

Can you find the one for me?

Privacy-Preserving Matchmaking via Threshold PSI

Yongjun Zhao and Sherman S.M. Chow

Department of Information Engineering
The Chinese University of Hong Kong, Hong Kong
{zy113, sherman}@ie.cuhk.edu.hk

Abstract. Private set-intersection (PSI) allows a client to only learn the intersection between his/her set C and the set S of another party, while this latter party learns nothing. We aim to enhance PSI in different dimensions, motivated by the use cases of increasingly popular online matchmaking — Meeting “the one” who possesses all desired qualities and free from any undesirable attributes may be a bit idealistic. Meanwhile, the criteria should be expressed in a succinct form. In this paper, we realize *over-* (resp. *below-*) threshold PSI, such that the client learns the intersection (or other auxiliary private data) only when $|C \cap S| > t$ (resp. $\leq t$). The threshold corresponds to tunable criteria for (mis-)matching, without marking all possible attributes as desired or not. To the best of our knowledge, our constructions are the very first solution for these two open problems posed by Bradley *et al.* (SCN '16) and Zhao and Chow (PoPETS '17), without resorting to the asymptotically less efficient generic approach from garbled circuits.

Moreover, we consider an “outsourced” setting with a service provider coordinating the PSI execution, instead of having two strangers to be online simultaneously for executing a highly-interactive PSI directly with each other. Outsourcing our two protocols are arguably optimal, namely, the two users perform $O(|C|)$ and $O(1)$ decryptions, for unlocking the private set C and the outcome whether a match has been found.

1 Introduction

In this big-data era, information is an asset. Sharing of information often leads to a win-win situation. The key issue is how to share selectively and strategically. People nowadays tend to share their information over the online social network (OSN). Usually, the sharing decision is based on whether the “subscribing user” has been admitted into a certain “circle” or not. The admission decision can be easy to make if we can rely on real-world friendship. Yet, people often reach out and expand their networks to enjoy the real benefits brought by OSN. That will be desirable if this decision can be made in a more systematic and intelligent way, *e.g.*, if a user possesses a sufficient number of common interests/attributes.

Private set intersection (PSI) is a handy cryptographic primitive which allows two parties P_1 and P_2 , traditionally referred to as a client and a server,

to compute the intersection of their respective private sets. For example, two companies can learn who are their common customers without sharing their databases. Yet, apart from hiding elements which are absent from the set of the counter-party, PSI offers no more control. For example, if two companies do not share a high number of common customers, they may not bother to discuss any joint campaign, not to say revealing any common customers to each other.

Recently, Zhao and Chow [39] initiated the study of PSI with access structure, *i.e.*, the client gets to know the intersection set only if its private set satisfies some policy specified by the server. As a special case, they consider threshold PSI which only reveals the intersection if its size is greater than a threshold agreed by both parties. While the vision of incorporating a sharing strategy to a vanilla PSI is great, the actual protocols realized by Zhao and Chow are a bit unsatisfactory. Without resorting to indistinguishability obfuscation or anonymous ABE, their design *always leak* how many elements “contribute” to the satisfaction of the policy. Specifically, for threshold policy, they only achieved a weaker variant under the name of *threshold private set-intersection cardinality* (t -PSI-CA), which realizes the following functionality: if the number of common elements from the two sets is less than an agreed threshold t , the size of the intersection will be revealed; otherwise, the intersection will be revealed. They argue that this level of privacy protection is enough for applications such as online dating, which the matching criteria are sensitive, and revealing the degree of overlapping even in a mis-match is a nice feature. Nevertheless, threshold PSI protocol based on standard cryptographic assumptions remains open [39].

From the perspective of finding common interests, the “over threshold” policy discussed above appears to be a natural choice. Yet, considering matchmaking or access control in general, it is equally interesting to realize the complementary notion of *below-threshold private set-intersection*. To unify the notion, we rename the two functionalities above as t^{\geq} -PSI [39] and t^{\leq} -PSI [7] respectively.

The benefits of supporting both kinds of policy are apparent. For the application of online dating, using t^{\geq} -PSI alone only allows the search for desired quality. Users probably also want to match with others who do not possess a certain set of undesired attributes (*e.g.*, smoking, over-gregarious). With t^{\leq} -PSI, users can simply make sure that their number of occurrence is below an acceptable threshold instead of specifying every possible negated attributes (*e.g.*, non-smoking). Unfortunately, t^{\leq} -PSI is also posed as an open question [7].

1.1 Technical Overview

Despite the conceptual similarity, it is fair to say that t^{\leq} -PSI and t^{\geq} -PSI are two *different* problems. We have a weakened version of t^{\geq} -PSI (namely t -PSI-CA) [39], but the corresponding weakened form for t^{\leq} -PSI does not exist in the literature [7]. To the best of our knowledge, no one has been able to give any solution for these two problems so far (perhaps except the generic approach of using garbled circuit [28]). The fundamental issue is that, t^{\geq} -PSI (resp. t^{\leq} -PSI) fall within the framework of private set-intersection with *monotone* (resp. *non-*

monotone) access structure [39]. It is not clear how to construct non-monotone access structure from monotone one (and vice versa).

In this work, we unify the design of both protocols. We take an innovative approach to realize both kinds of threshold PSI protocol, without the deficiency of leaking the intersection size. To better understand the difficulty of hiding the intersection size, we briefly go over the design of the protocol of Zhao and Chow [39]. Roughly, it works by generating “secret shares” to the participant. With enough shares, the intersection set can be “unlocked”. The difficulty faced by Zhao and Chow in hiding the size of the intersection appears to be the following. On one hand, there should be a way to quickly identify what shares can be used to reconstruct the unlocking key; for otherwise one needs to exhaust an exponential number of possible combinations among the shares. On the other hand, revealing whether an individual share is potentially useful or not is related to and hence reveals the size of the intersection.

Interestingly, we got inspiration from an apparently even more restrictive variant of PSI proposed by Carpent *et al.* [10], which is known as existential PSI (or PSI-X in short). PSI-X only outputs a single bit instead of a set. The output denotes whether the two private sets have any overlapping. As minimal as it may seem, we “upgrade” our own design of PSI-X protocol (Π_X , another contribution discussed in Appendix C) to encode more information. Specifically, we build a protocol $\Pi_{\text{ePSI-CA}}$ which we call *encrypted private set-intersection cardinality* (ePSI-CA). In this design, cardinality is no longer an unintended leakage, but intentionally encrypted to provide a basis useful for realizing the threshold functionality.

With ePSI-CA, we obtain a t^{\geq} -PSI protocol by a simple modification of our t^{\leq} -PSI protocol $\Pi_{t^{\leq}\text{-PSI}}$. Underlying both designs is a technique for computation over encrypted data realized by *oblivious polynomial evaluation* (OPE) [22].

1.2 Merits of Our Constructions

Our proposed protocols ($\Pi_{\text{ePSI-CA}}$, $\Pi_{t^{\leq}\text{-PSI}}$, $\Pi_{t^{\geq}\text{-PSI}}$, Π_X , *etc.*) are of both practical and theoretical interest. From the efficiency perspective, the complexities of our constructions are linear in the set size n . This beats the state-of-the-art generic approach based on garbled circuit with $O(n \log n)$ complexity [28].

From the design perspective, we demonstrate how to use old techniques in the PSI literature in a novel way to realize PSI functionalities which no efficient solutions are known. Specifically, many PSI protocols are built by using OPE to encode the private set [14, 22, 25, 26, 33, 39]. Another idea which uses Bloom filter related technique to realize PSI was first brought by Dong *et al.* [20] only recently, and has a shorter history [17, 18, 39]. To the best of our knowledge, for the first time we combine these two techniques in a non-trivial way. It also suggests a new avenue for addressing other open problems inside or even outside the field of PSI. (See Appendix C for more details.)

Lastly, our protocols remain conceptually simple and modular. Both desideratum greatly simplify the security analysis. Any more efficient $\mathcal{F}_{\text{ePSI-CA}}$ instantiation in the future immediately results in more efficient $\Pi_{t^{\leq}\text{-PSI}}$, $\Pi_{t^{\geq}\text{-PSI}}$ protocols.

1.3 Outsourced Threshold PSI Protocols

Although our protocols are asymptotically efficient, they still rely on public-key techniques to ensure security, which are difficult to avoid and are not as efficient as symmetric-key primitives (say, hash functions or blockciphers) which often fail to provide algebraic structure for fancy functionalities. PSI protocols are also often highly interactive. For our motivating applications of matchmaking, the interactive nature and the heavy use of public key cryptography hinder the practical usage of PSI.

Our final contribution lies in outsourcing the heavy computations in our protocols to an oblivious cloud. Beyond simply following the increasingly popular trend of leveraging cloud service, we believe that most of the popular mobile applications nowadays, no matter privacy-preserving or not, are often executed with the help of some central servers operated by the service provider. As such, outsourcing PSI not only leads to better efficiency, but also better matches the business model and the usage habits of mobile applications (where the user who is a potential match may respond to notification of the smartphone from time to time but not permanently staying online).

While there is a server which mediates requests between clients, the privacy guarantees of PSI still carry over. In other words, with our outsourced extensions, not only we can enjoy the richer functionalities on top of the privacy provided by our PSI protocols, but also a more deployable protocol which is closer to the real-world model from the perspective of both users and the service provider.

1.4 Related Work

Freedman *et al.* [22] first proposed a PSI protocol based on oblivious polynomial evaluation. Dong *et al.* [20] initiated the pursue of PSI protocol using oblivious transfer extension. Subsequently, more efficient PSI protocols using similar technique are proposed [34, 35, 37, 38].

A branch of work aims to restrict the output or the leakage of PSI. PSI-CA reveals only the cardinality of the intersection but not the set itself [16, 17, 19, 21, 22, 27, 33]. Some of them also uses Bloom filter technique [16, 19, 21], but none of them can be directly adapted to $\Pi_{\text{ePSI-CA}}$ (see discussion in Sec. 4). Ateniese *et al.* [3] and D’Arco *et al.* [15] proposed (input-)size-hiding PSI. Bradley *et al.* [7] further enables imposing an upper bound on the input set size. Another possible way is to use authorized PSI [8, 12, 13] with a trusted party.

