8$  or  $2^{12}$ ) and the items that will be calculated PRFs are too many (more than  $2^{16}$ ,  $2^{20}$  or  $2^{24}$ ), we adjust the parameters of Cuckoo hashing to 3 hash functions and  $1.27 \cdot n$  bins with stash size  $s = 0$  according to the results in [33]. Table 9 shows the performances of  $\Pi_{\text{PSU}}^{\text{R}}$  and  $\Pi_{\text{PSU}}^{\text{S}}$  with 8 threads in WAN/LAN settings.

		LAN						WAN					
		$2^8$			$2^{12}$			$2^8$			$2^{12}$		
$n_2$ (resp. $n_1$ )		$2^{16}$	$2^{20}$	$2^{24}$	$2^{16}$	$2^{20}$	$2^{24}$	$2^{16}$	$2^{20}$	$2^{24}$	$2^{16}$	$2^{20}$	$2^{24}$
$\Pi_{\text{PSU}}^{\text{R}}$		0.487	3.17	47.788	0.524	3.648	51.513	1.266	5.101	64.109	1.396	5.341	67.802
$\Pi_{\text{PSU}}^{\text{S}}$		0.511	2.918	44.606	0.581	2.958	48.379	1.042	3.759	51.043	1.406	3.789	57.352

**Table 9.** Runtime (in seconds) of  $\Pi_{\text{PSU}}^{\text{R}}$  for unbalanced set sizes ( $n_1 \in \{2^{16}, 2^{20}, 2^{24}\}$ ,  $n_2 \in \{2^8, 2^{12}\}$ ) and  $\Pi_{\text{PSU}}^{\text{S}}$  for unbalanced set sizes ( $n_1 \in \{2^8, 2^{12}\}$ ,  $n_2 \in \{2^{16}, 2^{20}, 2^{24}\}$ ) with 8 threads in WAN/LAN settings.

When the receiver’s set is much smaller than the sender’s set (i.e.,  $n_2 \ll n_1$ ), we perform  $\Pi_{\text{PSU}}^{\text{R}}$  to obtain the union. We can see that for the sender’s set of size  $2^{24}$ ,  $\Pi_{\text{PSU}}^{\text{R}}$  only needs about 50 seconds in LAN setting, and about 65 seconds in WAN setting, which is reasonable in practice. As for the opposite unbalanced cases (i.e.,  $n_1 \ll n_2$ ),  $\Pi_{\text{PSU}}^{\text{S}}$  can obtain a better performance since it replaces OT related to the larger set with Permute + Share just related to the smaller set.

## 5 Leakage Analysis on [18]

In this section, we first recall the optimization via bucketing in [18], and then explain in detail why the usage of bucketing technique will leak the intersection information. Please refer to Appendix C.1 for more details of the protocol in [18]. Also, we further explain why the protocol in [18] cannot benefit from Cuckoo hashing in Appendix C.2.




Fig. 14. The bucketing technique in [18].

**Optimization via bucketing.** In order to improve the performance, Kolesnikov et al. [18] proposed to optimize their protocol by using the bucketing technique, as shown in Figure 14. More specifically, the sender and receiver in [18] first assign their items in  $X$  and in  $Y$ , into two simple hash tables with the same number of bins, and the maximum bin sizes are assumed to be  $\rho_1$  and  $\rho_2$ , respectively. Then they perform the  $(\rho_1, \rho_2)$ -PSU sub-protocol on the items of each bin separately. As pointed out by Kolesnikov et al. in [18], however, *the bucketing technique will leak the information “which bins contain items in  $X \cap Y$ ” to the receiver.* To avoid this leakage, in [18] the receiver is required to put a special item  $\perp$  into each bin, and to pad the bins with different dummy items  $d$ , while the sender pads his bins with the special item  $\perp$ . For example, in Figure 14, the items  $\{x_6, x_2, x_{10}\}$  of  $X$  are mapped to the first bin of the sender’s simple hash table, and the items  $\{y_3, y_8\}$  of  $Y$  are mapped to the first bin of the receiver’s hash table. Without the special item  $\perp$ , if  $x_2 = y_3$ , the receiver can learn that an item belonging to  $X \cap Y$  is in  $\{y_3, y_8\}$  after executing the  $(\rho_1, \rho_2)$ -PSU. By adding the special item  $\perp$  to both sides, if the receiver learns that an item from the sender belongs to  $\{y_3, \perp, y_8, d\}$ , it seems that the receiver cannot know whether the item is a real item (namely, in  $X$ ) or the special item  $\perp$ . Unfortunately, we observe that this strategy is insufficient to avoid the leakage incurred by the bucketing technique. A detailed analysis is given below.

**Leakage analysis.** For ease of exposition, we take the 4th  $(\rho_1, \rho_2)$ -PSU sub-protocol in Figure 14 as an example to explain why the optimization in [18] fails to hide the intersection information. After the execution of the sub-protocol over the 4th bins, if the receiver does not obtain any items from the sender (that is, all items in the sender’s 4th bin belong to the subset in the receiver’s 4th bin i.e.,  $\{d, \perp, y_5, y_7\}$ ), then the receiver could obtain additional information about the intersection. Concretely, one of the following will occur:

- **Case<sub>1</sub>**: all the real items that are mapped to the sender’s bin (say  $x_4$  in Figure 14) belong to  $\{y_5, y_7\}$ ;
- **Case<sub>2</sub>**: no real items are mapped to the sender’s bin (i.e., all items are special item  $\perp$ ).

