

# Nothing Refreshes Like a RePSI: Reactive Private Set Intersection<sup>\*†</sup>

Andrea Cerulli<sup>1</sup>, Emiliano De Cristofaro<sup>1</sup>, and Claudio Soriente<sup>2</sup>

<sup>1</sup>University College London    <sup>2</sup> NEC Laboratories Europe

**Abstract.** Private Set Intersection (PSI) is a popular cryptographic primitive that allows two parties, a client and a server, to compute the intersection of their private sets, so that the client only receives the output of the computation, while the server learns nothing besides the size of the client’s set. A common limitation of PSI is that a dishonest client can progressively learn the server’s set by enumerating it over different executions. Although these “oracle attacks” do not formally violate security according to traditional secure computation definitions, in practice, they often hamper real-life deployment of PSI instantiations, especially if the server’s set does not change much over multiple interactions.

In a first step to address this problem, this paper presents and studies the concept of Reactive PSI (RePSI). We model PSI as a reactive functionality, whereby the output depends on previous instances, and use it to limit the effectiveness of oracle attacks. We introduce a general security model for RePSI in the (augmented) semi-honest model and a construction which enables the server to control how many inputs have been used by the client across several executions. In the process, we also present the first practical construction of a Size-Hiding PSI (SHI-PSI) protocol in the standard model, which may be of independent interest.

## 1 Introduction

Private Set Intersection (PSI) lets two parties compute the intersection of their private sets, drawn from a common universe, without disclosing items outside the intersection. In its most common formulation, only one party, usually referred to as the *client*, obtains the intersection, while the other, aka *server*, only learns the size of the client’s set. Over the past few years, PSI has been used in numerous privacy-friendly applications, including ridesharing [HOS17], collaborative threat mitigation [FDCB15], genomic testing [BBD<sup>+</sup>11], and online advertising [IKN<sup>+</sup>17].

Nonetheless, there are some challenging issues limiting the adoption of PSI in practice. In particular, if two parties run the protocol several times, the server is vulnerable to *oracle attacks*. In such an attack, a dishonest client progressively learns the server’s set by enumerating it over different executions. Although this does not formally violate security definitions of two-party computation [Gol04], it may hamper real-life deployment of PSI, especially if the server’s set is mostly static. Moreover, in the *Size-Hiding*

---

<sup>\*</sup>A preliminary version of this paper appears in the Proceedings of the 16th International Conference on Applied Cryptography and Network Security (ACNS 2018). This is the full version.

<sup>†</sup>The first author was supported by the European Research Council under the European Union’s Seventh Framework Programme (FP/2007-2013) / ERC Grant Agreement n. 307937.

variant of PSI [ADT11], where the server does not learn the size of client’s set, the problem is further compounded as the server cannot limit the size of client’s input.

Aiming to mitigate oracle attacks in PSI protocols, we start reasoning about the security of this cryptographic primitive across multiple runs. To this end, we introduce the notion of Reactive PSI (RePSI), along with a general security model in the augmented semi-honest model [Gol04], and set to propose provably secure instantiations.

Let us first consider a naïve solution. In the non size-hiding setting, using certain PSI protocols, e.g., [DT10], one could in theory let the client re-use the randomness for the elements in its input set that do not change across runs. This way, the server learns how many of the client’s elements are “fresh” in the current run and imposes an upper-bound. However, this approach at the very least makes two protocol executions linkable as it reveals the patterns of the client’s inputs. Moreover, if the distribution of client’s elements is somewhat predictable, this might actually reveal too much information. By contrast, our goal is to provide stronger definitions whereby the client does not reveal its input patterns, but only proves that number of unique elements input from the first run up to the current one is below a given threshold.

## 1.1 Roadmap

In this paper, we tackle the issue of oracle attacks in Private Set Intersection (PSI) by extending security definitions to account for reactive functionalities, whereby the output of the current execution can depend on previous executions.

First, we introduce the notion of Reactive PSI (RePSI), along with a general security model in the augmented semi-honest model [Gol04]. In this model, the adversary is assumed to follow the specifications of the protocol (as in the standard semi-honest model) but it is allowed to adaptively modify the inputs used by the controlled party at each protocol run. We argue that the augmented semi-honest model can effectively model oracle attacks in PSI, whereas, the standard semi-honest model cannot, since it prevents the adversary to change the input of the corrupted party between protocol executions. That is, the adversary can only leverage honestly generated transcripts. Also, although we do not yet provide security in the fully malicious setting, we believe that ours is an important first step towards the development of efficient protocols. In fact, there exist general transformations [GMW87, Gol04] allowing to compile a semi-honest secure protocol into one secure against malicious adversaries, and efficient PSI-like protocols are also traditionally in semi-honest settings (see Section 1.2). Moreover, our definitions are general enough to capture various types of reactive functionalities and they cover the sequential composition of standard (i.e., stateless) PSI protocols.

Then, we provide two constructions, one static and one reactive. We focus on the size-hiding setting since, as mentioned above, the fact that the server cannot even check and limit the number of client’s inputs in a single execution, makes oracle attacks significantly worse. Our static construction, named *Bounded-Input PSI* limits the size of the client’s input set at every protocol run. We achieve this by adapting the Bounded Size-Hiding PSI recently presented by Bradley, Faber, and Tsudik [BFT16], which provided security in the Random Oracle Model (ROM). As an additional contribution, we instantiate Bounded Size-Hiding PSI in the standard model, thus also presenting the first practical Size-Hiding PSI protocol not in ROM. Our reactive construction, called

***Input Controlling RePSI***, enables the server to control how many inputs have been used by the other party across several executions. Specifically, it limits the size of the unions set stemming from the union of client’s input sets across all protocol runs. Input Controlling RePSI, therefore, addresses oracle attacks in practical scenarios where a client and a server engage in multiple PSI executions.

By modeling PSI as a reactive functionality, we require that client and server keep *state* across protocol executions. Nevertheless, the amount of state information kept by the two parties in our constructions is small and independent of the number of runs.

## 1.2 Related Work

To the best of our knowledge, the problem of Reactive PSI has not been studied in literature. Standard security definitions for semi-honest and malicious two-party and multiparty computation can be extended to model security of generic protocols computing reactive functionalities. The augmented semi-honest model was introduced by Goldreich [Gol04] to bridge the semi-honest model and the malicious model and used it as an intermediate step in the compilation of secure protocols from the semi-honest to the malicious settings. Hazay and Lindell [HL10a] observed that security in the malicious settings sometime does not imply security in the semi-honest settings, while this anomaly does not happen in the augmented semi-honest model.

Overall, prior work on PSI can be grouped in protocols using special-purpose constructions [FNP04, DT10], oblivious transfer and its extensions [PSZ14, PSSZ15], and/or generic garbled circuits [PSSZ15]. Most protocols are secure against semi-honest adversaries [FNP04, DT10, PSZ14, PSSZ15], with fewer, less efficient ones, against malicious ones [DKT10, JL10, RR17]. Also, protocols by Hazay et al. [HL08] operate in the covert model (i.e., a malicious adversary may be able to cheat but it can get caught with at least a certain probability).

There are also a few variants to the standard PSI functionality. Besides the size-hiding one discussed above [ADT11, BFT16], Authorized PSI [CZ09, DT10] partially mitigates malicious behavior by introducing a trusted party that authorizes (i.e., signs) the elements that a client can use as input. However, finding a common trusted party may be hard in most practical use cases.

More closely related to our work are the protocols proposed in [BFT16] and [DMV13]. Bradley, Faber, and Tsudik [BFT16] introduce the concept of Bounded Size-Hiding PSI, which allows the client to hide the size of its input, and the server to impose an upper-bound on the size of the client’s set *for the current run*. We start from the protocol of [BFT16] and cast it within the framework of RePSI to counter oracle attacks *across multiple runs*. Furthermore, while [BFT16] works in the random oracle model, we instantiate it in the standard model. Dagdelen, Mohassel, and Venturi [DMV13] introduce the concept of rate-limited Secure Function Evaluation (SFE), whereby protocol participants can monitor and limit the number of distinct inputs (i.e., rate) used by their counterparts in multiple executions of an SFE. They present compilers by which any SFE scheme can be turned into a rate-limited one. In particular, the “rate-hiding” compiler [DMV13] may be applied to a PSI protocol to achieve the same provisions of our Input Controlling RePSI. We take a less general approach and focus

on PSI, by incorporating reactiveness in the functionality and achieving a more efficient construction (see Section 5.3).

Finally, a theoretical construction based on fully-homomorphic encryption for size-hiding PSI in the standard model was recently presented in [COV15].

### 1.3 Paper Organization

Next section introduces some preliminaries, then, Section 3 provides security definitions for the Reactive PSI primitive in the augmented semi-honest model. Next, in Sections 4 and 5, we present our constructions of Bounded Input RePSI and Input Controlling RePSI, respectively. Finally, the paper concludes in Section 6.

## 2 Preliminaries

In this section, we introduce notation, cryptographic assumptions and building blocks used later on in the paper.

We write  $y \leftarrow \mathcal{A}(x)$  for a probabilistic algorithm returning output  $y$  given as input  $x$ . In case we want to specify the randomness  $r$  used, we write  $y = \mathcal{A}(x; r)$ . We implicitly assume all the algorithms considered in this paper to receive as input the security parameter  $\lambda$ . For functions  $f, g : \mathbb{N} \rightarrow [0, 1]$  we write  $f(\lambda) \approx g(\lambda)$  if  $|f(\lambda) - g(\lambda)| = \lambda^{\omega(1)}$ . We say a function  $f$  is *overwhelming* if  $f(\lambda) \approx 1$  and *negligible* if  $f(\lambda) \approx 0$ .