Recently researchers also consider PSI in the outsourced setting [1, 30–32]. They do not support the advanced threshold set operations considered here.

Concurrent Work.¹ Recently, Hallgren *et al.* [24] also study over-threshold private set-intersection with application in ride sharing. They also consider security in the semi-honest model, but the complexity of their protocol is of or-

¹ The first public presentation of this work was delivered on December 5, 2017 in the rump session of AsiaCrypt 2017. We have submitted this work to some conference before our first upload to IACR Cryptology ePrint Archive, which is made on February 14, 2018.

der $O(n^2)$, which is worse than generic garbled circuit approach. A very recent manuscript by Ciampi and Orlandi [11] studies how to perform secure post-processing of the output of PSI protocol in the semi-honest model. Combining their solution with two party computation techniques allows computation of $f(C \cap S)$ for arbitrary function f . Our protocols are specific protocols which are more efficient ($O(n)$ vs. $O(n^2)$). Also, our design naturally allows revelation of some auxiliary data (*e.g.*, a session key K encrypting the contact information for matchmaking) when the threshold policy is satisfied. Extending their two party computation approach [11] in a straightforward way to also support this feature requires a more complicated function f which embeds an encryption of the auxiliary data since its basic usage is for deriving $f(C \cap S)$.

2 Preliminary

For a finite set S , $|S|$ denotes its size and $s \xleftarrow{\$} S$ denotes picking an element uniformly at random from S . We denote $[i] = \{1, \dots, i\}$. We write $\{s_i\}_n$ as a shorthand for the set $S = \{s_1, \dots, s_n\}$ of n elements. We drop the subscript n if it is clear from context. We use \mathcal{F}_ρ to denote an ideal functionality that implements the protocol ρ , and use Π_ρ to denote a concrete construction of the protocol ρ .

2.1 Homomorphic Encryption and Oblivious Polynomial Evaluation

We will use CPA-secure additive homomorphic encryption (AHE) (KeyGen, Enc, Dec) such as Paillier encryption [36]. Given two ciphertexts $\text{Enc}(\text{pk}, m_0)$ and $\text{Enc}(\text{pk}, m_1)$, one can efficiently compute their addition $\text{Enc}(\text{pk}, m_0 + m_1)$ without using private key sk . As a corollary, given one ciphertext $\text{Enc}(\text{pk}, m)$ and a constant c , one can perform repeated addition and obtain $\text{Enc}(\text{pk}, c \cdot m)$. With addition and constant-multiplication, we can build an OPE protocol. A polynomial $p(x)$ can be hidden by encrypting its coefficients a_0, \dots, a_k . With these encrypted polynomial, anyone holding a plaintext s can then compute an encryption of $p(s)$.

2.2 Bloom Filters

A Bloom filter [4] is a compact array of m bits that represents a set S of n elements for efficient set membership testing. It consists of a set of k independent hash functions $H = (h_1, \dots, h_k)$, h_i uniformly maps elements to index in $[m]$.

All bits in the array are initialized to 0. To insert an element $x \in S$, x is hashed by the k hash functions to get k index numbers. All the bits at these indexes in the array are set to 1, regardless of its original value. To check if an item y is in S , y is hashed by the k hash functions to get k indexes. If any of the bits at these indexes is 0, we conclude that y is certainly not in S (no false negative). Otherwise, y is probably in S . So there is only a small fraction of false positives. The upper bound of the false positive probability [6]

$$\text{is: } \epsilon = p^k \left(1 + O \left(\frac{k}{p} \sqrt{\frac{\ln m - k \ln p}{m}} \right) \right) \text{ where } p = 1 - \left(1 - \frac{1}{m} \right)^{kn}.$$

If we set the false positive rate to be less than a threshold ϵ_0 , it can be shown that the length of the bit array size m should be at least $m \geq n \log_2 e \cdot \log_2 1/\epsilon_0$, where e is the base of the natural logarithm. Equality is achieved when $k = (m/n) \cdot \ln 2 = \log_2 1/\epsilon_0$. We will stick with these optimal values as follows: the false positive probability is $\epsilon = 2^{-\omega(\log \lambda)}$ so that ϵ is negligible in the security parameter λ . As a result, $k = \omega(\log \lambda)$ and $m = k \cdot n \log_2 e$.

3 PSI with Threshold Policy

3.1 Definitions

We begin with the formal definitions of two private set-intersection with threshold policy functionalities in the literature.

Definition 1 (Below-Threshold Private Set-Intersection (t^{\leq} -PSI) [7]). Let S and C be subsets of a predetermined domain, the functionality $\mathcal{F}_{t^{\leq}\text{-PSI}}$ is:

$$((C, |S|), (S, |C|)) \mapsto \begin{cases} (C \cap S, \perp) & \text{if } |C \cap S| \leq t \\ (\perp, \perp) & \text{otherwise} \end{cases}$$

Definition 2 (Over-Threshold Private Set-Intersection (t^{\geq} -PSI) [24, 39]). Let S and C be subsets of a predetermined domain, the functionality $\mathcal{F}_{t^{\geq}\text{-PSI}}$ is:

$$((C, |S|), (S, |C|)) \mapsto \begin{cases} (C \cap S, \perp) & \text{if } |C \cap S| \geq t \\ (\perp, \perp) & \text{otherwise} \end{cases}$$

To realize both $\mathcal{F}_{t^{\leq}\text{-PSI}}$ and $\mathcal{F}_{t^{\geq}\text{-PSI}}$, our key insight is to leverage a primitive call *encrypted private set-intersection cardinality* (ePSI-CA) functionality in a novel way. This new functionality is formally defined below.

Definition 3 (Encrypted Private Set-Intersection Cardinality (ePSI-CA)). Let S and C be subsets of a predetermined domain. Let $(\text{pk}_1, \text{sk}_1)$ be a public / secret key pair of a homomorphic encryption scheme. The functionality $\mathcal{F}_{\text{ePSI-CA}}$ is:

$$((C, |S|, \text{pk}_1, \text{sk}_1), (S, |C|, \text{pk}_1)) \mapsto (\perp, \text{Enc}(\text{pk}_1, |C \cap S|))$$

We choose to single out $\mathcal{F}_{\text{ePSI-CA}}$ not only for a more compact presentation of $\Pi_{t^{\leq}\text{-PSI}}$ and $\Pi_{t^{\geq}\text{-PSI}}$, but also because we believe that $\mathcal{F}_{\text{ePSI-CA}}$ itself is an interesting primitive. In Appendix C, we use it to construct the most efficient *existential private set-intersection protocol* (PSI-X) and *private projection* [10] to date. Motivations and applications of these two protocols can be found in Appendix C.

3.2 Intuition

We will describe our $\Pi_{t^{\leq}\text{-PSI}}$ and $\Pi_{t^{\geq}\text{-PSI}}$ constructions in the $(\mathcal{F}_{\text{ePSI-CA}}, \mathcal{F}_{\text{PSI}})$ -hybrid model, while the concrete instantiation for $\mathcal{F}_{\text{ePSI-CA}}$ is deferred to Sec. 4.

In what follows, we use t^{\leq} -PSI as an example to show case how it is readily achievable with $\mathcal{F}_{\text{ePSI-CA}}$ as a building block. The crux of our novel combination of ePSI-CA and oblivious polynomial evaluation (OPE) technique is that it transforms $\text{Enc}(\text{pk}_1, |C \cap S|)$ (namely, the output of $\mathcal{F}_{\text{ePSI-CA}}$) to an encryption of a session key K if and only if $|C \cap S| \in [0, t]$. If $|C \cap S| \notin [0, t]$, the evaluation result will be random and contains no information about K .

In more details, P_2 re-randomizes the encrypted cardinality obtained from $\mathcal{F}_{\text{ePSI-CA}}$ by a random number r as $\text{Enc}(\text{pk}_1, |C \cap S| + r)$ such that P_1 knows nothing about $|C \cap S|$ but can put the randomized cardinality to the OPE. For OPE, P_2 first prepares a polynomial $p'(\cdot)$ whose roots are $r, r+1, \dots, r+t$, and chooses a random symmetric key K . P_2 then sends encrypted coefficients of polynomial $p''(\cdot) = r' \cdot p'(\cdot) + K$ under its own public key pk_2 . In this way, when P_1 obviously evaluates $p''(\cdot)$, the result will be an encryption of K if and only if $|C \cap S| \leq t$; otherwise it will be encrypting a random number that reveals no information about K (because $r' \cdot p'(\cdot)$ serves as a one-time pad encrypting K).

To retrieve the evaluation of $p''(\cdot)$ in plaintext, P_1 randomizes it in the same way as P_2 re-randomizes $|C \cap S|$, and asks P_2 for decryption. What P_1 eventually obtains is a value K' which equals to K if and only if $|C \cap S| \leq t$. This K' serves as a token showing if the intersection $|C \cap S|$ is below the threshold. The final step is to have P_1 and P_2 engage in a normal Π_{PSI} , in which P_1 and P_2 uses $C^{K'} = \{c_i \| K'\}$ and $S^K = \{s_i \| K\}$ as input respectively. Hence P_1 can recover the intersection if it possesses the same key $K' = K$.

3.3 t^{\leq} -PSI Protocol: $\Pi_{t^{\leq}\text{-PSI}}$

We describe our t^{\leq} -PSI protocol $\Pi_{\text{ePSI-CA}}$ in the $(\mathcal{F}_{\text{ePSI-CA}}, \mathcal{F}_{\text{PSI}})$ -hybrid model in Fig. 1, following to the intuition above. Some points to note are in order.