We denote the probabilities that **Case<sub>1</sub>** and **Case<sub>2</sub>** occur by  $\Pr[\text{Case}_1]$  and  $\Pr[\text{Case}_2]$ , respectively. Clearly, if the receiver is able to determine that **Case<sub>1</sub>** occurs with certain (high) probability, she will know that items belonging to  $X \cap Y$  are in  $\{y_5, y_7\}$  with the same probability. Next we provide our estimation of the probability that **Case<sub>1</sub>** occurs. Note  $\Pr[\text{Case}_1] = 1 - \Pr[\text{Case}_2]$ ; we now bound  $\Pr[\text{Case}_2]$ . In general, assuming that there are  $\alpha n$  bins and  $n$  items and the hash mapping is a random oracle, the probability **Case<sub>2</sub>** occurs is  $P = \Pr[\text{Case}_2] = (1 - \frac{1}{\alpha n})^n \approx e^{-1/\alpha}$ .

**Table 10.** The probability of **Case<sub>2</sub>** for different set sizes

parameters	set size $n$							
	$2^8$	$2^{10}$	$2^{12}$	$2^{14}$	$2^{16}$	$2^{18}$	$2^{20}$	$2^{22}$
$\alpha$	0.043	0.055	0.05	0.053	0.058	0.052	0.06	0.051
$\Pr(\times 10^{-11})$	7.946	1270	206.1	639.4	3252	444.8	5778	305.1

Based on the parameters in [18], we calculate the probability  $P$  for different set sizes as shown in Table 10. From the results, we can see that the probability  $P$  is very small. For example, when the set size is  $n = 2^{20}$ ,  $P = 5.778 \times 10^{-8}$ . This means that when the receiver finds that all items in a bin belong to the intersection, she can learn that this bin has at least one real item with probability  $1 - 5.778 \times 10^{-8}$ , and that her corresponding bin contains at least an item in  $X \cap Y$  with the same probability. Hence, their approach is insufficient to avoid the leakage incurred by the bucketing technique.

On the contrary, in our  $\Pi_{\text{PSU}}^{\text{R}}$  shown in Figure 12, for an item  $x_i$  in  $X \cap Y$ , the receiver will find a  $F(k, a_j) \in I_i$  where  $a_j$  is the share of  $y_{\pi^{-1}(j)}$ , which means that  $y_{\pi^{-1}(j)}$  is the item equal to  $x_i$ . But from the receiver’s point of view, any item in  $Y$  may correspond to  $F(k, a_j)$  as she does not know  $\pi$ , and thus any item in  $Y$  may be the item in  $X \cap Y$ . In our  $\Pi_{\text{PSU}}^{\text{S}}$  shown in Figure 13, the receiver’s set  $Y$  is inserted into a simple hash table as in [18], but our protocol does not suffer from the leakage analyzed before. This is because we use the shuffling technique to hide which bins contain items in  $X \cap Y$ . To sum up, in  $\Pi_{\text{PSU}}^{\text{R}}$  and  $\Pi_{\text{PSU}}^{\text{S}}$ , any item in  $Y$  may be the item in  $X \cap Y$ , whereas in the protocol [18], the receiver can know that items belonging to  $X \cap Y$  are in a bin (namely, a subset of  $Y$ ) with a overwhelming probability.

## 6 Discussion

In this work, we focus on designing efficient PSU protocols for both balanced and unbalanced datasets. Somewhat surprisingly, our techniques can also be used for designing PSI protocols with only slight modification. With our techniques, it is extremely convenient to design fast protocols when both functionalities, set intersection and set union, are required; details are below.

Recall that in  $\Pi_{\text{PSU}}^{\text{R}}$ , the receiver sets  $b_i$  depending on whether the sender’s item  $x_i$  belongs to her set  $Y$ ; if  $x_i \in Y$ , sets  $b_i = 1$ , otherwise  $b_i = 0$ . Then through  $\mathcal{F}_{\text{OT}}$  the receiver obtains  $x_i$  if  $b_i = 0$ , and nothing otherwise. To obtain a PSI protocol from  $\Pi_{\text{PSU}}^{\text{R}}$ , the receiver only needs to set  $b'_i = b_i \oplus 1$  and obtains the sender’s items through  $\mathcal{F}_{\text{OT}}$  according to  $b'_i$ , rather than  $b_i$ . In this way, the receiver will obtain the sender’s items belonging to  $Y$ . Thus we obtain a PSI protocol, denoted by  $\Pi_{\text{PSI}}^{\text{R}}$ . In  $\Pi_{\text{PSU}}^{\text{S}}$ , the sender

sends the shares of the items in  $X \setminus Y$  according to the bit string  $U$ , then the receiver uses them to recover the items in  $X \setminus Y$ . Therefore, the sender can send the shares associated with the items in  $X \cap Y$  by flipping each bit in  $U$ , and then the receiver will obtain items in  $X \cap Y$ . Thus we obtain a PSI protocol, denoted by  $\Pi_{\text{PSI}}^{\text{S}}$ .