### 2.1 Bilinear Groups

A bilinear group is a tuple  $(p, \mathbb{G}, \mathbb{G}_T, e, g)$  s.t.  $\mathbb{G}$  and  $\mathbb{G}_T$  are groups of prime order  $p$  and  $g \in \mathbb{G}$  generates the group  $\mathbb{G}$ . The function  $e$  is an efficiently computable bilinear map  $e : \mathbb{G} \times \mathbb{G} \rightarrow \mathbb{G}_T$  such that  $e(g, g)$  is a generator of  $\mathbb{G}_T$ . We assume there are probabilistic polynomial time generators  $\mathcal{G}$  and  $\mathcal{BG}$  that, given as input the security parameter, return the description of a group  $(p, \mathbb{G}, g) \leftarrow \mathcal{G}(\lambda)$  and bilinear group  $(p, \mathbb{G}, \mathbb{G}_T, e, g) \leftarrow \mathcal{BG}(\lambda)$ , respectively. In the constructions of Sections 4 and 5, we rely on the exponent Strong Diffie-Hellman (Exponent  $q$ -SDH) and the Decisional Bilinear Diffie-Hellman Inversion problem ( $q$ -DBDHI) recalled in Appendix A.

### 2.2 Bilinear Accumulators

A cryptographic accumulator is a primitive that allows to give a compact representation of a set and that enables to efficiently prove membership of an element into the accumulated set. Accumulators were firstly introduced by Benaloh and de Mare [BDM94] and were later extended and provided with additional properties [BP97, CL02, Ngu05, DHS15, GOP<sup>+</sup>16, CKS09].

A (static) accumulator consists of four algorithms (KeyGen, Eval, WitGen, Verify). The key generation algorithm KeyGen takes as input the security parameter and generates a secret and an evaluation key pair  $(sk, ek)$  for the accumulator. The evaluation algorithm Eval gets as input the evaluation key  $ek$  and a set  $A$  of values and returns

an accumulator  $\text{acc}_A$ . The  $\text{WitGen}$  and  $\text{Verify}$  are deterministic algorithms for, respectively, producing and verifying a witness  $\text{wit}$  for the membership of an element  $a \in A$  in a given accumulator  $\text{acc}_A$ . We follow [DHS15] on modelling  $\text{Eval}$  and  $\text{WitGen}$  to optionally get as input the secret key  $sk$ , since this makes the algorithms more efficient. We denote the optional input by writing  $\boxed{sk}$ .

The main security properties required from accumulators are: *correctness*, i.e. honestly generated witnesses should verify; *collision-freeness*, i.e. that it is unfeasible to compute a witness for elements not included in the accumulated set; and *indistinguishability*, i.e. the accumulator does not reveal any information on the accumulated set. We recall security definitions for cryptographic accumulator in Appendix B which we adapted from [DHS15] to our settings.

In our constructions of PSI we will later use the accumulator introduced by Nguyen in [Ngu05] based on bilinear pairings. Since we will not require the possibility of removing elements from an accumulator, we restrict Nguyen's construction [Ngu05] to a static accumulator, description of which can be found in Figure 1.

<p><math>\text{KeyGen}(\lambda) \rightarrow (sk, ek)</math>:</p> <ul style="list-style-type: none"> <li>• <math>(p, \mathbb{G}, \mathbb{G}_T, e, g) \leftarrow \mathcal{BG}(\lambda)</math></li> <li>• <math>x \leftarrow \mathbb{Z}_p^*</math></li> <li>• <math>sk := x</math></li> <li>• <math>ek := (g, g^x, g^{x^2}, \dots, g^{x^q})</math></li> </ul>	<p><math>\text{Eval}(\boxed{sk}, ek, A) \rightarrow \text{acc}_A</math>:</p> <ul style="list-style-type: none"> <li>• Parse <math>A = (a_1, \dots, a_n)</math> for <math>a_i \in \mathbb{Z}_p</math></li> <li>• <math>\text{Ch}_A(X) = \sum_{i=0}^n c_i X^i</math></li> <li>• <math>r \leftarrow \mathbb{Z}_p^*</math></li> <li>• <math>\text{acc}_A := g^{r \text{Ch}_A(x)} = \left( \prod_{i=0}^n (g^{x^i})^{c_i} \right)^r</math></li> </ul>
<p><math>\text{WitGen}(\boxed{sk}, ek, \text{acc}_A, r, A, a) \rightarrow \text{wit}</math>:</p> <ul style="list-style-type: none"> <li>• Parse <math>A = (a_1, \dots, a_n)</math> for <math>a_i \in \mathbb{Z}_p</math></li> <li>• <math>\text{Ch}_{A \setminus \{a\}}(X) = \sum_{i=0}^n d_i X^i</math></li> <li>• <math>\text{wit} := \text{acc}_A^{\frac{1}{a+x}} = g^{r \text{Ch}_{A \setminus \{a\}}(x)} = \left( \prod_{i=0}^n (g^{x^i})^{d_i} \right)^r</math></li> </ul>	<p><math>\text{Verify}(ek, \text{acc}_A, a, \text{wit}) \rightarrow 0/1</math>:</p> <ul style="list-style-type: none"> <li>• If <math>e(\text{acc}_A, g) = e(\text{wit}, g^x \cdot g^a)</math> : Return 1</li> <li>• Else: Return 0</li> </ul>

**Fig. 1:** Bilinear Accumulators.

Let  $A = \{a_1, \dots, a_n\}$  be a set of elements  $a_i \in \mathbb{Z}_p$  that we wish to include into an accumulator. We first start by computing the *characteristic polynomial* representation of set  $A$ . This is the monic polynomial  $\text{Ch}_A(X) \in \mathbb{Z}_p[X]$  which has roots in the elements contained in the set  $A$ , namely  $\text{Ch}_A(X) = \prod_{i=1}^n (X + a_i)$ . In order to efficiently evaluate the accumulator, it will be convenient to express such polynomial using its coefficient representation, i.e. computing  $c_j$  such that  $\text{Ch}_A(X) = \prod_{i=1}^n (X + a_i) = \sum_{j=0}^n c_j X^j$ . We stress that given  $A$  it is always possible to efficiently compute the coefficient  $c_j$  of  $\text{Ch}_A(X)$ .

The evaluation key of [Ngu05] bilinear accumulator consists of  $ek = (g, g^x, g^{x^2}, \dots, g^{x^q}) \in \mathbb{G}^{q+1}$ , where  $g$  is a generator of the group  $\mathbb{G}$  and  $x \in \mathbb{Z}_p$  is a secret value.

Given the evaluation key and polynomial  $\text{Ch}_A(X)$  of degree at most  $q$ , it is possible to compute  $g^{\text{Ch}_A(x)}$ . This is done by first expanding  $\text{Ch}_A(X)$  into its coefficient representation, i.e  $\text{Ch}_A(X) = \sum_{i=0}^q c_i X^i$ , and then computing  $g^{\text{Ch}_A(x)} = \prod_{i=0}^q (g^{x^i})^{c_i}$ . An accumulator  $\text{acc}_A$  to a set  $A$  is computed by picking a random  $r \leftarrow \mathbb{Z}_p$  and setting  $\text{acc}_A = (g^{\text{Ch}_A(x)})^r$ .

We recall the following result from [Ngu05, DHS15].

**Lemma 1.** *Under the  $q$ -SDH assumption, the accumulator described in Figure 1 is collision-free and indistinguishable.*

**Subset Queries.** The WitGen algorithm described in Figure 1 is used to compute witnesses for the membership of single elements in  $\text{acc}_A$ . We now extend it to compute witnesses for multiple elements, namely to show that a set  $B \subseteq A$  is included in  $\text{acc}_A$ . We write  $\text{WitGen}^*(\boxed{sk}, ek, A, \text{acc}_A, r, B)$  for the computation of the witness  $\text{wit} = \text{acc}_A^{\frac{1}{\text{Ch}_B(x)}} = g^{r \text{Ch}_{A \setminus B}(x)}$ . Similarly, we let  $\text{Verify}^*(ek, \text{acc}_A, B, \text{wit})$  to return 1 in case  $e(\text{acc}_A, g) = e(\text{wit}, g^{\text{Ch}_B(x)})$  holds, and 0 otherwise.

Furthermore, we can extend WitGen to compute witnesses for an accumulator  $\text{acc}_B$  to accumulate a subset of the set accumulated into  $\text{acc}_A$ . Let  $r$  and  $r'$  be the randomness used to generate  $\text{acc}_A$  and  $\text{acc}_B$ , respectively. We define the following

- $\text{WitGen}^*(\boxed{sk}, ek, (\text{acc}_A, r, A), (r', B))$  : it computes the witness  $\text{wit}^* = \text{acc}_A^{\frac{1}{r' \text{Ch}_B(x)}} = g^{\frac{r}{r'} \text{Ch}_{A \setminus B}(x)}$
- $\text{Verify}^*(ek, \text{acc}_A, \text{acc}_B, \text{wit}^*)$  : it returns 1 if  $e(\text{acc}_A, g) = e(\text{acc}_B, \text{wit}^*)$  holds, and 0 otherwise.

### 2.3 Hard Relations

Let  $p$  be a polynomial and  $\mathcal{R}_{\text{pp}} \subseteq \{0, 1\}^{p(\lambda)} \times \{0, 1\}^{p(\lambda)}$  be a binary relation indexed by some public parameters  $\text{pp}$ . We call  $(u, w) \in \mathcal{R}$  *instance* and *witness*, respectively. We assume the public parameters  $\text{pp} \leftarrow \mathcal{G}(\lambda)$  to be efficiently computable given as input the security parameter. Also, let  $L_{\text{pp}} := \{u : \exists w \text{ s.t. } (u, w) \in \mathcal{R}\}$  to be the NP language corresponding to  $\mathcal{R}_{\text{pp}}$ . We require the language  $L$  to be efficiently sampleable and denote with  $u \leftarrow \mathcal{D}(L)$  the process of picking a random element from  $L$ . A relation  $(\mathcal{G}, \mathcal{R}, \mathcal{D})$  is said to be *hard* if for any probabilistic polynomial time adversary  $\mathcal{A}$  the following probability is negligible

$$\Pr \left[ \text{pp} \leftarrow \mathcal{G}(\lambda); u \leftarrow \mathcal{D}(L_{\text{pp}}); w \leftarrow \mathcal{A}(\text{pp}, u) : (u, w) \in \mathcal{R}_{\text{pp}} \right] \approx 0$$

More concretely, we are interested in relations corresponding to hard search problems associated with cryptographic accumulators. For example the relation corresponding to the following language

$$L_{ek}(a) := \{(\text{acc}_A, a) \in \mathbb{G} \times \mathbb{Z}_p : \exists \text{wit} \in \mathbb{G} \text{ s.t. } \text{Verify}(ek, \text{acc}_A, a, \text{wit}) = 1\}$$

The above language consists of all accumulators  $\text{acc}_A$  for which there exists a witness for the accumulation of  $a \in \mathbb{Z}_p$ . We note that the above language is efficiently sampleable by letting  $\text{acc}_A \leftarrow \text{Eval}(ek, a)$ . We now state the following straightforward Lemma and refer to Appendix B for the proof.