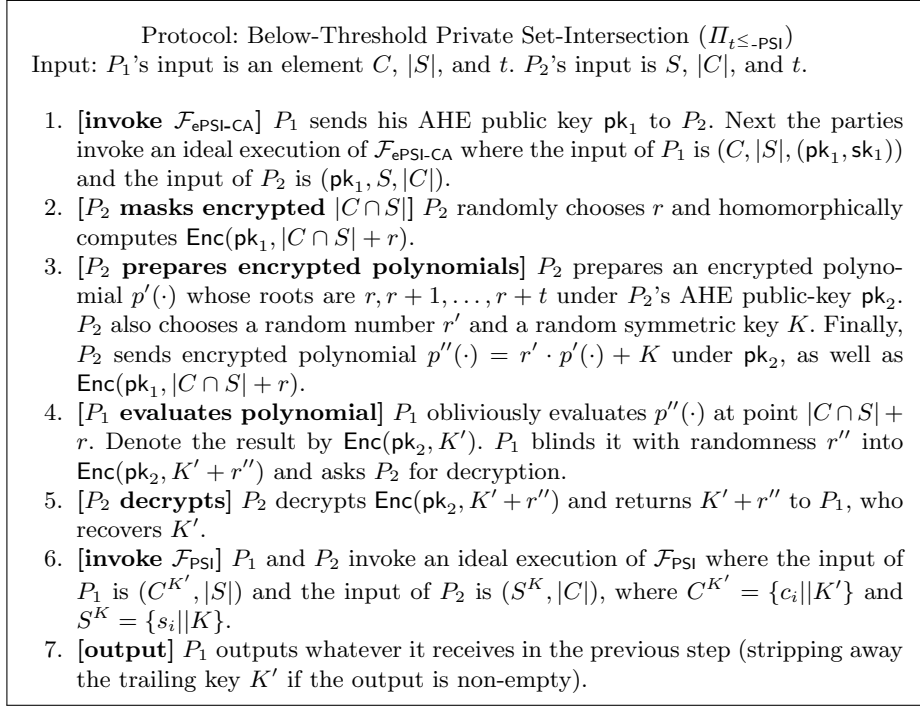
In Step (2), P_2 blinds the encrypted cardinality $\text{Enc}(\text{pk}_1, |C \cap S|)$ obtained from $\mathcal{F}_{\text{ePSI-CA}}$ by a uniformly random number r as $\text{Enc}(\text{pk}_1, |C \cap S| + r)$, which can be viewed as the result of double encryption: firstly encrypt $|C \cap S|$ using a one-time pad r , and then further encrypt under pk_1 .

In Step (3), the polynomial $p''(\cdot)$ and $p'(\cdot)$ are both degree $(t+1)$, so the number of coefficients to be encrypt and transmit is $(t+1) \in O(|C| + |S|)$. Also note that these coefficients are independent of the set S . Looking ahead, such independence allows outsourcing computation to a cloud server.

In Step (4), P_1 needs to blind the evaluation of polynomial by a one-time pad r'' similar to what P_2 does in Step (2). This is because if P_2 sees K' in plaintext, it can check if $K \stackrel{?}{=} K'$ and learn $|C \cap S| \stackrel{?}{>} t$, violating the requirement of $\mathcal{F}_{t^{\leq}\text{-PSI}}$.

3.4 Analysis

By the correctness of $\mathcal{F}_{\text{ePSI-CA}}$ functionality, P_2 obtains $\text{Enc}(\text{pk}_1, |C \cap S|)$ in Step 1. If the size of intersection $|C \cap S| \leq t$, then in Step 3 the polynomial $p'(\cdot)$ will be evaluated to 0, and hence the evaluation of $p''(\cdot)$ will be $K' = K$. On the

Fig. 1: Below-Threshold Private Set-Intersection ($\Pi_{t \leq \text{PSI}}$)

other hand, if $|C \cap S| > t$, the evaluation of $p''(\cdot)$ will be $K' \neq K$. Then by the correctness of \mathcal{F}_{PSI} , P_1 obtains $C \cap S$ if and only if $|C \cap S| \leq t$.

For efficiency, since public-key operations are much slower than symmetric-key ones, we only count the total number of public-key operations, including encryption, decryption, and homomorphic operations (addition and multiplication by a constant). We assume using $\Pi_{\text{ePSI-CA}}$ to instantiate $\mathcal{F}_{\text{ePSI-CA}}$ in Step 1. The number of public-key operations of $\Pi_{\text{ePSI-CA}}$ will be presented in Table 2. We also assume an efficient Π_{PSI} construction. Note that the state-of-the-art Π_{PSI} protocols [34, 37, 38] under the semi-honest model requires linear computation and communication complexities. Table 1 summarizes the overall result, showing that the $\Pi_{t \leq \text{PSI}}$ protocol features linear computational complexity.

Theorem 1. *Assuming the existence of CPA-secure additive homomorphic encryption scheme (KeyGen, Enc, Dec), whose plaintext space is super polynomial in the security parameter; then the protocol $\Pi_{t \leq \text{PSI}}$ in Fig. 1 securely implements the functionality $\mathcal{F}_{t \leq \text{PSI}}$ in Definition 1 in the presence of semi-honest adversaries under the $(\mathcal{F}_{\text{ePSI-CA}}, \mathcal{F}_{\text{PSI}})$ -hybrid model.*

Proof. **Simulating the view of P_1 using $\text{Sim}_1^{t \leq \text{PSI}}$.** The view of P_1 contains $\text{Enc}(\text{pk}_1, |C \cap S| + r)$, pk_2 , encryptions of coefficients of the polynomial $p''(\cdot)$

Table 1: Computational Efficiency of $\Pi_{t \leq \text{PSI}}$
 $(\mathcal{F}_{\text{ePSI-CA}}$ is instantiated by $\Pi_{\text{ePSI-CA}}$ with complexity in Table 2)

		Enc	Dec	addition	multiplication	Π_{PSI}
Step 1	(P_1)	$O(\log^2(\lambda)(C + S))$	0	$O(\log^2(\lambda)(C + S))$	0	0
	(P_2)	$O(\log^2(\lambda)(C + S))$	$O(C)$	$O(\log^2(\lambda)(C + S))$	$O(\log^2(\lambda) C)$	0
Step 2	(P_2)	1	0	1	0	0
Step 3	(P_2)	$t + 3$	0	0	0	0
Step 4	(P_1)	1	1	$t + 2$	$t + 1$	0
Step 5	(P_2)	0	1	0	0	0
Step 6	(P_1)	0	0	0	0	$O(C + S)$
	(P_2)	0	0	0	0	$O(C + S)$
Total		$O(\log^2(\lambda)(C + S))$	$O(C)$	$O(\log^2(\lambda)(C + S))$	$O(\log^2(\lambda) C)$	$O(C + S)$

(denoted by $\text{Enc}(\text{pk}_2, p''(\cdot))$), $K' + r''$ which are messages sent by P_2 , and \perp from $\mathcal{F}_{\text{ePSI-CA}}$, the output of \mathcal{F}_{PSI} .

Let \mathcal{A} be a probabilistic polynomial time (PPT) adversary corrupting party P_1 . We design a PPT simulator $\text{Sim}_1^{t \leq \text{PSI}}$ that invokes \mathcal{A} by playing the role of the honest party P_2 and it emulates the ideal functionalities $\mathcal{F}_{\text{ePSI-CA}}$, \mathcal{F}_{PSI} . The simulator will generate a view indistinguishable from a hybrid one. $\text{Sim}_1^{t \leq \text{PSI}}$ has different simulation strategies for different outputs of P_1 . We consider two disjoint cases.

(1) the output is \perp :

1. Given input $((C, |S|, t), \perp)$, $\text{Sim}_1^{t \leq \text{PSI}}$ invokes \mathcal{A} on input $(C, |S|, t)$, and receives \mathcal{A} 's first message pk_1 .
2. $\text{Sim}_1^{t \leq \text{PSI}}$ plays as the trusted party and emulates the ideal calls to $\mathcal{F}_{\text{ePSI-CA}}$.
3. $\text{Sim}_1^{t \leq \text{PSI}}$ generates a random public/private key pair $(\text{pk}_2, \text{sk}_2)$, a random symmetric key K , just as what P_2 will do, and continues the protocol emulation by sending encryptions of 0 under pk_2 (representing an all-zero polynomial) instead of encryption of coefficients of the polynomial $p'(\cdot)$. It also encrypts a random number R_1 under pk_1 to emulate the intended message $\text{Enc}(\text{pk}_1, |C \cap S| + r)$.
4. When \mathcal{A} obliviously evaluates the zero-polynomial at point R in Step 4, the result will be an encryption of 0 under pk_2 . \mathcal{A} randomizes this encryption of 0 into an encryption of r'' and asks $\text{Sim}_1^{t \leq \text{PSI}}$ for decryption. $\text{Sim}_1^{t \leq \text{PSI}}$ returns yet another random value R_2 .
5. Finally \mathcal{A} will compute $K' = R_2 - r''$ ($\neq K$ with overwhelming probability). Then \mathcal{A} uses $(C' = \{c_i || K'\}, |S|)$ as input to the \mathcal{F}_{PSI} functionality emulated by $\text{Sim}_1^{t \leq \text{PSI}}$, who will return \perp as intended.

We argue that this simulated view is indistinguishable from the real one. First, notice that the simulated messages $\text{Enc}(\text{pk}_1, R_1)$ and R_2 , are identically distributed as the real ones $\text{Enc}(\text{pk}_1, |C \cap S| + r)$ and $K' + r''$. It is because in the real protocol, r is selected by P_2 uniformly at random and K' is distributed

uniformly at random when $|C \cap S| > t$. Second, notice that the other simulated messages are encryptions under pk_2 . Any distinguisher of these two views can be transformed to an adversary breaking CPA security of the encryption scheme.