It can be seen that  $\Pi_{\text{PSI}}^{\text{R}}$  and  $\Pi_{\text{PSI}}^{\text{S}}$  are obtained from the PSU protocols with almost no extra overhead. Therefore, it is believed that they have nearly the same performance as the proposed PSU protocols. Due to the page limit, we leave their formal descriptions and security analysis to future work.

In addition, as stated in Section 1.3, it is desirable to have better PSU protocols designed. Notice that our protocols,  $\Pi_{\text{PSU}}^{\text{R}}$  and  $\Pi_{\text{PSU}}^{\text{S}}$  are designed in a modular manner; thus, a natural way to achieve better performance is to improve the performance of the underlying building blocks. More concretely, if the underlying `Permute + Share` protocol can be designed with linear complexity, the cost of  $\Pi_{\text{PSU}}^{\text{R}}$  and  $\Pi_{\text{PSU}}^{\text{S}}$  can be reduced to be linear. Finally, it is also interesting to design PSU with better security (e.g., defending against malicious adversaries) and/or with better functionalities (e.g., new variants of PSU including multi-party PSU and PSU with payload).

## References

1. SSL blacklist. <https://sslbl.abuse.ch/blacklist/>.
2. 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.
3. 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.
4. Martin Burkhart, Mario Strasser, Dilip Many, and Xenofontas A. Dimitropoulos. SEPIA: Privacy-preserving aggregation of multi-domain network events and statistics. In *USENIX Security 2010*, pages 223–240. USENIX Association, August 2010.
5. Melissa Chase, Esha Ghosh, and Oxana Poburinnaya. Secret-shared shuffle. In Shihor Moriai and Huaxiong Wang, editors, *Advances in Cryptology - ASIACRYPT 2020 - 26th International Conference on the Theory and Application of Cryptology and Information Security, Daejeon, South Korea, December 7-11, 2020, Proceedings, Part III*, volume 12493 of *Lecture Notes in Computer Science*, pages 342–372. Springer, 2020.
6. Melissa Chase and Peihan Miao. Private set intersection in the internet setting from lightweight oblivious PRF. In Daniele Micciancio and Thomas Ristenpart, editors, *Advances in Cryptology - CRYPTO 2020 - 40th Annual International Cryptology Conference, CRYPTO 2020, Santa Barbara, CA, USA, August 17-21, 2020, Proceedings, Part III*, volume 12172 of *Lecture Notes in Computer Science*, pages 34–63. Springer, 2020.
7. Alex Davidson and Carlos Cid. An efficient toolkit for computing private set operations. In Josef Pieprzyk and Suriadi Suriadi, editors, *ACISP 17, Part II*, volume 10343 of *LNCS*, pages 261–278. Springer, Heidelberg, July 2017.
8. Michael J. Freedman, Yuval Ishai, Benny Pinkas, and Omer Reingold. Keyword search and oblivious pseudorandom functions. In Joe Kilian, editor, *TCC 2005*, volume 3378 of *LNCS*, pages 303–324. Springer, Heidelberg, February 2005.
9. 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.
10. 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.
11. Carmit Hazay and Yehuda Lindell. Efficient protocols for set intersection and pattern matching with security against malicious and covert adversaries. In Ran Canetti, editor, *TCC 2008*, volume 4948 of *LNCS*, pages 155–175. Springer, Heidelberg, March 2008.
12. K. Hogan, N. Luther, N. Schear, E. Shen, D. Stott, S. Yakoubov, and A. Yerukhimovich. Secure multiparty computation for cooperative cyber risk assessment. In *2016 IEEE Cybersecurity Development (SecDev)*, pages 75–76, 2016.