**Lemma 2.** *Assuming the accumulator is collision-free and indistinguishable, then the above the binary relation corresponding to  $L_{ek}(a)$  is hard for any  $a \in \mathbb{Z}_p$ .*

## 2.4 Smooth Projective Hash Function

Smooth Projective Hash Functions (SPHF) were introduced by Cramer and Shoup [CS02] (with the name of hash proof system) as a kind of designated-verifier proof systems for certain classes of algebraic languages. These found great applications towards the development of several primitives such as CCA2 secure public key encryption [CS02] and password authenticated key exchange [GL03, KV09]. Here we define a simpler hash proof system for the language of elements accumulated using the above bilinear accumulator.

An SPHF consists of three algorithms (HGen, Hash, PHash). The key generation algorithm HGen takes as input the security parameter and returns a relation<sup>1</sup>, and a pair of secret and public keys  $(hsk, hpk)$ ; we sometimes refer to  $hpk$  as the projection key. The keys specify an hash function from the relation  $\mathcal{R}$  to an abelian group  $\mathbb{G}$ . The hash function can be privately evaluated using  $hsk$  on any instance in  $L_{\mathcal{R}}$ , namely  $\text{Hash}_{hsk} : L_{\mathcal{R}} \rightarrow \mathbb{G}$ . The hash function allows also for public evaluation given  $hpk$  but only on instances for which a witness is known, namely  $\text{PHash}_{hpk} : \mathcal{R} \rightarrow \mathbb{G}$ . An SPHF satisfies two main properties: correctness and smoothness.

- Correctness: for any  $(u, w) \in R$ , the private and public evaluation algorithms Hash, PHash of the SPHF return the same result, i.e.

$$\text{Hash}_{hsk}(u) = \text{PHash}_{hpk}(u, w)$$

- (Computational) Smoothness: for any instance for which a witness is not known, the evaluation of the hash function is (computationally) indistinguishable from random. Namely, we say that an SPHF on a relation  $(\mathcal{G}, \mathcal{R}, \mathcal{D})$  is smooth if for any probabilistic polynomial time adversary  $\mathcal{A}$ , the following advantage is negligible

$$\left| \Pr \left[ \begin{array}{l} (L_{pp}, hsk, hpk) \leftarrow \text{HGen}(\lambda); \\ u \leftarrow \mathcal{D}(L_{pp}); \\ H \leftarrow \text{Hash}_{hsk}(u); \end{array} : \mathcal{A}(hpk, u, H) = 1 \right] - \Pr \left[ \begin{array}{l} (L_{pp}, hsk, hpk) \leftarrow \text{HGen}(L_{pp}); \\ u \leftarrow \mathcal{D}(L_{pp}); \\ H \leftarrow \mathbb{G}; \end{array} : \mathcal{A}(hpk, u, H) = 1 \right] \right| \approx 0$$

We now show the construction for an SPHF defined on the relation specified by  $L_{ek}(a)$ , for any  $a \in \mathbb{Z}_p$ , to the target group  $\mathbb{G}_T$  of a bilinear group. The construction of the SPHF is described in Figure 2 and is a simple combination of the bilinear accumulators of [Ngu05] and the verifiable random function (VRF) constructed by Dodis and Yampolskiy [DY05]. A VRF is a pseudorandom function which admits proofs of correct evaluation that can be publicly verified. In our SPHF we apply the VRF to an accumulator and an element accumulated in it. The proof of evaluation for the function corresponds to the accumulation witness, and the secret key of the SPHF is the secret key of the accumulator. Since the secret key of the accumulator allows to compute witnesses for every element in  $\mathbb{Z}_p$  it also allows to evaluate the SPHF in every pair  $(acc, a) \in \mathbb{G} \times \mathbb{Z}_p$ .

The security of the SPHF constructed in Figure 2 follows from the security of the verifiable random function of [DY05], based on the  $q$ -DHDBI assumption. We refer to Appendix B for a proof of the following Lemma.

<sup>1</sup>The original definition of SPHF was introduced for languages related with hard subset membership problems, while here we define SPHF for languages related with a hard search problem.

$\text{HGen}(\lambda) \rightarrow (L_{ek}, hsk, hpk):$ <ul style="list-style-type: none"> <li>• <math>(sk, ek) \leftarrow \text{KeyGen}(\lambda)</math></li> <li>• <math>L_{ek} := \cup_{a \in \mathbb{Z}_p} L_{ek}(a)</math></li> <li>• <math>z \leftarrow \mathbb{Z}_p^*</math></li> <li>• <math>hsk := (sk, z)</math></li> <li>• <math>hpk := (ek, g^z)</math></li> </ul>	$\text{Hash}_{hsk}(\text{acc}, a) \rightarrow H:$ <ul style="list-style-type: none"> <li>• <math>H := e(\text{acc}, g^z)^{\frac{1}{sk+a}}</math></li> </ul> $\text{PHash}_{hpk}((\text{acc}, a), \text{wit}) \rightarrow H:$ <ul style="list-style-type: none"> <li>• If <math>((\text{acc}, a), \text{wit}) \in \mathcal{R}_{L_{ek}}</math>:  <ul style="list-style-type: none"> <li>• <math>H := e(\text{wit}, g^z)</math></li> </ul> </li> <li>• Else: Return <math>H := \perp</math>.</li> </ul>
---	--

**Fig. 2:** SPHF for accumulators.

**Lemma 3.** *Under the  $q$ -DBDHI assumption over a bilinear group  $(p, \mathbb{G}, \mathbb{G}_T, e, g)$ , the construction in Figure 2 is a smooth projective hash function.*

### 3 Reactive PSI in the Augmented Semi-Honest Model

Aiming to prevent oracle attacks in scenarios where two parties engage in several PSI executions, we consider stateful PSI protocols computing reactive functionalities, whereby their outputs can depend on previous instances of the protocol.

We set our security definitions in the augmented semi-honest model of [Gol04]. In this model, the adversary is restricted to follow the specifications of the protocol as in the standard semi-honest settings. In addition, the adversary is allowed to adaptively modify the inputs used by the controlled party before each instance of the protocol. Apart from being more natural [HL10b] to give semi-honest adversaries this capability, we argue that the augmented model is more appropriate than the standard one to study composition of protocols.

Let  $t = t(\lambda)$  be a polynomial. We define the *reactive functionality*  $\text{ReF} = (F_1, F_2, \dots, F_t)$  as a sequence of stateful functionalities<sup>2</sup>  $F_i$  each taking as input a client set  $C_i$  and a server set  $S_i$  and returning a pair  $\text{ReF}_i(C_i, S_i) = (I_i, b_i)$ . These correspond to the outputs of the client and server should have at the end of each execution, respectively.

Next, we state our security definitions in terms of a generic reactive functionality and refer to the end of the section for specific instantiations of  $\text{ReF}$  for private set intersection protocols.

**Definition 1 (RePSI).** *A private set intersection protocol is a tuple  $(\text{Setup}, \Pi)$  s.t.*

- $\text{Setup}(\lambda) \rightarrow (\text{param}_C; \text{param}_S)$ : *it takes as input the security parameter and returns a pair of initial parameters for the client and the server. These can include public parameters and secret keys for the client and the server. If a specific protocol does not require a setup algorithm, this can be simply regarded as copying the security parameter into the initial parameters.*
- $\Pi(\mathcal{C}(C; St_C); \mathcal{S}(S; St_S)) \rightarrow ((\text{out}_C; St_C); (\text{out}_S; St_S))$ : *this is a stateful probabilistic polynomial time interactive protocol between a client  $\mathcal{C}$  and a server  $\mathcal{S}$ . Each party takes as input a set and a state information (initialised to  $St_C := \text{param}_C, St_S := \text{param}_S$  in the first instance of the protocol) and returns an output and an updated the state.*

<sup>2</sup>In this paper we restrict our attention to the case of deterministic functionalities.



We say that private set intersection protocol  $(Setup, \Pi)$  is a RePSI if it securely realizes a reactive functionality  $\text{ReF}$  in the augmented semi-honest model, i.e. if it satisfies correctness, server privacy and client privacy as defined below.

Correctness is defined by the security game  $\text{Exp}_A^{\text{Corr}}(\lambda)$  described in Figure 3. Informally, a protocol is correct if at the end of each instance both parties return their prescribed outputs.

**Definition 2 (Correctness).** Let  $t = t(\lambda)$  a polynomial in the security parameter  $\lambda$ , and  $\text{ReF}$  defined as above. A protocol  $(Setup, \Pi)$  is correct if for any probabilistic polynomial time adversary  $A$

$$\Pr \left[ \text{Exp}_A^{\text{Corr}}(\lambda) = 1 \right] \approx 1$$

**Exp<sub>A</sub><sup>Corr</sup>(λ):**

- $(\text{param}_C; \text{param}_S) \leftarrow \text{Setup}(\lambda)$
- $(C_1, S_1, \dots, C_t, S_t) \leftarrow \mathcal{A}(\text{param}_C; \text{param}_S)$
- $St_C := \text{param}_C, St_S := \text{param}_S$
- For  $i = 1$  to  $t$ :
  - $((\text{out}_{C,i}; St_C); (\text{out}_{S,i}; St_S)) \leftarrow \Pi(\mathcal{C}(C_i; St_C); \mathcal{S}(S_i; St_S))$
  - $(I_i, b_i) = \text{ReF}_i(C_i, S_i)$
- If  $(\text{out}_{C,i}, \text{out}_{S,i}) = (I_i, b_i)$  for all  $i \in [1, \dots, t]$ : Return 1
- Else: Return 0

**Fig. 3:** Correctness Game

Client privacy is specified by two oracles  $O_\Pi, O_{Sim}$  described in Figure 4. The oracle  $O_\Pi$  allows the adversary to run the next interaction between client and server on server's inputs of her choice. The oracle then returns the server's view in the protocol  $\text{view}_{i,S}((C_i; St_C); (S^*; St_S^*))$ , which contains the server's input, random coins and messages received from the client in the execution of the protocol. Oracle  $O_{Sim}$  returns instead a simulated view, based only on the input and output of the server. Informally, we say that the protocol achieves client privacy if an adversary is not able to distinguish which oracle she is interacting with.