(2) the output is a subset $\hat{C} \subseteq C$ whose size $|\hat{C}|$ is less than t :

$\text{Sim}_1^{t \leq \text{-PSI}}$ works as the previous case, except in Step (4) it decrypts the ciphertext to get r'' . In Step (5) it computes $K' = R_2 - r''$, converts the set \hat{C} into $\hat{C}^{K'}$ by appending K' to each element in \hat{C} . Finally it uses $\hat{C}^{K'}$ to emulate the output of \mathcal{F}_{PSI} for \mathcal{A} .

This simulated view is indistinguishable from the real view, because everything is the same as the other case with the only exception being the output of \mathcal{F}_{PSI} .

Simulating the view of P_2 using $\text{Sim}_2^{t \leq \text{-PSI}}$.

This part is easy because P_2 's view contains only $\text{Enc}(\text{pk}_1, |C \cap S|)$ from $\mathcal{F}_{\text{ePSI-CA}}$, \perp from \mathcal{F}_{PSI} and $\text{Enc}(\text{pk}_2, K' + r'')$. The third element is an encryption of a truly random value (because r'' is chosen by P_1 uniformly at random), which can be perfectly simulated using $\text{Enc}(\text{pk}_2, R)$ where R is also chosen uniformly at random. The first element can be simulated by $\text{Enc}(\text{pk}_1, 0)$. A straightforward reduction shows that any distinguisher who can distinguish $\text{Enc}(\text{pk}_1, |C \cap S|)$ from $\text{Enc}(\text{pk}_1, 0)$ can be used to break the CPA-security of the encryption scheme.

3.5 $t \geq$ -PSI Construction and Generalizations

To construct $\Pi_{t \geq \text{-PSI}}$, we modify $\Pi_{t \leq \text{-PSI}}$ given in Fig. 1. The modification is simple and straightforward: in Step 3, P_2 prepares a polynomial whose roots are $r + t, r + (t + 1), \dots, r + \min(|S|, |C|)$, where the function $\min(x, y)$ returns the smaller value of x and y . The rest of the protocol remains exactly the same. Note that the degree of this polynomial is $\min(|S|, |C|) - t + 1$, which remains to be $O(|S| + |C|)$, and hence the efficiency analysis in Table 1 also holds for $\Pi_{t \geq \text{-PSI}}$.

In general, P_2 can specify the roots of the polynomial at will. For example, the roots could be integers within a certain interval $\{r + a, r + (a + 1), \dots, r + b\}$. It means P_1 only learns the intersection if $|C \cap S|$ falls within the range $[a, b]$. Further generalizing, we can change the criteria to be $|C \cap S| \in \{m_1, m_2, \dots, m_q\}$, *i.e.*, an arbitrary set of numbers instead of consecutive numbers.

Proving the security of $\Pi_{t \geq \text{-PSI}}$ (and its generalizations) just takes a very straightforward adaption of Theorem 1 and thus we omit the repetitive details.

4 Encrypted PSI-Cardinality

Both our $\Pi_{t \leq \text{-PSI}}$ and $\Pi_{t \geq \text{-PSI}}$ protocol heavily rely on $\mathcal{F}_{\text{ePSI-CA}}$, which is a novel variant of PSI and PSI-CA protocol in the literature. In this section, we describe how to instantiate $\mathcal{F}_{\text{ePSI-CA}}$ efficiently by combining the use of oblivious polynomial evaluation in the previous section with Bloom filter techniques.

Before detailing our construction, it should be noted that $\mathcal{F}_{\text{ePSI-CA}}$ is not a trivial extension of $\mathcal{F}_{\text{PSI-CA}}$ or \mathcal{F}_{PSI} . It is tempting to instantiate $\mathcal{F}_{\text{ePSI-CA}}$ from

$\mathcal{F}_{\text{PSI-CA}}$ generically as follows: P_1 and P_2 executes $\mathcal{F}_{\text{PSI-CA}}$, and P_1 encrypts the output $|C \cap S|$ under its public key pk_1 , forwards the ciphertext to P_2 . Unfortunately this approach does not satisfy the security requirement of $\mathcal{F}_{\text{ePSI-CA}}$, namely P_1 should output \perp instead of learning $|C \cap S|$. Moreover, all $\Pi_{\text{PSI-CA}}$ protocols that we are aware of proceed by transforming elements in the two parties' respective set via some one-way transformation, and let P_1 counts the number of common elements in the transformed domain. If one follows this paradigm, it seems impossible to hide the size of intersection $|C \cap S|$ from P_1 . After all, in $\mathcal{F}_{\text{PSI-CA}}$ and \mathcal{F}_{PSI} , it is P_1 who has non-trivial output while in $\mathcal{F}_{\text{ePSI-CA}}$, it is P_2 instead. Such inherent inconsistency suggests that new techniques are required to design ePSI-CA.

4.1 ePSI-CA Protocol: $\Pi_{\text{ePSI-CA}}$

Our construction is inspired by the *existential private set-intersection* (denoted by \mathcal{F}_X) protocol by Carpent *et al.* The core component of their protocol implements an encrypted private membership test protocol, a special case of $\mathcal{F}_{\text{ePSI-CA}}$ where P_1 's input set consists of a single element. Unfortunately their construction is too inefficient, making it unsuitable for our application. We significantly improve their construction by a novel combination of oblivious polynomial evaluation and Bloom filter. Interested readers are referred to Appendix C for a detailed discussion of \mathcal{F}_X . Here we only highlight our design principle.

Recall that a Bloom filter BF_S encoding a set S supports efficient membership test by hashing the test element x into k locations using k hash functions. If the value of BF_S at those locations are *all* "1", we conclude that $x \in S$. Namely, the predicate " $x \in S$ or not" is transformed to determining the number of "1"s.

$$P(x, BF_S) = \begin{cases} 1 & x \in S \\ 0 & x \notin S \end{cases} = \begin{cases} 1 & \text{exactly } k \text{ "1"s in } \{BF_S[h_1(x)], \dots, BF_S[h_k(x)]\} \\ 0 & \text{less than } k \text{ "1"s in } \{BF_S[h_1(x)], \dots, BF_S[h_k(x)]\} \end{cases}$$

Suppose that P_2 's Bloom filter BF_S is encrypted under its public-key pk_2 , then P_1 can obviously compute an encryption of the number of "1"s under pk_2 by adding the ciphertexts of those k locations for element x . Denote such number of "1"s as n_X , and its encryption under pk_2 as e_{n_X} . Our task becomes how to transform e_{n_X} to an encryption of 0 if $n_X \in [0, k-1]$; or to an encryption of 1 if $n_X = k$. This task is quite similar to what we have seen in Sec. 3, in which we transform an encryption of $|C \cap S|$ to an encryption of 0 if $|C \cap S| \in [0, t]$; or to an encryption of a non-zero number otherwise.

Fig. 2 gives the full details of $\Pi_{\text{ePSI-CA}}$ protocol. Some notes are in order.

In Step (1), $k = \omega(\log \lambda)$ and $N = \omega(\log \lambda)|S| \log_2 e$, which are optimal values for a false positive rate of ϵ that is negligible in λ (see Sec. 2.2).

In Step (3), the resulting ciphertext is encrypting an integer $\hat{n}_i \in \{r_i, \dots, r_i + k\}$ because the sum consists of a random r_i and k encrypted numbers in $\{0, 1\}$.

In Step (4), the coefficients of the polynomials $p_i(\cdot)$ are independent of the private set $|C|$, so P_1 can outsource this step to a cloud server.

In Step (5), each evaluation of polynomial $p_i(\hat{n}_i)$ equals 0 if $c_i \notin S$; or equals to a constant $k \times (k-1) \times \dots \times 1 = k!$ if $\hat{n}_i \in S$. Hence we multiply by a factor $(k!)^{-1}$ outside the summation to normalize the number.

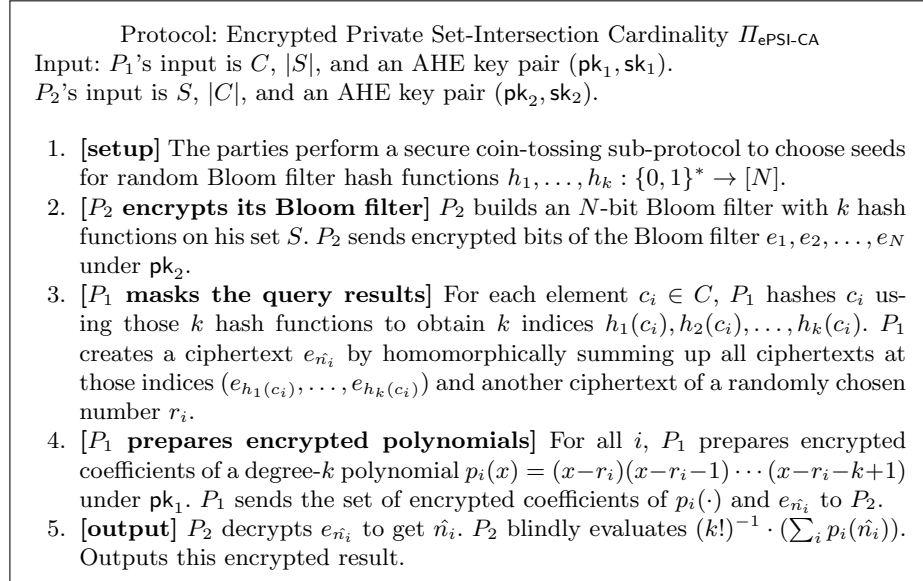


Fig. 2: Encrypted Private Set-Intersection Cardinality $\Pi_{\text{ePSI-CA}}$

4.2 Analysis

We first argue for the correctness. For any element $c_i \notin S$, P_1 can only collect less than k encryptions of “1” in Step 3 (otherwise a false positive of Bloom filter has occurred). Therefore, the polynomials $p_i(\hat{n}_i)$ will be evaluated to 0 in the ciphertext domain in Step 5. On the other hand, for $c_j \in S$, then the polynomial evaluation $p_j(\hat{n}_j) = k \times (k-1) \times \dots \times 1 = k!$. After normalizing by a factor of $(k!)^{-1}$, the output of the protocol will be an encryption of $|C \cap S|$ as intended.