13. Yan Huang, David Evans, and Jonathan Katz. Private set intersection: Are garbled circuits better than custom protocols? In *NDSS 2012*. The Internet Society, February 2012.
14. Russell Impagliazzo and Steven Rudich. Limits on the provable consequences of one-way permutations. In *21st ACM STOC*, pages 44–61. ACM Press, May 1989.
15. 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.
16. 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.
17. 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.
18. 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.
19. Yehuda Lindell. How to simulate it - A tutorial on the simulation proof technique. Cryptology ePrint Archive, Report 2016/046, 2016. <http://eprint.iacr.org/2016/046>.
20. Catherine 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, 1986.
21. 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, *EUROCRYPT 2013*, volume 7881 of *LNCS*, pages 557–574. Springer, Heidelberg, May 2013.
22. Rajeev Motwani and Prabhakar Raghavan. *Randomized algorithms*. Cambridge university press, 1995.
23. Moni Naor and Benny Pinkas. Oblivious transfer and polynomial evaluation. In *31st ACM STOC*, pages 245–254. ACM Press, May 1999.
24. Michele Orrù, Emanuela Orsini, and Peter Scholl. Actively secure 1-out-of-N OT extension with application to private set intersection. In Helena Handschuh, editor, *CT-RSA 2017*, volume 10159 of *LNCS*, pages 381–396. Springer, Heidelberg, February 2017.
25. Rasmus Pagh and Flemming Friche Rodler. Cuckoo hashing. In Friedhelm Meyer auf der Heide, editor, *Algorithms - ESA 2001, 9th Annual European Symposium, Aarhus, Denmark, August 28-31, 2001, Proceedings*, volume 2161 of *Lecture Notes in Computer Science*, pages 121–133. Springer, 2001.
26. Benny Pinkas, Mike Rosulek, Ni Trieu, and Avishay Yanai. SpOT-light: Lightweight private set intersection from sparse OT extension. In Alexandra Boldyreva and Daniele Micciancio, editors, *CRYPTO 2019, Part III*, volume 11694 of *LNCS*, pages 401–431. Springer, Heidelberg, August 2019.
27. Benny Pinkas, Mike Rosulek, Ni Trieu, and Avishay Yanai. PSI from PaXoS: Fast, malicious private set intersection. In Anne Canteaut and Yuval Ishai, editors, *EUROCRYPT 2020, Part II*, volume 12106 of *LNCS*, pages 739–767. Springer, Heidelberg, May 2020.
28. 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.
29. Benny Pinkas, Thomas Schneider, Nigel P. Smart, and Stephen C. Williams. Secure two-party computation is practical. In Mitsuru Matsui, editor, *ASIACRYPT 2009*, volume 5912 of *LNCS*, pages 250–267. Springer, Heidelberg, December 2009.
30. 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.
31. Benny Pinkas, Thomas Schneider, Christian Weinert, and Udi Wieder. Efficient circuit-based PSI via cuckoo hashing. In Jesper Buus Nielsen and Vincent Rijmen, editors, *EU-*

- ROCRYPT 2018, Part III*, volume 10822 of *LNCS*, pages 125–157. Springer, Heidelberg, April / May 2018.
32. Benny Pinkas, Thomas Schneider, and Michael Zohner. Faster private set intersection based on OT extension. In Kevin Fu and Jaeyeon Jung, editors, *USENIX Security 2014*, pages 797–812. USENIX Association, August 2014.
  33. Benny Pinkas, Thomas Schneider, and Michael Zohner. Scalable private set intersection based on ot extension. *ACM Trans. Priv. Secur.*, 21(2), January 2018.
  34. Martin Raab and Angelika Steger. “balls into bins” — a simple and tight analysis. In Michael Luby, José D. P. Rolim, and Maria Serna, editors, *Randomization and Approximation Techniques in Computer Science*, pages 159–170, Berlin, Heidelberg, 1998. Springer Berlin Heidelberg.
  35. Michael O. Rabin. How to exchange secrets with oblivious transfer. Cryptology ePrint Archive, Report 2005/187, 2005. <http://eprint.iacr.org/2005/187>.
  36. Sivaramakrishnan Ramanathan, Jelena Mirkovic, and Minlan Yu. BLAG: improving the accuracy of blacklists. In *27th Annual Network and Distributed System Security Symposium, NDSS 2020, San Diego, California, USA, February 23-26, 2020*. The Internet Society, 2020.
  37. Peter Rindal. libote: an efficient, portable, and easy to use oblivious transfer library. <https://github.com/osu-crypto/lib0Te>.
  38. Peter Rindal and Mike Rosulek. Malicious-secure private set intersection via dual execution. In Bhavani M. Thuraisingham, David Evans, Tal Malkin, and Dongyan Xu, editors, *ACM CCS 2017*, pages 1229–1242. ACM Press, October / November 2017.
  39. Peter Rindal and Phillipp Schoppmann. Vole-psi: Fast oprf and circuit-psi from vector-ole. In Anne Canteaut and François-Xavier Standaert, editors, *Advances in Cryptology – EUROCRYPT 2021*, pages 901–930, Cham, 2021. Springer International Publishing.

## A Additional Materials for Section 3: Simplified Protocols

In this section, we present the simplified version of protocol  $\Pi_{\text{PSU}}^{\text{R}}$  and of protocol  $\Pi_{\text{PSU}}^{\text{S}}$ , respectively. We note that, these two simplified versions, denoted by  $\hat{\Pi}_{\text{PSU}}^{\text{R}}$  and  $\hat{\Pi}_{\text{PSU}}^{\text{S}}$  respectively, have not been optimized by using Cuckoo hashing yet, and they are only efficient for datasets with small size.