**Definition 3 (Client Privacy).** Let  $t = t(\lambda)$  and  $\text{ReF}$  defined as above. A protocol  $(Setup, \Pi)$  has client privacy if for any probabilistic polynomial time adversary  $A$  there exists a probabilistic polynomial time simulator  $Sim$ , such that for every sequence  $S_1, \dots, S_t$  the following advantage is negligible

$$\text{Adv}_A^{\text{CPriv}}(\lambda) = \left| \Pr \left[ (\text{param}_C; \text{param}_S) \leftarrow \text{Setup}(\lambda) : \mathcal{A}^{O_\Pi}(\text{param}_S) = 1 \right] - \Pr \left[ (\text{param}_C; \text{param}_S) \leftarrow \text{Setup}(\lambda) : \mathcal{A}^{O_{Sim}}(\text{param}_S) = 1 \right] \right|$$

Server privacy is also specified in terms of two oracles  $O_\Pi, O_{Sim}$  described in Figure 5. The oracle  $O_\Pi$  allows the adversary to run the next interaction between client

$O_{\Pi}(S^*, St_S^*):$ <ul style="list-style-type: none"> <li>• If <math>i = 0</math>: <math>St_C := \text{param}_C</math></li> <li>• <math>i = i + 1</math></li> <li>• <math>((\text{out}_{C,i}; St_C); (\text{out}_{S,i}; St_S)) \leftarrow \Pi(\mathcal{C}(C_i; St_C); \mathcal{S}(S^*; St_S^*))</math></li> <li>• Return <math>\text{view}_{i,S}((C_i; St_C); (S^*; St_S^*))</math></li> </ul> $O_{Sim}(S^*, St_S^*):$ <ul style="list-style-type: none"> <li>• <math>i = i + 1</math></li> <li>• <math>(I_i; b_i) = \text{ReF}_i(C_i, S^*)</math></li> <li>• <math>\text{view}_{i,S,Sim} \leftarrow \text{Sim}((S^*, St_S^*), b_i, \text{param}_S,  C_i )</math></li> <li>• Return <math>\text{view}_{i,S,Sim}</math></li> </ul>
---

**Fig. 4:** Oracles used in the Client Privacy game.

$O_{\Pi}(C^*, St_C^*):$ <ul style="list-style-type: none"> <li>• If <math>i = 0</math>: <math>St_S := \text{param}_S</math></li> <li>• <math>i = i + 1</math></li> <li>• <math>((\text{out}_{C,i}; St_C); (\text{out}_{S,i}; St_S)) \leftarrow \Pi(\mathcal{C}(C^*; St_C^*); \mathcal{S}(S_i; St_S))</math></li> <li>• Return <math>\text{view}_{i,C}((C^*; St_C^*); (S_i; St_S))</math></li> </ul> $O_{Sim}(C^*, St_C^*):$ <ul style="list-style-type: none"> <li>• <math>i = i + 1</math></li> <li>• <math>(I_i; b_i) = \text{ReF}_i(C^*, S_i)</math></li> <li>• <math>\text{view}_{i,C,Sim} \leftarrow \text{Sim}((C^*, St_C^*), I_i, \text{param}_C,  S_i )</math></li> <li>• Return <math>\text{view}_{i,C,Sim}</math></li> </ul>
---

**Fig. 5:** Details of the oracles used in the Server Privacy game.

and server on client's inputs of her choice. The oracle then returns the client's view  $\text{view}_{i,C}$  in the protocol. Oracle  $O_{Sim}$  returns instead a simulated view, based only on the input and output of the client. Informally, we say that the protocol achieves server privacy if an adversary is not able to distinguish which oracle she is interacting with.

**Definition 4 (Server Privacy).** Let  $t = t(\lambda)$  and  $\text{ReF}$  defined as above. A protocol  $(Setup, \Pi)$  has server privacy if for any probabilistic polynomial time adversary  $\mathcal{A}$  there exists a probabilistic polynomial time simulator  $Sim$ , such that for every sequence  $C_1, \dots, C_t$  the following advantage is negligible

$$Adv_{\mathcal{A}}^{\text{SPriv}}(\lambda) = \left| \Pr \left[ (\text{param}_C; \text{param}_S) \leftarrow Setup(\lambda) : \mathcal{A}^{O_{\Pi}}(\text{param}_C) = 1 \right] \right. \\ \left. - \Pr \left[ (\text{param}_C; \text{param}_S) \leftarrow Setup(\lambda) : \mathcal{A}^{O_{Sim}}(\text{param}_C) = 1 \right] \right|$$

**Size-Hiding.** In the previous definitions of client and server privacy we gave the simulator the size of the honest party's input set. This captures the security of most protocols in which participants learn information about the size of the other party's input. However, in certain cases the size of the inputs represents confidential information which should not be leaked in a protocol execution. Protocols achieving this stronger property

are usually referred as *size-hiding* [ADT11]. To formalise size-hiding variants of client and server privacy it is sufficient to remove the size of the honest party’s input from the input of the simulator. Looking ahead to the next sections, our protocols achieve size-hiding only in the case of client privacy.

The above definitions are general enough to capture various types of reactive functionalities ReF. Moreover, they can also be used to formalise security for the sequential composition of standard PSI stateless protocols. In this case it is sufficient to replace ReF with  $t$  copies of the same functionality  $F$  and replace protocol  $\Pi$  with a stateless protocol that does not update the states  $(St_C, St_S)$ , which are initialised as  $(\text{param}_C, \text{param}_S)$ .

Next, we specify two functionalities, one static (PSI) and one reactive (RePSI), which we call *Bounded Input PSI* and *Input Controlling RePSI*, respectively.

**Bounded Input PSI.** A Bounded Input PSI limits the maximum size of the set the client can use in each instance of the protocol. More precisely, let  $R$  be a polynomial in the security parameter  $\lambda$ , a Bounded Input PSI  $(\text{PSI}_1, \text{PSI}_2, \dots, \text{PSI}_t)$  is defined as

$$\text{PSI}_i(C_i, S_i) = \begin{cases} (C_i \cap S_i; \perp) & \text{If } |C_i| \leq R \\ (\perp; \perp) & \text{Otherwise} \end{cases}$$

**Input Controlling RePSI.** An Input Controlling RePSI limits the number of maximum distinct elements a client can include in its sets across all the executions. In this case the server’s outputs is a predicate on whether the client has exceeded the allowed bound. More precisely, let  $R$  be a polynomial in the security parameter  $\lambda$ , an Input Controlling RePSI  $(\text{RePSI}_1, \text{RePSI}_2, \dots, \text{RePSI}_t)$  is defined as

$$\text{RePSI}_i(C_i, S_i) = \begin{cases} (C_i \cap S_i; 1) & \text{If } |\cup_{j \leq i} C_j| \leq R \\ (\perp; 0) & \text{Otherwise} \end{cases}$$

## 4 Bounded Input PSI

In this section we introduce our construction for a Bounded Input PSI. Bounded Input PSI allows client and server to compute the intersection of their private sets while imposing a bound  $R$  on the size of the client set at each execution of the protocol. Bounded Input PSI is not a reactive RePSI but we will use it as a stepping stone for constructing our Input Controlling RePSI in the next section.

We notice that in several PSI protocols the size of the client set is naturally revealed during the interaction. Hence, a Bounded Input PSI variant can be easily achieved with simple modifications. The server can check the number of inputs used by the client and abort in case it exceeds the bound. This strategy is not viable in size-hiding PSI protocols [ADT11] where the use of cryptographic accumulators hides the size of the client set. We also leverage cryptographic accumulators, thus the server cannot directly check the number of inputs used by the client as just explained. Moreover, we will start from the Bounded Input PSI introduced in this section to construct our Input Controlling RePSI. In the latter, apart from hiding the size of the client set and reducing

the communication, the use of accumulators will enable to use compact states for the server whose size does not depend on the number of protocol executions.

Our Bounded Input protocol is a modification of the bounded size-hiding protocol of Bradley et al. [BFT16], whose security is based on the  $q$ -SDH assumption in the random oracle model. The idea behind the protocol of Bradley et al. [BFT16] is to have the client to accumulate its input set using a cryptographic accumulator and send it to the server. The server would then use the accumulator secret key to *remove* her elements from the accumulator. This amounts to compute witnesses for elements in the server set. Then, the server hashes the witnesses using a random oracle and sends back the hash values to the client. The client is able to compute witnesses for each accumulated element and then hash them on the random oracle. The intersection can be then retrieved by checking matches between the two sets of hash values.

The protocol of [BFT16], as well as ours, relies on the boundedness of the underlying accumulator to limit the size of the sets that can be accumulated.

Informally, the protocol of [BFT16] fulfils server privacy because the random oracle hides all the information about the witnesses computed by the server, apart from the ones the client can compute on its own. We modify the protocol and remove the need of random oracles. The idea is to replace it with a function that can be efficiently computed by the client given a witness, but for which the evaluation looks random if a witness is not known. This is exactly the smoothness property of the SPHF we introduced in Section 2.4. Thus, we are able to remove the random oracle assumption and reduce the security to the  $q$ -DBDHI assumption, on which the SPHF relies on. We note that our Bounded Input PSI is, to the best of our knowledge, the first instantiation of size-hiding PSI in the standard model.