Next we analyze the efficiency by counting the total number of public-key operations from Step (2) to Step (5). We assume that the false positive rate of Bloom filter is set to be $\epsilon = 2^{-\omega(\log \lambda)}$, so $k = \omega(\log \lambda)$. The other parameters of Bloom filter are set to the optimal values accordingly. Moreover, we assume that Horner’s rule is applied in evaluation of $p_i(\cdot)$ in Step 5, which requires k additions and k multiplications for a degree- k polynomial. Table 2 summarizes the result, which shows that the complexity of our construction is only linear in the set size. Note that the estimation in the last row is quite loose when we replace $k = \omega(\log \lambda)$ with $O(\log^2(\lambda))$. Nevertheless, we achieve computational and communicational complexity linear in the size of respective sets.

In terms of security, we have the following theorem:

Table 2: Computational Complexity of $\Pi_{\text{ePSI-CA}}$
 (false positive rate $\epsilon = 2^{-\omega(\log \lambda)}$, # of hash $k = \omega(\log \lambda)$)

	Enc	Dec	addition	multiplication
Step 2 (P_2)	$(\log_2 e)k S $	0	0	0
Step 3 (P_1)	$ C $	0	$k C $	0
Step 4 (P_1)	$(k+1) C $	0	0	0
Step 5 (P_2)	0	$ C $	$(k+1) C - 1$	$k C + 1$
P_1 total	$O(\log^2(\lambda)(C + S))$	0	$O(\log^2(\lambda)(C + S))$	0
P_2 total	$O(\log^2(\lambda)(C + S))$	$O(C)$	$O(\log^2(\lambda)(C + S))$	$O(\log^2(\lambda) C)$

Theorem 2. *Assuming the existence of a CPA-secure additive homomorphic encryption scheme (KeyGen, Enc, Dec), whose plaintext space is super-polynomial in the security parameter; then the protocol $\Pi_{\text{ePSI-CA}}$ in Fig. 2 securely implements the functionality $\mathcal{F}_{\text{ePSI-CA}}$ under the semi-honest model.*

Proof. We consider two corruption cases.

Simulating the view of P_1 using $\text{Sim}_1^{\text{ePSI-CA}}$. The view of P_1 only contains its view in the coin-tossing protocol $\text{View}_1^{\text{coin}}$, pk_2 , and (e_1, \dots, e_N) (encryptions of binary numbers under pk_2). $\text{Sim}_1^{\text{ePSI-CA}}$ can generate the first two using the $\text{Sim}_1^{\text{coin}}$, KeyGen algorithm, while the third one can be simulated by encryptions of zero due to the CPA-security of the encryption scheme. Assume for contradiction that there exists a distinguisher \mathcal{D} for the simulated view and the real view. One can build a distinguisher \mathcal{D}' breaking the CPA-security of the encryption scheme. In the CPA-security game, \mathcal{D}' is given a public-key pk . \mathcal{D}' submits two vectors of plaintexts $\mathbf{m}_0, \mathbf{m}_1$ where \mathbf{m}_0 is an all-zero bit vector as constructed in the simulated view, and \mathbf{m}_1 is the Bloom filter as in the real execution. \mathcal{D}' receives a vector of ciphertext \mathbf{c} corresponding to an encryption of either \mathbf{m}_0 or \mathbf{m}_1 , and directly forwards (pk, \mathbf{c}) to \mathcal{D} . Finally, \mathcal{D}' outputs what \mathcal{D} outputs. It is easy to see that the advantages of \mathcal{D} and \mathcal{D}' are the same.

Simulating the view of P_2 using $\text{Sim}_2^{\text{ePSI-CA}}$. P_2 's view can also be simulated in an analogous way. In particular, P_2 's view contains $\text{View}_2^{\text{coin}}$, P_1 's public-key pk_1 , encryptions of random numbers $\hat{n}_i = r_i + n_i$ under pk_2 , where n_i is a number in $[0, k]$, encryptions of coefficients of a polynomials $p_i(\cdot)$ whose roots are $r_i, r_i + 1, \dots, r_i + k - 1$ under pk_1 . The first two elements can be simulated using $\text{Sim}_2^{\text{coin}}$ and KeyGen. Encryptions under pk_2 can be generated by encrypting random numbers, while encryptions under pk_1 can be emulated by encryptions of 0. By a similar argument as above, the simulation will be indistinguishable from the real view.

5 Outsourcing to an Untrusted Server

In this section, we elaborate how to outsource some of the heavy computations in our protocols to an oblivious cloud. As to be discussed in Sec. 1.4, there are

quite a few outsourced PSI protocols in the literature [1, 2, 30–32]. However, they only implement the basic PSI protocol with different degrees of outsourceability. None of them supports flexible control as our $t^<$ -PSI and $t^>$ -PSI do. We extend the PSI model by introducing an additional cloud server, denoted by *CSP*. This party serves as an oblivious helper in our ePSI-CA, $t^<$ -PSI, and $t^>$ -PSI protocols: it helps P_1 and P_2 to perform some heavy computations but remains oblivious to both P_1 and P_2 's inputs, and the outcome of the protocol. The *CSP* is not trusted in the sense that it does not share any secret information with P_1 or P_2 , but we assume that it will follow the protocol specification faithfully. Such a semi-honest cloud server is widely accepted in the literature [1, 30–32].

5.1 Outsourcing $\Pi_{\text{ePSI-CA}}$

As the major building block of the bigger protocols $\Pi_{t^< \text{-PSI}}$ and $\Pi_{t^> \text{-PSI}}$, we first discuss the outsourceability of each step of $\Pi_{\text{ePSI-CA}}$.

Step 1 P_1 and P_2 run a coin-tossing protocol per execution to obtain the random seed for hash functions. The seed should remain hidden from *CSP*.

Step 2 P_2 bitwise-encrypts its Bloom filter. Notice that a Bloom filter is always a binary string, regardless of the set that it is encoding. As a result, P_2 can prepare a set of ciphertext encrypting “0”s and “1”s under pk_2 *offline* before any protocol execution.

Recall that P_2 uses an N -bit Bloom filter to represent its set. Hence it suffices to let *CSP* prepare $N_0 = \frac{1+\delta}{2} \times N$ encryptions of “0”s and the same number $N_1 = N_0$ for “1”s. (δ is a small constant and N is the same as that in Sec. 4.1). P_2 permutes the order of the ciphertexts according to a pseudorandom permutation generated from a secret random seed, and uploads the $(N_0 + N_1)$ ciphertexts to the *CSP*.

These ciphertexts can be reused for different protocol executions as follows: after obtaining the random seeds for the hash functions in Step 1, P_2 locally generates its *plaintext* Bloom filter. For each bit $BF[i]$, P_2 randomly selects one of the N_0 (or N_1) ciphertexts stored in *CSP*. P_2 informs *CSP* its choice by sending N indices in total, so that *CSP* can prepare an *encrypted* Bloom filter for P_2 . Note that only P_2 knows whether these ciphertexts are “0”s or “1”s. Hence *CSP* remain totally oblivious to the content of the Bloom filter.

Step 3 P_1 hashes its elements according to the hash functions agreed in Step 1, and obtains k indices for each elements. P_1 sends these locations to *CSP* so that *CSP* can homomorphically calculate the sum of these ciphertexts for P_1 . Since *CSP* does not know the random seeds for the hash functions generated in Step 1, these indices are complete random numbers from *CSP*'s point of view. Moreover, the random number r_i is independent of its corresponding element c_i , so *CSP* can choose r_i on behalf of P_1 .

Step 4 The coefficients of the polynomial $p_i(x) = (x - r_i) \cdots (x - r_i - k + 1)$ are completely determined by r_i , which are now selected by *CSP* in Step 3. Hence *CSP* can perform the whole Step 4.

Step 5 The polynomial evaluation step homomorphically computes encryption of $a_k \cdot \hat{n}_i^k + \dots + a_1 \cdot \hat{n}_i^1 + a_0$ in the ciphertext domain using \hat{n}_i from decryption. Note that we cannot reveal $\hat{n}_i = n_i + r_i$ to *CSP* because *CSP* knows r_i . The knowledge of n_i leaks whether the i -th element is in the intersection or not. Still, P_2 can locally compute ciphertexts of $a_j \cdot \hat{n}_i^j$ for all $j \in [0, k]$, and then ask *CSP* to add them together (which saves some computation). In this way, P_2 can still outsource $(k + 1)|C| - 1$ homomorphic additions to *CSP*.

Putting these together, Fig. 3 presents the outsourced $\Pi_{\text{ePSI-CA}}$ protocol. Table 3 shows its online computational complexity, with the saving highlighted in red. In short, P_1 can outsource *all* public key operations to the *CSP* while P_2 can outsource some. Sec. 6 will show that such improvement is significant.