### A.1 $\hat{\Pi}_{\text{PSU}}^{\text{R}}$ : Simplified version of $\Pi_{\text{PSU}}^{\text{R}}$



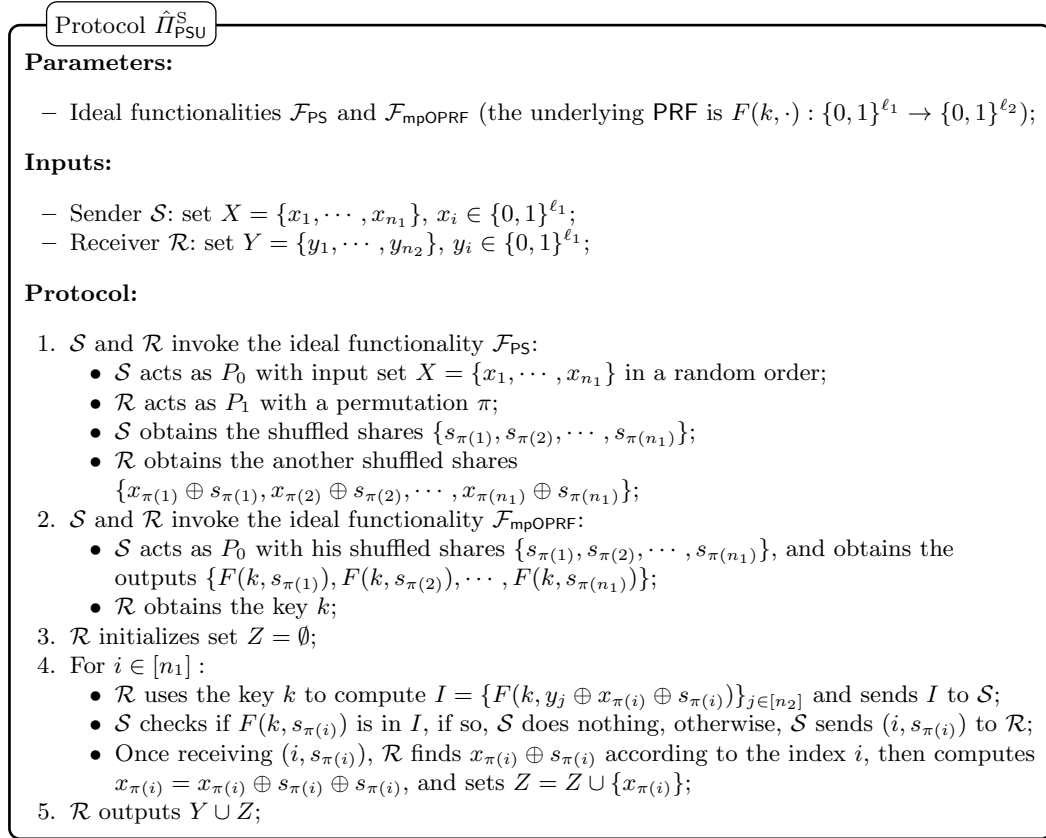
**Fig. 15.**  $\hat{\Pi}_{\text{PSU}}^{\text{R}}$ : Simplified version of  $\Pi_{\text{PSU}}^{\text{R}}$ .

The specification of  $\hat{\Pi}_{\text{PSU}}^{\text{R}}$  is shown in Figure 15. We assume that the sender's set is  $X = \{x_1, \dots, x_{n_1}\}$  and the receiver's set is  $Y = \{y_1, \dots, y_{n_2}\}$ . We first shuffle the receiver's set  $Y$  while generating their shares, by a permutation  $\pi$  only known to the sender. By invoking **Permute + Share**, the sender can obtain the share set  $\{s_{\pi(1)}, s_{\pi(2)}, \dots, s_{\pi(n_2)}\}$ , and the receiver obtains the other share set  $\{s'_{\pi(1)}, s'_{\pi(2)}, \dots, s'_{\pi(n_2)}\} = \{y_{\pi(1)} \oplus s_{\pi(1)}, y_{\pi(2)} \oplus s_{\pi(2)}, \dots, y_{\pi(n_2)} \oplus s_{\pi(n_2)}\}$ . After executing the **Permute + Share**, the two parties proceed to perform multi-point OPRF, through which the receiver obtains  $\{F(k, s'_{\pi(1)}), \dots, F(k, s'_{\pi(n_2)})\}$  and the sender obtains the OPRF key  $k$ . Then, for each  $x \in X$ , the sender computes



$I_x = \{F(k, x \oplus s_{\pi(1)}), \dots, F(k, x \oplus s_{\pi(n_2)})\}$  and sends it to receiver. Then, if there is no  $F(k, s'_{\pi(i)})$  belonging to  $I_x$ , the receiver obtains  $x$  through  $\mathcal{F}_{\text{OT}}$ . At last, the receiver outputs the items from the sender, as well as  $Y$ .