#### 4.1 Bounded Input PSI Without Random Oracles

The setup of the protocol of the bounded-size PSI consists of generating a pair of secret key and evaluation key for a bilinear accumulator, as shown in Figure 1. The length of the evaluation key  $ek$  of the accumulator matches the input bound  $R$  allowed to the client input size. The setup algorithm then sets the initial parameters for the client to be the evaluation key of the accumulator, and the initial parameters for the server to include both the secret key and evaluation key. The complete description of the *Setup* algorithm of our Bounded Input PSI is described in Figure 6. Note that since we are in semi-honest settings, we can allow the server to run the setup and send the initial parameters to the client in a preliminary interaction with the client.

In the first move of the protocol, the client starts by computing an accumulator  $acc_C$  of its input set  $C$  and send it to the server. The evaluation of the accumulator can be done efficiently by first computing the characteristic polynomial of the set  $C$ , expanding its coefficients, and then performing a multi-exponentiation of the evaluation key  $ek$ , using the coefficients of  $Ch_C(X)$  as exponents.

In the second move of the protocol, the server then picks a pair of keys  $(hsk, hpk)$  for a SPHF associated with the witness relation of the accumulator. The secret key for the SPHF consists of the secret key  $x$  of the accumulator and a random element  $z \leftarrow \mathbb{Z}_p^*$ . The projective key of the hash function corresponds to the accumulator  $acc_S$  of the server set, using randomness  $z$ . Then for every  $s_i$  contained in its input set  $S$ ,

$\text{Setup}(\lambda) \rightarrow (\text{param}_C; \text{param}_S):$ <ul style="list-style-type: none"> <li>• <math>(p, \mathbb{G}, \mathbb{G}_T, e, g) \leftarrow \mathcal{BG}(\lambda)</math></li> <li>• Set <math>R := R(\lambda)</math></li> <li>• <math>x \leftarrow \mathbb{Z}_p^*</math></li> <li>• <math>sk := x</math></li> <li>• <math>ek := (g, g^x, g^{x^2}, \dots, g^{x^R})</math></li> <li>• <math>\text{param}_C := ((p, \mathbb{G}, \mathbb{G}_T, e, g), ek)</math></li> <li>• <math>\text{param}_S := ((p, \mathbb{G}, \mathbb{G}_T, e, g), sk, ek)</math></li> </ul>
---

**Fig. 6:** Setup algorithm for Bounded Input PSI.

the server evaluates the SPHF on instances  $(\text{acc}_C, s_i)$  using the secret key  $x$ . The server ends its move by forwarding the projective key  $\text{acc}_S$  to the client together with the set of SPHF evaluations. Without loss of generality we assume the server to sort the set of evaluations in lexicographic order before sending it to the client. Note that the server is not strictly required to know the secret key of the accumulator. However this can be used to speed up computation. For example, the server can avoid to accumulate its own set and simply set  $\text{acc}_S = g^z$ .

In the last move of the protocol, the client computes a witnesses  $\text{wit}_i$  for the accumulation of his input elements  $c_i \in \mathbb{C}$  in the accumulator  $\text{acc}_C$ . Then, the client evaluates the SPHF using the projective key  $\text{acc}_S$  on each witness  $\text{wit}_i$  for  $(\text{acc}_C, c_i)$ . The clients then compares the set of its evaluations of the SPHF with the evaluations received from the server, looking for matches. Finally, the client outputs the subset of elements in  $\mathbb{C}$ , for which the evaluation of the SPHF gave a match.

The full description of our Bounded Input PSI protocol is given in Figure 7. We discuss its security in the following Theorem and refer to Appendix C for the proof.

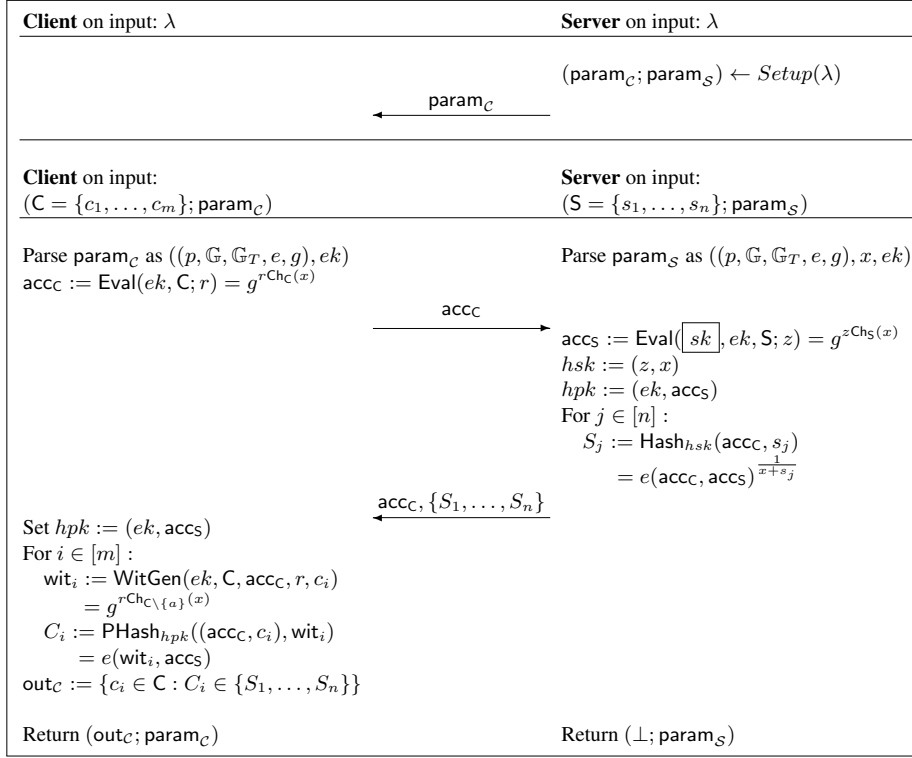
**Theorem 1.** *Under the  $R$ -SDH and  $n$ -DBDHI assumptions, the protocol (Setup, II) as described in Figures 6 and 7 is a secure instantiation of a Bounded Input PSI in the augmented semi-honest model.*

## 5 Input Controlling RePSI

We now introduce our Input Controlling RePSI protocol. The starting point is the Bounded Input protocol introduced in the previous section. The idea is to turn the previous protocol into a stateful one where both parties keep track of previous executions.

### 5.1 Description of the Protocol

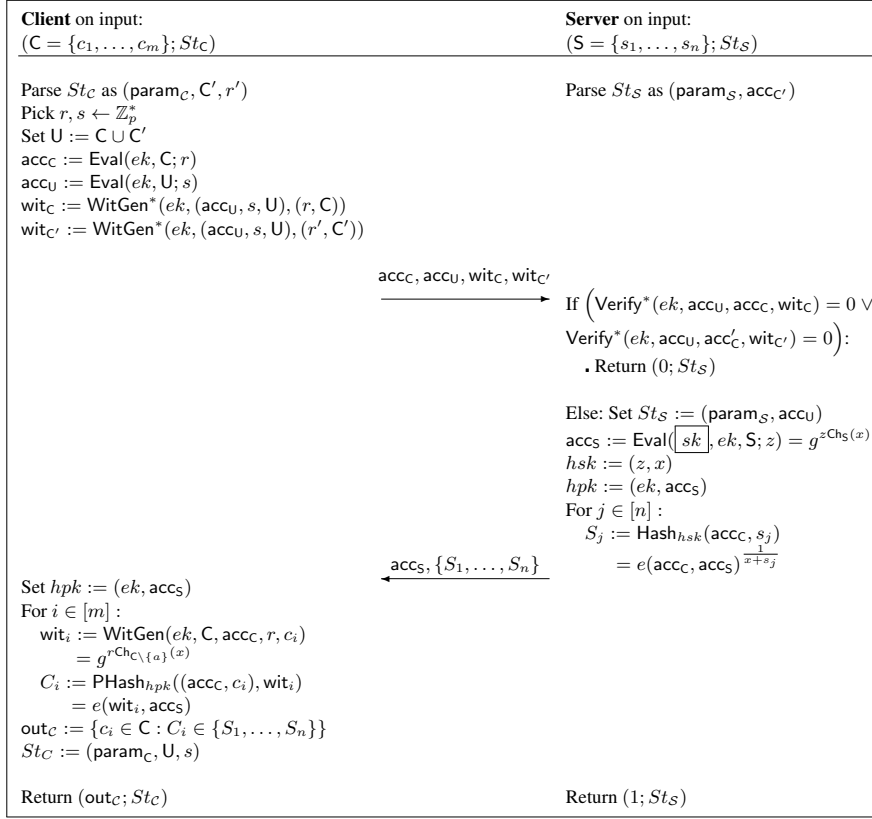
The *Setup* phase of the protocol is the same as the one described in Figure 6 for the Bounded Input protocol. We stress that in this case the bound  $R$  is not (only) the bound on the size of the client input of a single execution, but also a bound on the maximum number of elements the client can use across multiple executions. Again, since the client initial parameters only include public information we can allow the (semi-honest) server to run the *Setup* and forward the client the initial parameters  $\text{param}_C$ .



**Fig. 7:** Bounded Input PSI without random oracles.

The first instance of Input Controlling RePSI is similar to an execution of a Bounded Input RePSI described in Figure 7. The only difference here is that at the end of the first instance the client and the server update their output state. The client returns state  $St_C = (\text{param}_C, C, r)$ , which includes the initial parameter  $\text{param}_C$ , its current input set  $C$  and the randomness used to create the accumulator  $\text{acc}_C$ . The server returns state  $St_S = (\text{param}_S, \text{acc}_C)$ , which includes the initial parameters  $\text{param}_S$  as well as the accumulator  $\text{acc}_C$  received from the client. In the rest of the description we implicitly assume that initial parameters  $\text{param}_C$  and  $\text{param}_S$  are always part of the states  $St_C$  and  $St_S$ , respectively, and omit them from the notation to improve readability.