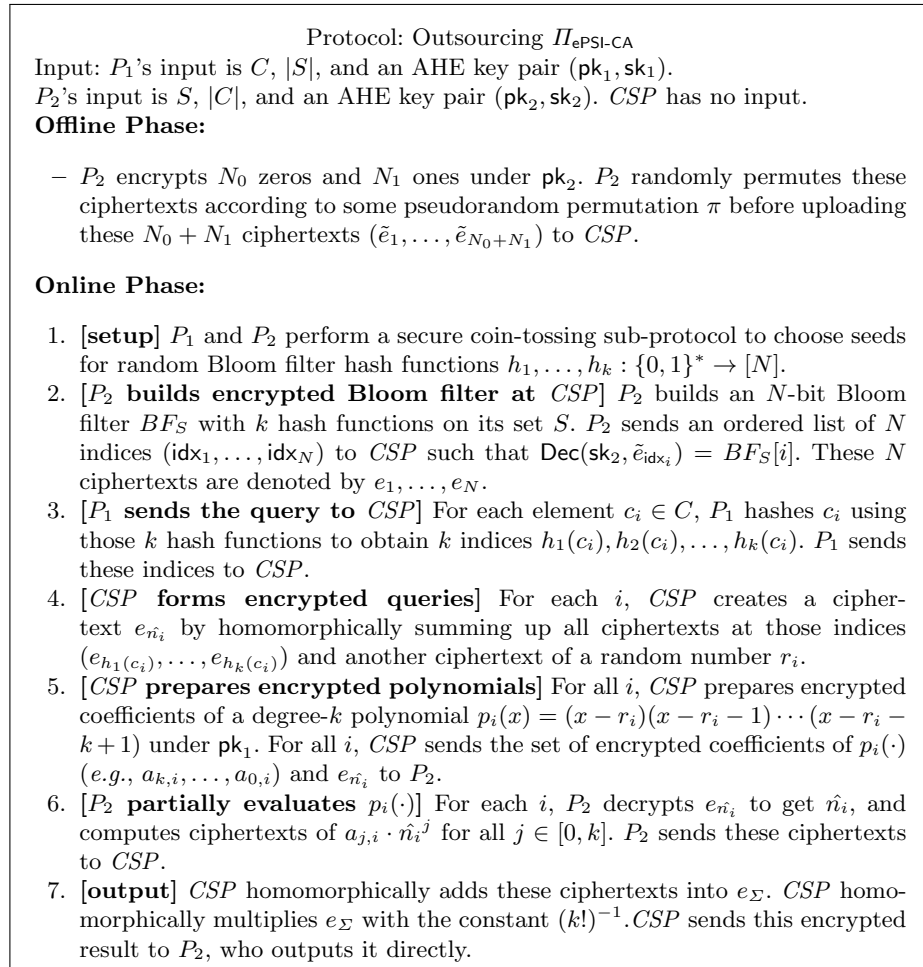
Table 3: Online Computational Complexity of Outsourced $\Pi_{\text{ePSI-CA}}$ (false positive rate $\epsilon = 2^{-\omega(\log \lambda)}$, # of hash $k = \omega(\log \lambda)$)

	Enc	Dec	addition	constant-multiplication
Step 2 (P_2)	0	0	0	0
Step 3 (P_1)	0	0	0	0
Step 4 (P_1)	0	0	0	0
Step 5 (P_2)	0	$ C $	0	$k C + 1$
P_1 total	0	0	0	0
P_2 total	0	$O(C)$	0	$O(\log^2(\lambda) C)$
<i>CSP</i> total	$O(\log^2(\lambda)(C + S))$	0	$O(\log^2(\lambda)(C + S))$	0

5.2 Outsourcing $\Pi_{t \leq \text{PSI}}$

We use $\Pi_{t \leq \text{PSI}}$ as an example to showcase outsourceability. It is very straightforward to apply the same technique to $\Pi_{t \geq \text{PSI}}$ and its generalizations. Basically most of the public key operations (except decryption) can be outsourced.

- Step 1 Invoking of $\mathcal{F}_{\text{ePSI-CA}}$. This can be (partially) outsourced as in Sec. 5.1.
- Step 2 The blinding factor r is independent of P_2 's private input. Hence the computation of $\text{Enc}(\text{pk}_1, |C \cap S| + r)$ can be delegated to the cloud.
- Step 3 The coefficients of the polynomial $p''(\cdot)$, like those in Step 4 of Fig. 2, are again independent of P_2 's private input. Therefore the encryption of coefficients can be outsourced to *CSP*, who will choose r', K on behalf of P_2 .
- Step 4 Since only P_1 knows sk_1 , the decryption of $\text{Enc}(\text{pk}_1, |C \cap S| + r)$ cannot be done by *CSP*. Moreover, the decryption result $|C \cap S| + r$ cannot be revealed to *CSP* because *CSP* knows r in Step 2. As a result, the evaluation of $p''(\cdot)$ cannot be fully outsourced to *CSP*, but still P_1 can locally compute encryptions of $a_j \cdot (|C \cap S| + r)^j$ for $j \in [0, t]$, as well as $\text{Enc}(\text{pk}_2, r'')$, and then ask *CSP* to homomorphically add them together.

Fig. 3: Outsourced Encrypted Private Set-Intersection Cardinality $\Pi_{\text{ePSI-CA}}$

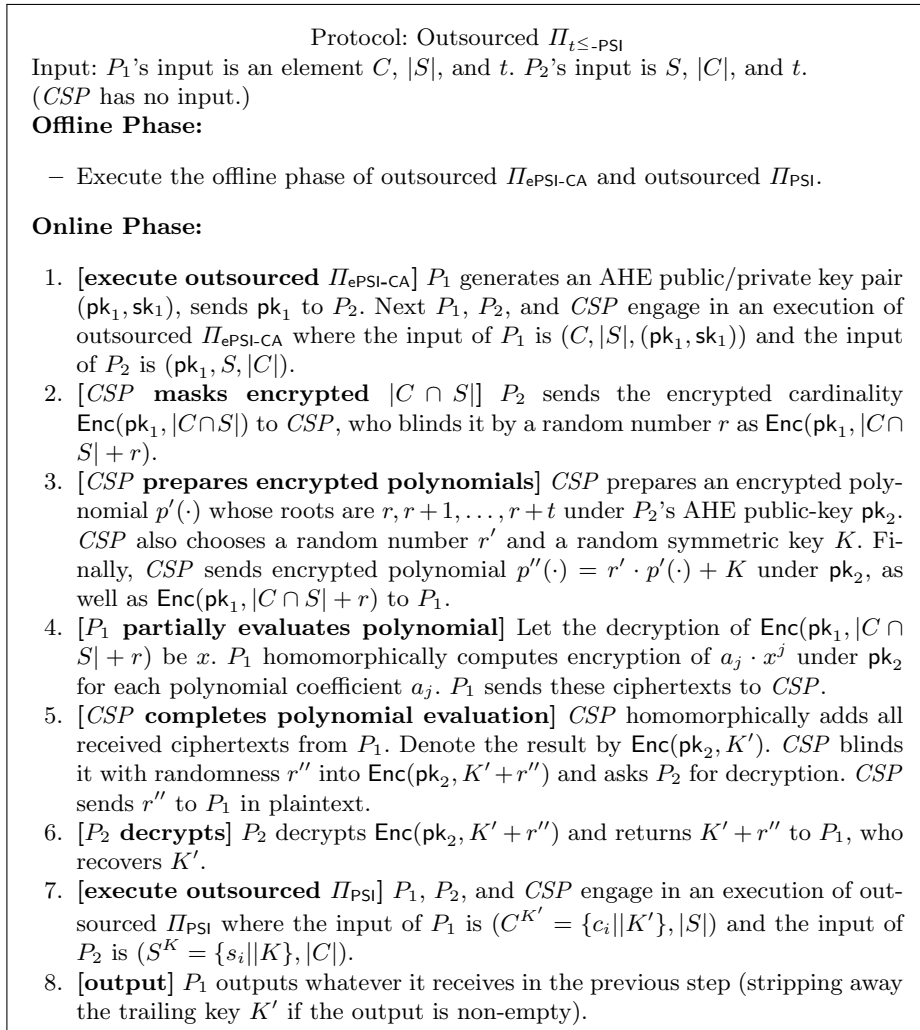
Step 5 Decryption cannot be outsourced.

Step 6 Intuitively outsourcing Π_{PSI} requires outsourceable PSI. There are quite a few potential solutions with different level of outsourceability in the literature [1, 2, 30–32]. We refer readers to these papers for more details.

Putting these pieces together, Fig. 4 presents the outsourced below-threshold private set-intersection protocol. Table 4 gives the online computational complexity of outsourced $\Pi_{t \leq \cdot \text{-PSI}}$.

6 Evaluation

We now examine the performance of our proposed protocols $\Pi_{\text{ePSI-CA}}$ and $\Pi_{t \leq \cdot \text{-PSI}}$. The experiment is conducted on a desktop machine running Windows 8.1, with

Fig. 4: Outsourced Below-Threshold Private Set-Intersection ($\Pi_{t \leq \text{PSI}}$)

2 Intel(R) Core(TM) i5-4590 3.30GHz CPUs, and 8GB RAM. We fix the size of the sets to be 100 and the threshold t is set to be half of set size, namely 50, as they should be sufficient for private-matching application in reality. Note that a dating site eHarmony recently only uses a couple questions that can be finished within 10 minutes to build up a model called “29 dimensions” to build up a user profile. The Bloom filter uses 30 hash functions instantiated by SHA-256, implemented by the OpenSSL library². This number of hash functions reflects a false negative rate $\epsilon = 2^{-30}$ and Bloom filter size 4500 bits. We note that one can achieve better performance at the cost of a larger false negative rate by reducing

² <https://www.openssl.org/>

Table 4: Online Computational Complexity of $\Pi_{t \leq \text{PSI}}$
 (using outsourced $\Pi_{\text{ePSI-CA}}$ (cf. Table 3) to instantiate of $\mathcal{F}_{\text{ePSI-CA}}$)

		Enc	Dec	addition	multiplication	Π_{PSI}
Step 1	(P_1)	0	0	0	0	0
	(P_2)	0	$O(C)$	0	$O(\log^2(\lambda) C)$	0
Step 2	(P_2)	0	0	0	0	0
Step 3	(P_2)	0	0	0	0	0
Step 4	(P_1)	1	1	0	$t + 1$	0
Step 5	(P_2)	0	1	0	0	0
Step 6	(P_1)	0	0	0	0	$O(C + S)$
	(P_2)	0	0	0	0	$O(C + S)$
P_1 Total		$O(1)$	$O(1)$	0	$O(t)$	$O(C + S)$
P_2 Total		0	$O(C)$	0	$O(\log^2(\lambda) C)$	$O(C + S)$
CSP Total		$O(\log^2(\lambda)(C + S))$	0	$O(\log^2(\lambda)(C + S))$	0	$O(C + S)$

the number of hash functions, which will lead to a smaller Bloom filter. We use existing Paillier encryption implementation³, and set the key length to be 2048. Experiments were measured in seconds via wall clock runtime, and the reported runtimes are the average of 100 trials.