## A.2 $\hat{\Pi}_{\text{PSU}}^{\text{S}}$ : Simplified version of $\Pi_{\text{PSU}}^{\text{S}}$




**Fig. 16.**  $\hat{\Pi}_{\text{PSU}}^{\text{S}}$ : Simplified version of  $\Pi_{\text{PSU}}^{\text{S}}$ .

The specification of  $\hat{\Pi}_{\text{PSU}}^{\text{S}}$  is shown in Figure 16. Similarly, we assume that the sender's set is  $X = \{x_1, \dots, x_{n_1}\}$  and the receiver's set is  $Y = \{y_1, \dots, y_{n_2}\}$ .  $\Pi_{\text{PSU}}^{\text{S}}$  first executes **Permute + Share** on  $X = \{x_1, \dots, x_{n_1}\}$ . After this, the sender obtains the shares  $\{s_{\pi(1)}, \dots, s_{\pi(n_1)}\}$  of his items, while the receiver obtains the other shares  $\{s'_{\pi(1)}, \dots, s'_{\pi(n_2)}\}$  and the permutation  $\pi$ . Note that  $\pi$  is hidden from the sender, so he cannot link his shares with his original items. Then for each item of the sender (e.g., the  $\pi(i)$ -th item), the receiver uses  $s'_{\pi(i)}$  and the OPRF key  $k$  to generate  $I_{\pi(i)} = \{F(k, y_1 \oplus s'_{\pi(i)}), \dots, F(k, y_{n_2} \oplus s'_{\pi(i)})\}$ . Once receiving  $I_{\pi(i)}$  for the  $\pi(i)$ -th item, the sender checks if  $F(k, s_{\pi(i)}) \in I_{\pi(i)}$ , if not, the sender will send the share  $s_{\pi(i)}$  to the receiver. Then, the receiver can obtain  $x_{\pi(i)} = s_{\pi(i)} \oplus s'_{\pi(i)}$ . Finally, the receiver outputs the items from the sender, as well as set  $Y$ .

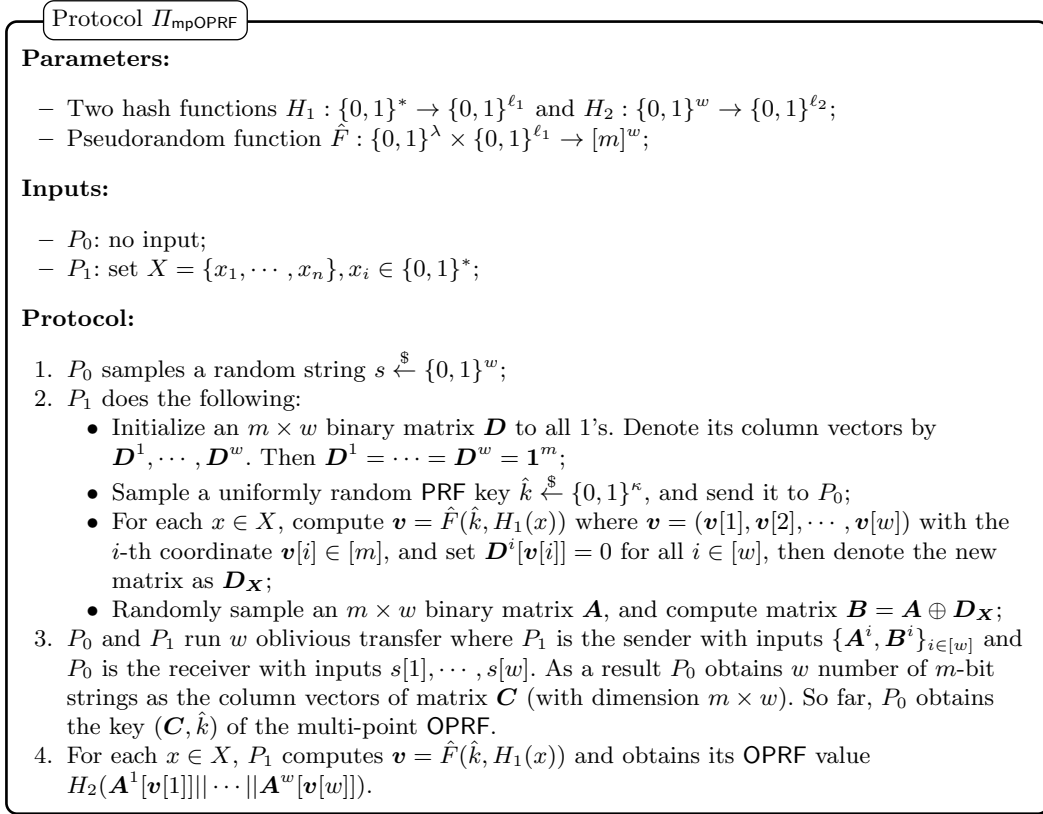
## B Additional Materials for Section 4: Constructions for the building blocks

In this section, we introduce the specific protocols  $\Pi_{\text{PS}}$  and  $\Pi_{\text{mpOPRF}}$  that securely realize  $\mathcal{F}_{\text{PS}}$  and  $\mathcal{F}_{\text{mpOPRF}}$ , respectively. As for  $\mathcal{F}_{\text{OT}}$ , we use the OT extension proposed by [15] that is commonly used in multi-party protocols and use the OT protocol in [23] as the base OT.



**Fig. 17.** The Permute + Share Protocol in [21].

*The Permute + Share protocol.* We use the protocol in [21] to realize  $\mathcal{F}_{\text{PS}}$ , and give the details in Figure 17. Generally speaking, the protocol leverages a switching network to realize the permutation, and the random labels of wires are used to form secret shares. A switching network consists of  $q$  switches and  $2q + n$  wires where  $n$  is the number of items to be permuted. The party  $P_0$  who inputs the permutation  $\pi$  will transfer  $\pi$  into a selection bit set  $S$ , in which each item is used to control a switch. The other party  $P_1$  randomly chooses the label of each wire, and uses the labels of input wires to mask the input set  $X = \{x_1, \dots, x_n\}$ . The masked values are then sent to  $P_0$  and taken as the switching network's input. Then, in topological order, the two parties jointly compute an atomic swap on each switch. And each atomic swap is implemented by using oblivious transfer according to the corresponding selection bit in the set  $S$ . After this, set  $X$  is permuted to  $\pi(X)$ , and each share is re-randomized by the labels in the path. At last,  $P_0$  obtains the blinded values for all the output wires as the share set, and  $P_1$  uses the labels of output wires as the other share set.