All the instances following the first one proceed as described in Figure 8. In the first move of the protocol, the client retrieves the set  $C'$  stored in state  $St_C$  which contains all the elements used in previous executions of the protocol and computes the union with its current input set, i.e.  $U = C \cup C'$ . Then, the client computes fresh accumulators for both the current input set  $C$  and the union set  $U$ . The client computes also witnesses for the accumulation in  $\text{acc}_U$  of subsets accumulated into  $\text{acc}_C$  and  $\text{acc}_{C'}$ . Here  $\text{acc}_{C'}$  corresponds to the accumulator of the union of all previous client input set, which was generated in the last execution. The client ends its move by sending the accumulators  $\text{acc}_C, \text{acc}_U$  and witnesses  $\text{wit}_C, \text{wit}_{C'}$  to the server.



**Fig. 8:** Input Controlling RePSI.

In the second move of the protocol, the server retrieves the accumulator  $\text{acc}_{C'}$  from its state, which contains the union of the sets of all previous client sets. Then, the server verifies the witnesses  $\text{wit}_C, \text{wit}_{C'}$  for the accumulation of  $C$  and  $C'$  in  $U$ . If any of these checks fails, then the server terminates the execution of the protocol with output  $(0; St_S)$ . Note that in this case both client and server do not update their states and might later enter a new instance of the protocol with different inputs. In case both checks pass, the server continues the execution as in the Bounded Input protocol: he sets a public and private key for the SPHF and evaluates it on instances  $(\text{acc}_C, s_i)$  for elements  $s_i \in S$ . The server ends its move by sending the public key for the SPHF and the set of evaluations to the client, updates its state with the accumulator  $\text{acc}_U$  and terminates its execution by outputting  $(1, St_S)$ .

In the last move of the protocol the client continues the execution as in the case of the Bounded Input RePSI. It computes witnesses for the accumulation of elements in  $C$  into  $\text{acc}_C$  and computes evaluations of the SPHF on these. Then it looks for matches between the sets of evaluations and includes the corresponding elements in  $C$  in the intersection  $\text{out}_C$ . The client updates its state with  $(U, s)$ , where  $s$  is the randomness

	Bounded Input	Input Controlling
Client Computation	$m\mathbb{P} + O(\frac{m^2}{\log m})\mathbb{E}$	$m\mathbb{P} + O(\frac{m^2}{\log m} + \frac{M}{\log M})\mathbb{E}$
Server Computation	$n\mathbb{P} + (n + 1)\mathbb{E}$	$(n + 4)\mathbb{P} + (n + 1)\mathbb{E}$
Communication	$n\mathbb{G}_T + 2\mathbb{G}$	$n\mathbb{G}_T + 5\mathbb{G}$
Client's State	n.a.	$1\mathbb{Z}_p$
Server's State	n.a.	$1\mathbb{G}$

**Table 1:** Efficiency. Computation is expressed in number of pairings  $\mathbb{P}$  and group exponentiations  $\mathbb{E}$ , communication in terms of the number of group elements  $\mathbb{G}$ , target group elements  $\mathbb{G}_T$  and field elements  $\mathbb{Z}_p$ . The size of the client and server sets are  $m$  and  $n$ , respectively. While the total size of the inputs used by the client up to the current iteration is denoted with  $M$ . The client's state does not include the total number of elements used by the client, i.e.  $M$ .

used in the generation of the accumulator  $\text{acc}_U$ , and terminates the instance execution with output  $(\text{out}_C; St_C)$ .

## 5.2 Security of Input Controlling RePSI

**Theorem 2.** *Under the R-SDH and n-DBDHI assumptions, the protocol (Setup, II) as described in Figures 6 and 8 is a secure instantiation of an Input Controlling RePSI in the augmented semi-honest model.*

*Proof.* Correctness as for the case of the Bounded Input RePSI follows from the correctness of accumulators and of SPHF.

Client privacy follows again from indistinguishability property of the accumulator. In the first instance of the protocol the only message the server receives from the client is the accumulator  $\text{acc}_C$ . In this case the simulator  $Sim$  picks  $r \leftarrow \mathbb{Z}_p$  and sets  $\text{acc}_C = g^r$ . As in the case of Bounded Input RePSI the simulated view is distributed identically to the real view. In the following calls of  $O_{Sim}$  the simulator picks  $s, t \leftarrow \mathbb{Z}_p$  and sets  $\text{acc}_U = g^s$  and  $\text{acc}_C = g^t$  and retrieves  $\text{acc}_{C'} = g^r$  from the previous instance. The simulator then sets the witnesses to be  $\text{wit}_C = g^{\frac{s}{t}}$  and  $\text{wit}_{C'} = g^{\frac{s}{r}}$ . The distribution of the simulated  $\text{acc}_C, \text{acc}_U, \text{wit}_C, \text{wit}_{C'}$  is uniformly random, conditioned on satisfying the two witness verification equations, as in a real distribution. Again, the simulator does not need the size of the client' set and thus client privacy is achieved with respect to the Size-Hiding variant.

The proof of server privacy unfolds as in the case of the Bounded Input RePSI (see Appendix C) since the messages sent from the server to the client are the same.

## 5.3 Efficiency

We summarize the efficiency of both our Bounded Input PSI and Input Controlling RePSI in Table 1. The dominant computational cost for the client in the Bounded Input PSI is  $O(m)$  pairings and multi-exponentiations of length at most  $m$ , where  $m$  is the size of the client's set. With respect to the Bounded Input PSI, the overhead incurred by the client in our Input Controlling RePSI is only of a single multi-exponentiation of length  $M$ , the total number of elements used so far in the protocol.



For both the Bounded Input PSI and the Input Controlling RePSI, the computational cost for the server is  $O(n)$  pairings and exponentiations, where  $n$  is the size of the server set in that run. The overall communication is linear in the size of the server set in that particular instance for both protocols. The communication overhead of the Input Controlling RePSI is of only 3 group elements more than the Bounded Input PSI.

The table also shows that both server and client keep constant state in case of Input Controlling RePSI. When computing the state size for the client, we *do not* consider the elements input thus far by the client. We argue that any instantiation of Input Controlling RePSI requires the client to include in its state the inputs thus far. This is because the client will have to tell whether the next input is “fresh” or not. This is also true for the instantiation of the protocol via a trusted third party. In this case, the trusted party will have to remember all inputs up to the current run, in order to tell whether the next input violates the bound. If we do not consider the inputs thus far as part of the client’s state, our instantiation of Input Controlling RePSI is optimal from the point of view of storage overhead since it only requires constant state at both the client and the server. In particular the state for each of them is independent of the number of runs.

**Comparison with [DMV13].** To the best of our knowledge, the only possible alternative to instantiate an Input Controlling RePSI would be to use the “rate-hiding” compiler of [DMV13] with a PSI protocol such as [HN10]. Since [DMV13] only hints at how to build a rate-hiding PSI,<sup>3</sup> we cannot compare its communication/computation complexity with the one of our Input Controlling RePSI. However, note that the rate-hiding compiler of [DMV13] requires the client to commit to the inputs of the current run and both parties to keep the commitments to the client’s inputs across all runs. (The client must prove that the number of unique inputs hidden by the commitments does not exceed the rate.). As such, even if we exclude the client’s input from its state, both the client and the server keep a state that is linear in the number of elements used by the client across all executions. Whereas, our protocol features constant state at both parties and communication complexity that is independent on the size of the client’s set.

## 6 Conclusions

Although a large number of Private Set Intersection (PSI) protocols have been proposed in recent years, their adoption in real-life applications is still hindered by a few challenges. In this paper, we focused on oracle attacks, whereby the client learns the server’s private set by enumerating it across several executions. To address this problem, we set out to model PSI as a reactive functionality, namely, Reactive PSI (RePSI), and provided a construction that allow the server to set an upper bound to the number of elements the client has input up to the current protocol run. Essentially, we made PSI a stateful protocol but provided a construction where the state kept by the two parties is small and independent of the number of runs thus far and, for the server only, independent on the number of elements in either input set.

---

<sup>3</sup>In [DMV13], the authors only show that the PSI protocol of [HN10] fulfills a property called “commit-first” and that it can be used in the rate-hiding compiler. However, the compiled protocol is not available.

To the best of our knowledge, our work is the first to formalize and instantiate Reactive PSI. In the process, we also presented the first size-hiding PSI protocol in the standard model, which may be of independent interest.

## References

- [ADT11] G. Ateniese, E. De Cristofaro, and G. Tsudik. (If) Size Matters: Size-Hiding Private Set Intersection. In *PKC*, pages 156–173, 2011.
- [BB04] D. Boneh and X. Boyen. Efficient selective-id secure identity-based encryption without random oracles. In *Eurocrypt*, pages 223–238, 2004.
- [BBD<sup>+</sup>11] P. Baldi, R. Baronio, E. De Cristofaro, P. Gasti, and G. Tsudik. Countering GAT-TACA: efficient and secure testing of fully-sequenced human genomes. In *ACM CCS*, pages 691–702, 2011.
- [BDM94] J. Benaloh and M. De Mare. One-way accumulators: A decentralized alternative to digital signatures. In *Eurocrypt*, pages 274–285, 1994.
- [BFT16] T. Bradley, S. Faber, and G. Tsudik. Bounded Size-Hiding Private Set Intersection. In *Security and Cryptography for Networks (SCN)*, pages 449–467, 2016.
- [BP97] N. Barić and B. Pfitzmann. Collision-free accumulators and fail-stop signature schemes without trees. In *Eurocrypt*, pages 480–494, 1997.
- [Che06] J. H. Cheon. Security analysis of the strong diffie-hellman problem. In *Eurocrypt*, pages 1–11, 2006.
- [CKS09] J. Camenisch, M. Kohlweiss, and C. Soriente. An accumulator based on bilinear maps and efficient revocation for anonymous credentials. In *PKC*, pages 481–500, 2009.
- [CL02] J. Camenisch and A. Lysyanskaya. Dynamic accumulators and application to efficient revocation of anonymous credentials. In *CRYPTO*, pages 61–76, 2002.
- [COV15] M. Chase, R. Ostrovsky, and I. Visconti. Executable Proofs, Input-Size Hiding Secure Computation and a New Ideal World. In *EUROCRYPT*, pages 532–560, 2015.
- [CS02] R. Cramer and V. Shoup. Universal hash proofs and a paradigm for adaptive chosen ciphertext secure public-key encryption. In *Eurocrypt*, pages 45–64, 2002.
- [CZ09] J. Camenisch and G. Zaverucha. Private intersection of certified sets. In *Financial Cryptography and Data Security*, 2009.
- [DHS15] D. Derler, C. Hanser, and D. Slamanig. Revisiting cryptographic accumulators, additional properties and relations to other primitives. In *Topics in Cryptology - CT-RSA 2015*, pages 127–144, 2015.
- [DKT10] E. De Cristofaro, J. Kim, and G. Tsudik. Linear-Complexity Private Set Intersection Protocols Secure in Malicious Model. In *Asiacrypt*, pages 213–231, 2010.
- [DMV13] Ö. Dagdelen, P. Mohassel, and D. Venturi. Rate-Limited Secure Function Evaluation: Definitions and Constructions. In *PKC*, pages 461–478, 2013.
- [DT10] E. De Cristofaro and G. Tsudik. Practical Private Set Intersection with Linear Complexity. In *Financial Cryptography and Data Security*, pages 143–159, 2010.
- [DY05] Y. Dodis and A. Yampolskiy. A verifiable random function with short proofs and keys. In *PKC*, pages 416–431, 2005.
- [FDCB15] J. Freudiger, E. De Cristofaro, and A. E. Brito. Controlled data sharing for collaborative predictive blacklisting. In *DIMVA*, pages 327–349, 2015.
- [FNP04] M. Freedman, K. Nissim, and B. Pinkas. Efficient private matching and set intersection. In *Eurocrypt*, pages 1–19, 2004.
- [GL03] R. Gennaro and Y. Lindell. A framework for password-based authenticated key exchange. In *Eurocrypt*, pages 524–543, 2003.