We report the computational time of each step of $\Pi_{\text{ePSI-CA}}$ and $\Pi_{t \leq \text{PSI}}$ in Table 5 and 6 respectively. The last column represents the online computation time for P_1 and P_2 in the outsourced setting. We do not include the running time of the last step of $\Pi_{t \leq \text{PSI}}$, because it relies on existing efficient Π_{PSI} protocol (and its outsourced version), which is not part of the contribution of this paper.

From the tables, we see that the outsourced version achieves a significant reduction in computational time for both P_1 and P_2 . There are several ways to further reduce it for P_2 . First note that our protocols are easily parallelizable, which means significant improvement can be achieved via multi-threading. Second, recently Jost *et al.* [29] reported optimizations on Paillier cryptosystem that improves naïve implementation by a factor of over 150. Taking these optimizations into consideration, our constructions can finish within 1s.

7 Private Matchmaking

We discuss how to utilize our (outsourcable) protocols in our motivating application. The service provider acts as *CSP*. When each user joins the system, apart from generating a public key pair for AHE, they randomly pick a symmetric key K and use it to encrypt their profile such as photos and contact information. The system suggests a set of attributes (*e.g.*, highly-educated, smoking). The user can mark a subset of attributes as desired, and mark another disjoint subset as undesired. The unmarked ones will be considered as “don’t care”, and they will not be part of the protocol input. The user also picks two thresholds:

³ <https://github.com/herumi/mie>

Table 5: Execution time of $\Pi_{\text{ePSI-CA}}$ (without the optimization of [29])

		time (s)	online (s)
Step 2 (P_2)	create Bloom filter	0.001	0
	encrypt Bloom filter	60.404	0
Step 3 (P_1)	query Bloom filter	2.685	0.001
Step 4 (P_1)	encrypt polynomial	81.268	0
Step 5 (P_2)	decryptions	1.322	1.322
	evaluate polynomials	76.627	76.562
P_1 total		83.953	0.001
P_2 total		138.353	77.884

Table 6: Execution time of $\Pi_{t^{\leq}\text{-PSI}}$, $t = 50$ (without the optimization of [29])

		time (s)	online (s)
Step 1 ($\Pi_{\text{ePSI-CA}}$)	P_1	83.953	0.001
	P_2	138.353	77.884
Step 2 (P_2)		0.026	0
Step 3 (P_2)		1.363	0
Step 4 (P_1)		1.325	1.324
Step 5 (P_2)		0.013	0.013
P_1 total(*)		85.278	1.325
P_2 total(*)		139.755	77.897

t_o which is for the least number of desired attributes, another one is t_b which is for the maximum number of undesired attributes.

To use both $t^{\leq}\text{-PSI}$ and $t^{\geq}\text{-PSI}$ simultaneously, the user splits the symmetric key K into two parts by a simple (2, 2) secret sharing based on XOR. Specifically, the user picks a key K_o which is as long as K , and outputs both K_o and $K_b = K \oplus K_o$. The user puts all the undesired attributes as the private set input of $t^{\leq}\text{-PSI}$. Likewise, the user puts all the desired attributes as the input of $t^{\geq}\text{-PSI}$.

Users who joined the service can either be passively matched by others or actively requests for matching. Here we discuss a typical protocol run from the perspective of an active user. The service provider will pick a potential user, called a passive user, and execute the PSI protocols on behalf. In other words, that is where the outsourced feature of our protocols come into the play. Suppose the number of undesired attributes of this passive user is below the threshold t_b after running $t^{\leq}\text{-PSI}$, and the number of desired attributes is over the threshold t_o after running $t^{\geq}\text{-PSI}$.

If our protocols are used directly, the active user will get the intersection of either kinds of attributes. This may not be the most privacy-preserving way for doing matchmaking since the passive user has no way to control about whether revealing any secret information (such as the profile) or not. Luckily, the intersection result can be easily removed from our protocols by removing the last

step of performing the (keyed-)PSI. As a result, even the passive user passed the matching criteria, the user only obtains a secret key generated by the passive user (from the second last step of the protocol).

Here, we utilize the idea of secret transfer with access structure from Zhao and Chow [39]. This key will serve as a proof of criteria satisfaction. Upon the presentation of the aforementioned secret key, the user can decide to reveal K_b (resp. K_o) or not. Such decision can be done after the service provider execute the PSI protocols on behalf of this passive user. If the interest is mutual, *i.e.*, the requesting user also satisfies the search criteria of the “passive” user, the passive user can finally reveal the encrypted profile (by releasing both K_b and K_o .)

Two remarks are in order. First, note that even with the help of the service provider who mediates the requests between two users, the outsourced PSI protocols are not entirely non-interactive. For this, the service provider still needs to relay message between the users. However, it matches with the workflow used by non-private matchmaking apps nowadays which the user cannot connect to another user until there are mutual interests. Second, some user may expect to assign different weighting to different attributes. A trivial approach is to replicate the attribute multiple times. Devising cleverer solutions which maintain a similar level of efficiency requires further twisting of our PSI protocols (perhaps by borrowing techniques from Zhao and Chow [39]) We left it as future work.

8 Conclusion

We propose efficient protocols for three important extensions of PSI, namely, *existential private set-intersection*, *over threshold private set-intersection*, and *below threshold private set-intersection*. The last two provide affirmative answers to two open problems posed very recently in the literature. All of our constructions achieve linear computational complexity by utilizing and extending existing building blocks in a novel way. We prove that our constructions are secure against semi-honest adversaries. Our constructions provide useful building blocks to realize privacy-preserving online matchmaking.

References

1. Aydin Abadi, Sotirios Terzis, and Changyu Dong. O-PSI: delegated private set intersection on outsourced datasets. In *SEC*, pages 3–17, 2015.
2. Aydin Abadi, Sotirios Terzis, and Changyu Dong. VD-PSI: verifiable delegated private set intersection on outsourced private datasets. In *FC*, 2016.
3. Giuseppe Ateniese, Emiliano De Cristofaro, and Gene Tsudik. (If) size matters: Size-hiding private set intersection. In *PKC*, 2011.
4. Burton H. Bloom. Space/time trade-offs in hash coding with allowable errors. *Commun. ACM*, 13(7):422–426, 1970.
5. Dan Boneh, Eu-Jin Goh, and Kobbi Nissim. Evaluating 2-dnf formulas on ciphertexts. In *TCC*, 2005.

6. Prosenjit Bose, Hua Guo, Evangelos Kranakis, Anil Maheshwari, Pat Morin, Jason Morrison, Michiel H. M. Smid, and Yihui Tang. On the false-positive rate of Bloom filters. *Inf. Process. Lett.*, 108(4):210–213, 2008.
7. Tatiana Bradley, Sky Faber, and Gene Tsudik. Bounded size-hiding private set intersection. In *SCN*, 2016.
8. Jan Camenisch and Gregory M. Zaverucha. Private intersection of certified sets. In *FC*, 2009.
9. Ran Canetti. Security and composition of multiparty cryptographic protocols. *J. Cryptology*, 13(1):143–202, 2000.
10. Xavier Carpent, Sky Faber, Tomas Sander, and Gene Tsudik. Private set projections & variants. In *WPES*, 2017.
11. Michele Ciampi and Claudio Orlandi. Combining private set-intersection with secure two-party computation. Cryptology ePrint Archive, Report 2018/105, 2018.
12. Emiliano De Cristofaro, Jihye Kim, and Gene Tsudik. Linear-complexity private set intersection protocols secure in malicious model. In *ASIACRYPT*, 2010.
13. Emiliano De Cristofaro and Gene Tsudik. Practical private set intersection protocols with linear complexity. In *FC*, 2010.
14. Dana Dachman-Soled, Tal Malkin, Mariana Raykova, and Moti Yung. Efficient robust private set intersection. In *ACNS*, 2009.
15. Paolo D’Arco, Maria Isabel Gonzalez Vasco, Angel L. Pérez del Pozo, and Claudio Soriente. Size-hiding in private set intersection: Existential results and constructions. In *AFRICACRYPT*, 2012.
16. Alex Davidson and Carlos Cid. An efficient toolkit for computing private set operations. In *ACISP Part II*, 2017.
17. Sumit Kumar Debnath and Ratna Dutta. Secure and efficient private set intersection cardinality using bloom filter. In *ISC*, 2015.
18. Sumit Kumar Debnath and Ratna Dutta. How to meet big data when private set intersection realizes constant communication complexity. In *ICICS*, 2016.
19. Sumit Kumar Debnath and Ratna Dutta. Provably secure fair mutual private set intersection cardinality utilizing bloom filter. In *Inscrypt*, 2016.
20. Changyu Dong, Liqun Chen, and Zikai Wen. When private set intersection meets big data: an efficient and scalable protocol. In *CCS*, 2013.
21. Rolf Egert, Marc Fischlin, David Gens, Sven Jacob, Matthias Senker, and Jörn Tillmanns. Privately computing set-union and set-intersection cardinality via bloom filters. In *ACISP*, 2015.
22. Michael J. Freedman, Kobbi Nissim, and Benny Pinkas. Efficient private matching and set intersection. In *EUROCRYPT*, 2004.
23. Oded Goldreich. *The Foundations of Cryptography - Volume 2, Basic Applications*. Cambridge University Press, 2004.
24. Per A. Hallgren, Claudio Orlandi, and Andrei Sabelfeld. Privatepool: Privacy-preserving ridesharing. In *CSF*, 2017.
25. Carmit Hazay. Oblivious polynomial evaluation and secure set-intersection from algebraic PRFs. In *TCC Part-II*, 2015.
26. Carmit Hazay and Kobbi Nissim. Efficient set operations in the presence of malicious adversaries. In *PKC*, 2010.
27. Susan Hohenberger and Stephen A. Weis. Honest-verifier private disjointness testing without random oracles. In *PET*, 2006.
28. Yan Huang, David Evans, and Jonathan Katz. Private set intersection: Are garbled circuits better than custom protocols? In *NDSS*, 2012.