**Fig. 18.** The Multi-Point OPRF Protocol in [6].

*The multi-point OPRF protocol.* Chase and Miao in [6] proposed a new construction for multi-point OPRF whose functionality is shown in Figure 10, and they leveraged the new construction to realize a lightweight private set intersection (PSI) protocol. In essence, the construction in [6] is an extension of BaRK-OPRF in [17], and Chase and Miao have pointed it out in their paper. To be self-contained, we rewrite their construction in Figure 18. The key of the multi-point OPRF is  $(\mathbf{C}, \hat{k})$  and the pseudorandom function is as follows:

$$\begin{aligned} \mathbf{v} &= \hat{F}(\hat{k}, H_1(x_i)) \\ F((\mathbf{C}, \hat{k}), x_i) &= H_2(\mathbf{C}^1[\mathbf{v}[1]] \parallel \dots \parallel \mathbf{C}^w[\mathbf{v}[w]]) \end{aligned}$$

Generally speaking, the two parties both know a pseudorandom function  $\hat{F}$  with key  $\hat{k}$  that maps a  $\ell_1$ -bit item into a vector  $\mathbf{v} \in [m]^w$ . Firstly,  $P_1$  prepares two  $m \times w$  binary matrices  $\mathbf{A}$  and  $\mathbf{B}$ . More concretely, all the items in  $X$  will be mapped to some positions in a  $m \times w$  matrix (i.e.,  $\mathbf{D}_X$  in Figure 18) by  $\hat{F}$ .  $\mathbf{A}$  is randomly chosen, and in  $\mathbf{B}$ , the 1-bit elements that are located in these position are equal to the corresponding values in  $\mathbf{A}$  while the other elements will be different.  $P_0$  picks a random string  $s \in \{0, 1\}^w$ , then  $P_0$  and  $P_1$  perform  $w$  OTs. For the  $i$ -th OT,  $P_0$  takes  $s[i]$  as input, and  $P_1$  takes  $\mathbf{A}^i$  and  $\mathbf{B}^i$  as input, then  $P_0$  will obtain output  $\mathbf{A}^i$  or  $\mathbf{B}^i$  according to  $s[i]$ . After OTs,  $P_0$  will obtain  $w$  column vectors, which will form matrix  $\mathbf{C}$ .  $P_1$  obtains PRFs of her items using  $\mathbf{A}$  and  $\hat{k}$ . Note that for  $\forall x \in X$ ,  $F((\mathbf{C}, \hat{k}), x) = F((\mathbf{A}, \hat{k}), x)$ . The protocol needs  $O(n)$  communication and computation cost, and only involves cheap symmetric-key and bitwise operations.

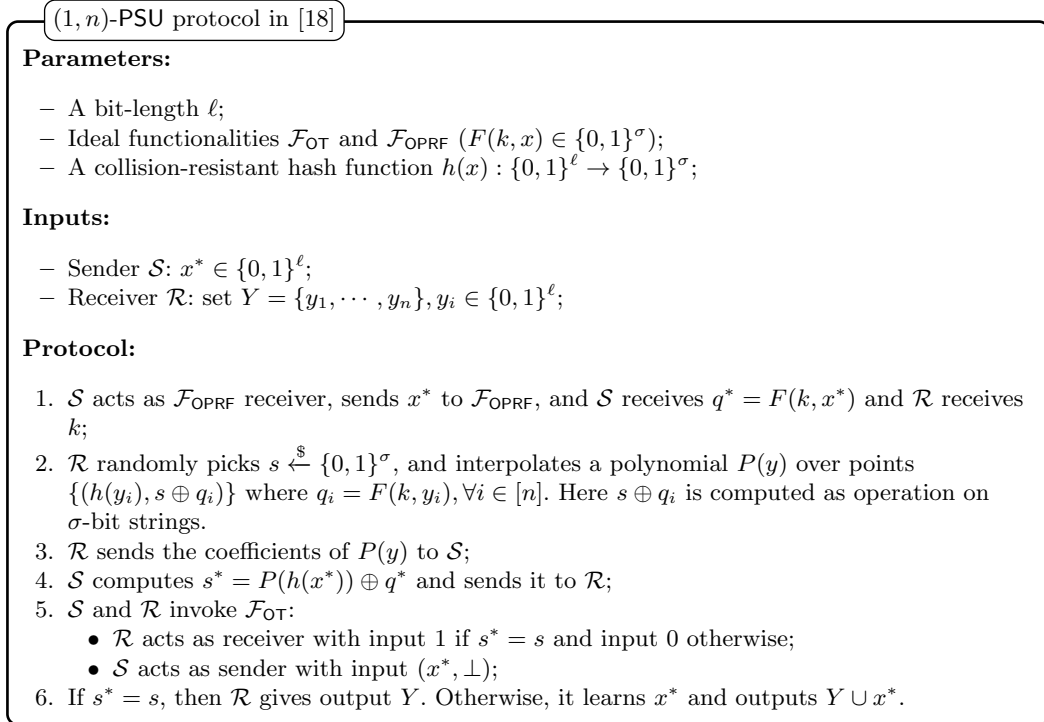
## C Additional Materials for Section 5

In this section, we first describe the protocol in [18], and then further explain why the protocol in [18] and earlier constructions [9,7] cannot benefit from Cuckoo hashing.