- [GMW87] O. Goldreich, S. Micali, and A. Wigderson. How to play any mental game or A completeness theorem for protocols with honest majority. In *Proceedings of the 19th Annual ACM Symposium on Theory of Computing*, pages 218–229, 1987.
- [Gol04] O. Goldreich. *The Foundations of Cryptography - Volume 2, Basic Applications*. Cambridge University Press, 2004.
- [GOP<sup>+</sup>16] E. Ghosh, O. Ohrimenko, D. Papadopoulos, R. Tamassia, and N. Triandopoulos. Zero-knowledge accumulators and set algebra. In *Eurocrypt*, pages 67–100, 2016.
- [HL08] C. Hazay and Y. Lindell. Efficient protocols for set intersection and pattern matching with security against malicious and covert adversaries. In *TCC*, pages 155–175, 2008.
- [HL10a] C. Hazay and Y. Lindell. *Efficient Secure Two-Party Protocols - Techniques and Constructions*. Information Security and Cryptography. Springer, 2010.
- [HL10b] C. Hazay and Y. Lindell. A note on the relation between the definitions of security for semi-honest and malicious adversaries. *ePrint*, 2010.
- [HN10] C. Hazay and K. Nissim. Efficient Set Operations in the Presence of Malicious Adversaries. In *PKC*, pages 312–331, 2010.
- [HOS17] P. Hallgren, C. Orlandi, and A. Sabelfeld. PrivatePool: Privacy-Preserving Ridesharing. In *CSF*, pages 276–291, 2017.
- [IKN<sup>+</sup>17] M. Ion, B. Kreuter, E. Nergiz, S. Patel, S. Saxena, K. Seth, D. Shanahan, and M. Yung. Private Intersection-Sum Protocol with Applications to Attributing Aggregate Ad Conversions. *ePrint 2017/738*, 2017.
- [JL10] S. Jarecki and X. Liu. Fast Secure Computation of Set Intersection. In *Security and Cryptography for Networks (SCN)*, pages 418–435, 2010.
- [KV09] J. Katz and V. Vaikuntanathan. Smooth projective hashing and password-based authenticated key exchange from lattices. In *Asiacrypt*, pages 636–652, 2009.
- [Ngu05] L. Nguyen. Accumulators from bilinear pairings and applications. In *CT-RSA*, pages 275–292, 2005.
- [PSSZ15] B. Pinkas, T. Schneider, G. Segev, and M. Zohner. Phasing: Private Set Intersection Using Permutation-based Hashing. In *USENIX Security Symposium*, 2015.
- [PSZ14] B. Pinkas, T. Schneider, and M. Zohner. Faster Private Set Intersection Based on OT Extension. In *USENIX Security Symposium*, pages 797–812, 2014.
- [RR17] P. Rindal and M. Rosulek. Improved private set intersection against malicious adversaries. In *CRYPTO*, pages 235–259, 2017.
- [ZSS04] F. Zhang, R. Safavi-Naini, and W. Susilo. An efficient signature scheme from bilinear pairings and its applications. In *PKC*, pages 277–290, 2004.

## Appendix

### A Assumptions

The exponent Strong Diffie-Hellman (Exponent  $q$ -SDH) [ZSS04] problem consists in computing the power  $g^{x^{q+1}}$  given all the powers  $g, g^x, \dots, g^{x^q}$ . A related problem is the Diffie-Hellman Inversion problem ( $q$ -DHI) [Che06] which asks to compute  $g^{1/x}$  given  $g, g^x, \dots, g^{x^q}$ . In [ZSS04], it is shown that the two problems are equivalent.

**Definition 5 (Exponent  $q$ -SDH).** We say that the Exponent  $q$ -SDH assumption holds in a group  $\mathbb{G}$  if for any probabilistic polynomial time algorithm  $\mathcal{A}$  the following probability is negligible

$$\Pr \left[ (p, \mathbb{G}, g) \leftarrow \mathcal{G}(\lambda); x \leftarrow \mathbb{Z}_p^* : \mathcal{A}(g, g^x, \dots, g^{x^q}) = g^{x^{q+1}} \right] \approx 0$$

The Decisional Bilinear Diffie-Hellman Inversion problem ( $q$ -DBDHI) [BB04, DY05] asks to distinguish  $e(g, g)^{\frac{1}{x}}$  from a random group element  $\Gamma \in \mathbb{G}_T$  having access to  $g, g^x, \dots, g^{x^q}$ .

**Definition 6 ( $q$ -DBDHI).** We say that the  $q$ -DBDHI assumption holds in  $(p, \mathbb{G}, \mathbb{G}_T, e, g)$  if for any probabilistic polynomial time algorithm  $\mathcal{A}$  the following advantage is negligible

$$\begin{aligned} Adv_{\mathcal{A}}^{q\text{-DBDHI}}(\lambda) = & \\ & \left| \Pr \left[ (p, \mathbb{G}, \mathbb{G}_T, e, g) \leftarrow \mathcal{BG}(\lambda); x \leftarrow \mathbb{Z}_p^* : \mathcal{A}(g, g^x, \dots, g^{x^q}, e(g, g)^{\frac{1}{x}}) = 1 \right] - \right. \\ & \left. \Pr \left[ (p, \mathbb{G}, \mathbb{G}_T, e, g) \leftarrow \mathcal{BG}(\lambda); x \leftarrow \mathbb{Z}_p^*, \Gamma \leftarrow \mathbb{G}_T : \mathcal{A}(g, g^x, \dots, g^{x^q}, \Gamma) = 1 \right] \right| \end{aligned}$$

We stress that the  $q$ -DBDHI assumption is stronger than the previous ones.

## B Accumulators and SPHF

In this Section we recall security definitions of cryptographic accumulators, Verifiable Random Functions and prove Lemmas 2 and 3.

### B.1 Security Definitions of Cryptographic Accumulators

The three main properties we require to a cryptographic accumulator are correctness, collision-freeness and indistinguishability. The following definitions are an adaptation of [DHS15] to our settings.

**Definition 7 (Static Accumulators).** An accumulator  $(\text{KeyGen}, \text{Eval}, \text{WitGen}, \text{Verify})$  is correct if for every probabilistic polynomial time adversary  $\mathcal{A}$  the following probability holds

$$\Pr \left[ \begin{array}{l} (sk, ek) \leftarrow \text{KeyGen}(\lambda); (A, a) \leftarrow \mathcal{A}(ek); \\ \text{acc}_A \leftarrow \text{Eval}(ek, A; r); \\ \text{wit} \leftarrow \text{WitGen}(ek, \text{acc}_A, r, A, a); \end{array} \begin{array}{l} a \in A, \\ \text{Verify}(ek, \text{acc}_A, a, \text{wit}) = 1 \end{array} \right] = 1$$

An accumulator  $(\text{KeyGen}, \text{Eval}, \text{WitGen}, \text{Verify})$  is collision-free if for every probabilistic polynomial time adversary  $\mathcal{A}$  the following probability is negligible

$$\Pr \left[ \begin{array}{l} (sk, ek) \leftarrow \text{KeyGen}(\lambda); (\text{acc}_A, r, A, a, \text{wit}) \leftarrow \mathcal{A}(ek); \\ \text{acc}_A = \text{Eval}(ek, A; r), \\ a \notin A \end{array} : \text{Verify}(ek, \text{acc}_A, a, \text{wit}) = 1 \right] \approx 0$$

An accumulator  $(\text{KeyGen}, \text{Eval}, \text{WitGen}, \text{Verify})$  is said indistinguishable if for every probabilistic polynomial time adversary  $\mathcal{A}$  the following advantage is negligible

$$\begin{aligned} Adv_{\mathcal{A}}^{\text{ind}}(\lambda) = & \\ & \left| \Pr \left[ \begin{array}{l} (sk, ek) \leftarrow \text{KeyGen}(\lambda); \\ (A_0, A_1) \leftarrow \mathcal{A}(ek); \\ \text{acc} \leftarrow \text{Eval}(ek, A_0) \end{array} : \mathcal{A}(\text{acc}) = 1 \right] - \Pr \left[ \begin{array}{l} (sk, ek) \leftarrow \text{KeyGen}(\lambda); \\ (A_0, A_1) \leftarrow \mathcal{A}(ek); \\ \text{acc} \leftarrow \text{Eval}(ek, A_1) \end{array} : \mathcal{A}(\text{acc}) = 1 \right] \right| \approx 0 \end{aligned}$$

## B.2 Proofs of Lemmas 2 and 3

We start showing that the language  $L_{ek}(a)$  is hard, i.e. that given an element  $a$  and an accumulator  $\text{acc}$  is (computationally) hard to compute a witness for the accumulation of  $a$  in  $\text{acc}$ .