29. Christine Jost, Ha Lam, Alexander Maximov, and Ben J. M. Smeets. Encryption performance improvements of the paillier cryptosystem. *Cryptology ePrint Archive*, Report 2015/864, 2015.
30. Seny Kamara, Payman Mohassel, Mariana Raykova, and Seyed Saeed Sadeghian. Scaling private set intersection to billion-element sets. In *FC*, 2014.
31. Florian Kerschbaum. Collusion-resistant outsourcing of private set intersection. In *SAC*, 2012.
32. Florian Kerschbaum. Outsourced private set intersection using homomorphic encryption. In *ASIACCS*, 2012.
33. Lea Kissner and Dawn Xiaodong Song. Privacy-preserving set operations. In *CRYPTO*, 2005.
34. Vladimir Kolesnikov, Ranjit Kumaresan, Mike Rosulek, and Ni Trieu. Efficient batched oblivious PRF with applications to private set intersection. In *CCS*, 2016.
35. Vladimir Kolesnikov, Naor Matania, Benny Pinkas, Mike Rosulek, and Ni Trieu. Practical multi-party private set intersection from symmetric-key techniques. In *CCS*, 2017.
36. Pascal Paillier. Public-key cryptosystems based on composite degree residuosity classes. In *EUROCRYPT*, 1999.
37. Benny Pinkas, Thomas Schneider, Gil Segev, and Michael Zohner. Phasing: Private set intersection using permutation-based hashing. In *USENIX Security*, 2015.
38. Benny Pinkas, Thomas Schneider, and Michael Zohner. Faster private set intersection based on OT extension. In *USENIX Security*, 2014.
39. Yongjun Zhao and Sherman S. M. Chow. Are you the one to share? Secret transfer with access structure. *PoPETs*, 2017(1):149–169, 2017.

A Secure Two-Party Computation

We use the simulation-based security definition for two-party computation (2PC). More details can be referred to [23]. A 2PC protocol π computes a function that maps a pair of inputs to a pair of outputs $f : \{0, 1\}^* \times \{0, 1\}^* \rightarrow \{0, 1\}^* \times \{0, 1\}^*$, where $f = (f_1, f_2)$. For every pair of inputs $x, y \in \{0, 1\}^*$, the output-pair is a random variable $(f_1(x, y), f_2(x, y))$. The first party obtains $f_1(x, y)$ and the second party obtains $f_2(x, y)$.

We first consider *static semi-honest* adversaries, which can control one of the two parties and assumed to follow the protocol specification exactly. However, it may try to learn more information about the other party’s input.

In the semi-honest model, a protocol π is secure if whatever can be computed by a party in the protocol can be obtained from its input and output only. This is formalized by the simulation paradigm. We require a party’s *view* in a protocol execution to be simulatable given only its input and output. The view of the party i during an execution of π on input (x, y) is denoted by $\text{View}_i^\pi(x, y) = (w, r^i, m_1^i, \dots, m_t^i)$, where $w \in (x, y)$ is the input of i , r^i is i ’s internal random coin tosses, and m_j^i denotes the j -th message that it received.

Definition 4 (Semi-honest Model). *Protocol π is said to securely compute a deterministic function $f = (f_1, f_2)$ in the presence of static semi-honest adversaries if there exists PPT algorithms $\text{Sim}_1, \text{Sim}_2$ such that*

$$\{\text{Sim}_1(x, f_1(x, y))\}_{x, y} \stackrel{c}{=} \{\text{View}_1^\pi(x, y)\}_{x, y}, \{\text{Sim}_2(y, f_2(x, y))\}_{x, y} \stackrel{c}{=} \{\text{View}_2^\pi(x, y)\}_{x, y}.$$

The \mathcal{F} -hybrid model. We will use some secure two-party protocols as sub-protocols in our constructions. We will describe our protocols in a “hybrid model” where the two parties both interact with each other and use trusted help. When constructing a protocol π that uses a sub-protocol that securely computes some functionality \mathcal{F} , we consider the case that the parties run π and use “ideal calls” to a trusted party for computing \mathcal{F} . Upon receiving the inputs from the parties, the trusted party computes \mathcal{F} and sends all parties the corresponding output. Then after receiving these outputs back from the trusted party, the protocol π continues. By the composition theorem [9], any protocol that securely implements \mathcal{F} can replace the ideal calls to \mathcal{F} .

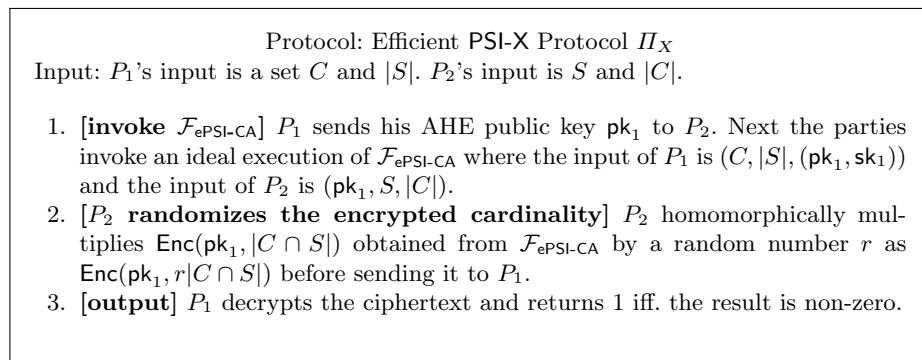
B Private Projection

Recently, Carpent *et al.* [10] started the study of private projection (PSI-P). Different from traditional PSI, the server in PSI-P has a database $DB = \{d_1, \dots, d_n\}$ (which may contain duplications) and a column of attributes $A = \{a_1, \dots, a_n\}$, while the client has a set of attributes $B = \{b_1, \dots, b_m\}$. After a PSI-P invocation, the server learns nothing while the client only learns $\{d_i | \exists(i, j) \text{ s.t. } b_j = a_i\}$. In particular, the client should not know which matching b_j corresponding to which d_i , nor how many matching b_j that d_i corresponds to.

PSI-P finds application in matching indicators of compromise (IOC), where attribute column A represents IOC while the database DB represents patches of known vulnerabilities and attacks. In such scenarios, the correspondence between IOC and patches can be sensitive. Attackers may slightly adapt their attack strategy to avoid being detected by the same IOC. Using PSI-P as a solution, the server (as a security expert) can protect the valuable information (IOC and DB of patches), yet provides a just-enough list of patches to the client, without letting the server know private set B (*e.g.*, network traffic).

C Existential PSI (PSI-X)

Carpent *et al.* recognize that none of the existing PSI protocols (or their variants) satisfies the security requirements of PSI-P. While outside their radar, oblivious transfer for a sparse array [39] can approximate PSI-P since it leaks the number of distinct data elements. They thus propose a series of protocols with different leakages, and finally construct a full-fledged PSI-P from any PSI-X. Unfortunately, their PSI-X is inefficient. For client and server set sizes being m and n , the computational complexity is of order $O(mn)$. Any improvement for PSI-X immediately leads to a better private projection protocol. Π_X can be easily realized by slightly modifying Step 5 of our $\Pi_{\text{ePSI-CA}}$. Recall that in Step 5 of $\Pi_{\text{ePSI-CA}}$ P_2 obtains the set-intersection cardinality encrypted under pk_1 . P_2 can rerandomize this ciphertext before sending it to P_1 .

Fig. 5: Efficient PSI-X Protocol Π_X

Definition 5 (Existential Private Set-Intersection (PSI-X)). Let S and C be subsets of a predetermined domain, the functionality \mathcal{F}_X is:

$$((C, |S|), (S, |C|)) \mapsto \begin{cases} (1, \perp) & \text{if } C \cap S \neq \phi \\ (0, \perp) & \text{otherwise} \end{cases}$$

We begin with a high-level description of the first (but inefficient) PSI-X construction by Carpent *et al.* Suppose party P_1 has a set C of m elements and party P_2 has a set S of size n . In the existing PSI-X [10], they first jointly choose a *single* 2-universal hash function $h(\cdot)$ mapping set elements to $[N]$ where $N \in O(mn)$. P_1 transforms set C into a bit string v_C of length N such that the $v_C[i] = 1$ if and only if $\exists x \in C : h(x) = i$. P_2 also performs similar operations on S to derive v_S . P_1 generates a BGN [5] public/private key pair (pk, sk) , publishes pk and $\text{Enc}(\text{pk}, v_C[i])$ for all $i \in [N]$. P_2 also encrypts v_S under pk and evaluates the 2-DNF $\phi = \bigvee_i (v_C[i] \wedge v_S[i])$ via the homomorphism of BGN. P_2 sends encryption of $r \cdot \phi$ to P_1 for a random r . If P_1 gets 0 after decryption, P_1 concludes that the intersection is *definitely* empty; otherwise it is *probably* non-empty. The uncertainty stems from the possible collision due to the hash function. One can reduce the error rate by increasing N , or repeating R independent instances of this protocol. Both increase the overall computational complexity in terms of the number of ciphertext multiplications. Carpent *et al.* [10] show that the optimal choice is $N = \frac{mn}{\log 2}$ for any error rate, resulting in $O(mn)$ complexity. PSI-X with linear computational complexity was an open problem before our paper.

The detailed protocol Π_X is in Fig. 5, which is very simple in the $\mathcal{F}_{\text{ePSI-CA}}$ model. Basically we only add one last step: P_2 blinds the encrypted cardinality using a random r before sending it to P_1 .

Corollary 1. Assuming the existence of a CPA-secure additive homomorphic encryption scheme (KeyGen, Enc, Dec), whose plaintext space is super polynomial in the security parameter; then the protocol Π_X in Fig. 5 securely implements the functionality \mathcal{F}_X in Definition 5 under the semi-honest model.