### C.1 Details of the protocol in [18]

According to the design idea in [18], we first give the protocol for  $(1, n)$ -PSU where the sender only holds an item  $x^*$  and the receiver holds a set  $Y = \{y_1, \dots, y_n\}$  in Figure 19. More specifically, their protocol works in the following way: the two parties first execute an OPRF sub-protocol for  $F$ , then the sender obtains  $F(k, x)$  without knowing  $k$  and the receiver obtains the PRF key  $k$ . After that, the receiver interpolates a polynomial  $P$  over points  $\{(y, s \oplus F(k, y))\}_{y \in Y}$ , where  $s$  is a random value chosen by the receiver, and sends  $P$  to the sender. Once receiving the polynomial  $P$ , the sender calculates  $s' = P(x) \oplus F(k, x)$  and sends  $s'$  to the receiver. Then the receiver checks if  $s' = s$ . If not, meaning that  $x \notin Y$ , the receiver obtains  $x$  through OT, otherwise obtains a dummy item.

Kolesnikov et al. [18] extend  $(1, n)$ -PSU to the general case (namely,  $(n_1, n_2)$ -PSU) by repeatedly using  $(1, n)$ -PSU. Then in Figure 20, we give the protocol for  $(n_1, n_2)$ -PSU where the sender and receiver hold  $X = \{x_1, \dots, x_{n_1}\}$  and  $Y = \{y_1, \dots, y_{n_2}\}$ , respectively. Note that for each sender's item, both the secret value  $s$  and the key  $k$  for OPRF need to be *refreshed*, otherwise the sender can learn information about the intersection.



**Fig. 19.**  $(1, n)$ -PSU protocol in [18].

## C.2 Discussion about Cuckoo hashing

As for Cuckoo hashing, Kolesnikov et al. [18] pointed out that "this hashing scheme (and the corresponding performance improvement) does not immediately fit in the PSU case." Recall that the protocols [9,7,18] share the same design framework as in Figure 1. We can see that the sender's set can be inserted into Cuckoo hash table. Then, for the item in each bin of Cuckoo hash table, the receiver will check if it belongs to the intersection, if not, the receiver will get the item. Note that the receiver also knows the item's position in Cuckoo hash table. Since the position of an item in Cuckoo hash table is also affected by other items, the receiver can obtain partial information about the sender's entire input set based on the received item and its position in Cuckoo hash table.

$(n_1, n_2)$ -PSU protocol in [18]

### Parameters:

- A bit-length  $\ell$  and  $n = \max(n_1, n_2)$ ;
- Number of bins  $\beta = \beta(n)$ , hash function  $H : \{0, 1\}^\ell \rightarrow [\beta]$ , and max bin size  $m$ ;
- A special item  $\perp \in \{0, 1\}^*$ ;

### Inputs:

- Sender  $\mathcal{S}$ :  $X = \{x_1, \dots, x_{n_1}\}, x_i \in \{0, 1\}^\ell$
- Receiver  $\mathcal{R}$ : set  $Y = \{y_1, \dots, y_{n_2}\}, y_i \in \{0, 1\}^\ell$ ;

### Protocol:

1.  $\mathcal{S}$  and  $\mathcal{R}$  hash items of their sets  $X$  and  $Y$  into  $\beta$  bins under hash function  $H$ . Let  $B_{\mathcal{S}}[i]$  and  $B_{\mathcal{R}}[i]$  denote the set of items in the sender's and receiver's  $i$ -th bin, respectively;
2.  $\mathcal{S}$  pads each bin  $B_{\mathcal{S}}[i]$  with the special item  $\perp$  up to the maximum bin size  $m + 1$ , and randomly permutes all items in this bin;
3.  $\mathcal{R}$  pads each bin  $B_{\mathcal{R}}[i]$  with one special item  $\perp$  and different dummy items to the maximum bin size  $m + 1$ ;
4.  $\mathcal{R}$  initializes set  $Z = \emptyset$ ;
5. For each bin  $i \in [\beta]$ , for each item  $x_j \in B_{\mathcal{S}}[i]$ :
  - $\mathcal{S}$  and  $\mathcal{R}$  invoke the  $(1, n)$ -PSU sub-protocol in Figure 19 and  $n = m + 1$ :
    - \*  $\mathcal{S}$  acts as sender with input  $x_j$ ;
    - \*  $\mathcal{R}$  acts as receiver with input set  $B_{\mathcal{R}}[i]$ ;
    - \*  $\mathcal{R}$  obtains output  $Z_{i,j}$  and sets  $Z = Z \cup Z_{i,j}$ ;
6.  $\mathcal{R}$  outputs  $Z$ .

**Fig. 20.**  $(n_1, n_2)$ -PSU protocol in [18].