**Lemma 2.** *Assuming the accumulator is collision-free and indistinguishable, then the above the binary relation corresponding to  $L_{ek}(a)$  is hard for any  $a \in \mathbb{Z}_p$ .*

*Proof.* We show that  $Adv_{\mathcal{A}}^{Hard} \leq Adv_{\mathcal{A}_1}^{CF} + Adv_{\mathcal{A}_2}^{ind}$ . An adversary  $\mathcal{A}$  for the hardness of the relation gets as input an instance consisting of an element  $a \in \mathbb{Z}_p$  and an accumulator  $\text{acc}_a \leftarrow \text{Eval}(ek, a)$  and returns a witness for the accumulation of  $a$  in  $\text{acc}$  with probability  $Adv_{\mathcal{A}}^{Hard}$ . If we replace the accumulator with  $\text{acc}_\emptyset \leftarrow \text{Eval}(ek, \emptyset)$ , the probability of success of  $\mathcal{A}$  can only be negligibly affected, otherwise the adversary can be used to break the indistinguishability of the accumulator. Since the adversary  $\mathcal{A}$  can be used to return a witness for an element that was not accumulated, we can construct an adversary  $\mathcal{A}_1$  that breaks collision freeness:  $\mathcal{A}_1$  simply accumulates the empty set using randomness  $r$  to obtain  $\text{acc}_\emptyset$ ; then it queries  $\mathcal{A}$  on producing a witness  $\text{wit}$  for the accumulation of an element  $a$  in  $\text{acc}_\emptyset$ , and returns  $(\text{acc}, r, \emptyset, a, \text{wit})$ . Anytime the adversary  $\mathcal{A}$  returns a witness, the adversary  $\mathcal{A}_1$  wins the collision-freeness game.

In the specific case of the accumulator of Figure 1, indistinguishability holds unconditionally and thus  $Adv_{\mathcal{A}_2}^{ind} = 0$ .

Next we show the construction of Figure 2 is an SPHF. The construction essentially applies a VRF of [DY05] (based on  $q$ -DBDHI assumption) to instances  $(\text{acc}, a) \in L_{ek}(a)$  and its security follows a similar argument. The security of the VRF imposes restrictions on the size of the domain of the function, which is required to be polynomially bounded. This inherently limits our SPHF which has to restrict the size of the domain of field elements we can evaluate our SPHF on. This is undesirable, however standard transformations can be applied to increase the size of the domain. For example [DY05] shows how a VRF can be combined with a collision resistance hash function to remove the restriction about the input size. This can be also adopted in our construction by requiring to first hash the values to be accumulated into a shorter domain and then to apply the SPHF to it. For application such as PSI, this does not necessarily represent a problem as actual inputs of the parties may be hashed at the beginning of the protocol. On the other hand, this requires to increase the size of the bilinear group beyond usual sizes to make up for the looseness of the reduction. We refer to [DY05] for a full discussion of the transformation as well as concrete suggestions for concrete parameters.

**Lemma 3.** *Under the  $q$ -DBDHI assumption over a bilinear group  $(p, \mathbb{G}, \mathbb{G}_T, e, g)$ , the construction in Figure 2 is a smooth projective hash function.*

*Proof.* Construction of Figure 2 is (perfectly) correct follows from the correctness of the accumulator. For smoothness, consider an adversary  $\mathcal{A}$  that for an element  $a$  is able to distinguish  $(\text{hpk}, (\text{acc}, a), \text{Hash}_{\text{hsk}}(\text{acc}, a))$  from  $(\text{hpk}, \text{acc}, a, H^*)$ , where  $H^* \leftarrow \mathbb{G}_T$ . We show how to use  $\mathcal{A}$  to construct an adversary  $\mathcal{B}$  against the  $q$ -DBDHI assumption. The input of  $\mathcal{B}$  consists of a  $q$ -DBDHI instance, i.e.  $(g, g^x, \dots, g^{x^q}, \Gamma)$ , where  $\Gamma$  is either a random element in  $\mathbb{G}_T$  or  $e(g, g)^{\frac{1}{x}}$ . The adversary starts setting

$y = x - a$  (for an unknown  $x$ ) and compute  $ek' = (g, g^y, g^{y^2}, \dots, g^{y^q})$ . Note that  $ek'$  can be computed without knowing  $y$  given  $ek$  and  $a$ , and that is identically distributed to a real evaluation key for an accumulator. The adversary then set  $hpk' = (ek', g^z)$  and sends  $\mathcal{A}$  the tuple  $(hpk', g^{z'}, a, \Gamma^{zz'})$  and returns the bit outputs by  $\mathcal{A}$ . Note that  $(hpk', g^{z'}, a)$  is identically distributed to  $(hpk', \text{acc} \leftarrow \text{Eval}(ek', a), a)$ . Moreover, in case  $\Gamma = e(g, g)^{\frac{1}{x}}$  then  $\Gamma^{zz'} = e(g^{z'}, g^z)^{\frac{1}{y+a}} = \text{Hash}_{hsk}(g^{z'}, a)$ , while in case  $\Gamma$  is a random element of  $\mathbb{G}_T$ , then so it is  $\Gamma^{zz'}$ .

## C Proof of Theorem 1 (Sketch)

We now reason about the security of the construction of our Bounded Input *PSI* presented in Section 4.

**Theorem 1.** *Under the  $R$ -SDH and  $n$ -DBDHI assumptions, the protocol (Setup,  $\Pi$ ) as described in Figures 6 and 7 is a secure instantiation of a Bounded Input *PSI* in the augmented semi-honest model.*

*Proof.* Correctness follows from the correctness of accumulators and of SPHF. For every element the client accumulates in its set the client can compute a witness of its accumulation. Since the two ways of evaluating the hash function return the same output, the client is able to recover any element in the intersection.

Client privacy follows from the indistinguishability property of the accumulator. The only message the server receives from the client is the accumulator  $\text{acc}_C$ . In each call of  $O_{Sim}$ , the simulator  $Sim$  can simulate the server's view by simply picking a random element  $\text{acc}_C \leftarrow \mathbb{G}$  and compute the rest of the view as an honest server would do. This includes the case where the adversary provides malformed input to the server which causes him to abort the execution. Perfect indistinguishability of the accumulator guarantees  $\text{acc}_C$  to be distributed uniformly at random, therefore the simulated and real views are identically distributed. Note that the simulator does not require the size of the client's set to compute the simulated view. Therefore client privacy holds with respect to the Size-Hiding variant of the definition.

We now argue about the security of server privacy based on the collision-free property of the accumulators and the smoothness of the SPHF. We recall that, from the security definition, in each call of  $O_{Sim}$  the simulator has access to the client's set  $C^*$  and state  $St_C^*$  as picked by the adversary  $\mathcal{A}$ , as well as  $I_i$ , the output of the (honest) computation. We distinguish the following cases for the simulator

- Case 1.** The client aborts on its first move: the simulator  $Sim$  runs the first move of the client on input the set  $C^*$  and state  $St_S^*$  provided by the adversary. If the client aborts, then the simulator terminates the execution with an abort message  $\perp$ . In this case the real and ideal distributions are identical. The client might abort in case the received inputs provided by the adversary are malformed. This includes for example the case where the size of the set is bigger than the length of the accumulator key, which is also provided by the adversary.
- Case 2.** The client does not abort on its first move: as in the previous case the simulator  $Sim$  runs the first move of the client, which produces an output message for the server. In this case we distinguish two sub-cases for the simulator.

- **Case 2a.** The output of the computation  $I_i \neq \perp$ : in this case we have that the size of  $|C^*| \leq R$ . The simulator  $Sim$  computes an accumulator  $acc_{I_i} = \text{Eval}(ek, I_i; z)$  by using the evaluation key  $ek$  as included in the initial parameters  $\text{param}_C$ . Then the simulator computes witnesses  $wit_j$  for the accumulation of element  $c_j \in I_i$  and the following target group elements  $S_k = e(acc_C, wit_j)$  for  $k \in \{1, \dots, |I_i|\}$ . The simulator then sets  $hpk = (ek, acc_{I_i})$  and picks the rest of the evaluations of the SPHF as random target group elements:  $S_k \leftarrow \mathbb{G}_T$  for  $k \in \{|I_i| + 1, \dots, n\}$ . The simulator then executes the second move of the client on input  $hpk, S_1, \dots, S_n$  and returns the entire view of the client. The simulated view is indistinguishable from the real view by the smoothness property of the SPHF: the  $S_k$  are distributed uniformly at random, conditioned on the ones for which the client holds a witness. Also the public key of the SPHF has the same distribution as an honestly generated key. Given an adversary which is able to distinguish between real and simulated view, one can use a standard hybrid argument to reduce the security to the smoothness of the SPHF. Which in turns gives a reduction to the  $n$ -DBDHI problem, where  $n$  is a bound on the size of the server set.
- **Case 2b.** The output of the computation is  $I_i = \perp$ : in this case  $|C^*| > R$ . While the input set used is bigger than the allowed bound, the client did not abort during the execution of the first move. This means that the evaluation key provided by the adversary  $\mathcal{A}$  in the state  $St_C^*$  is long enough for the set  $C^*$ . Using a different accumulator key would not allow the client to be able to compute witnesses for the SPHF. Thus in the simulation  $Sim$  simply returns random group elements from  $S_k \leftarrow \mathbb{G}_T$  and  $hpk = (ek, g')$  for a uniform  $g' \leftarrow \mathbb{G}$ . In this case the simulated view might differ from the real view, but the only way the adversary can distinguish between the two views is if she has a witness for evaluating the SPHF. However computing such a witness translates into a breach of the collision-free property of the accumulator scheme, which can then be used to construct an adversary against the  $R$ -SDH assumption.