

# Improved Private Set Intersection for Sets with Small Entries

Dung Bui<sup>1</sup> and Geoffroy Couteau<sup>2</sup>

<sup>1</sup> IRIF, Université Paris Cité  
bui@irif.fr

<sup>2</sup> CNRS, IRIF, Université Paris Cité  
couteau@irif.fr

**Abstract.** We introduce new protocols for private set intersection (PSI), building upon recent constructions of pseudorandom correlation generators, such as vector-OLE and ring-OLE. Our new constructions improve over the state of the art on several aspects, and perform especially well in the setting where the parties have databases with small entries. We obtain three main contributions:

1. We introduce a new semi-honest PSI protocol that combines subfield vector-OLE with hash-based PSI. Our protocol is the first PSI protocol to achieve communication complexity *independent* of the computational security parameter  $\kappa$ , and has communication lower than all previous known protocols for input sizes  $\ell$  below 70 bits.
2. We enhance the security of our protocol to the malicious setting, using two different approaches. In particular, we show that applying the *dual execution technique* yields a malicious PSI whose communication remains independent of  $\kappa$ , and improves over all known PSI protocols for small values of  $\ell$ .
3. As most previous protocols, our above protocols are in the random oracle model. We introduce a third protocol which relies on subfield ring-OLE to achieve maliciously secure PSI in the *standard model*, under the ring-LPN assumption. Our protocol enjoys extremely low communication, reasonable computation, and standard model security. Furthermore, it is *batchable*: the message of a client can be reused to compute the intersection of their set with that of multiple servers, yielding further reduction in the overall amortized communication.

## 1 Introduction

*Private Set Intersection (PSI)* is a cryptographic primitive that allows parties to jointly compute the set of all common elements between their datasets, without leaking any value outside of the intersection. It is a special case of secure multi-party computation (MPC). PSI enjoys a wide array of real-life applications; it is perhaps the most actively researched concrete functionality in secure computation, and has been the target of a tremendous number of works, see [10, 13, 18–24, 26–28] and references therein for a sample. As a consequence

of this intense research effort, modern PSI protocols now achieve impressive efficiency features, communicating only a few hundred bits per database items, and processing millions of items in seconds.

**Improving PSI with pseudorandom correlation generators.** Pseudorandom correlation generators (PCG) have been introduced in the works of [3, 5, 8] and have been the subject of a long and fruitful line of work [3–8, 11, 30, 32, 34]. At a high level, a PCG allows two parties to securely stretch long pseudorandom correlated strings from short, correlated seeds. Securely sharing correlated random strings is a crucial component in most modern secure computation protocols, which operate in the preprocessing model; PCG allows to realize this functionality with almost no communication. Among their many applications, PCGs allow to construct *silent oblivious transfer extension* protocols [4], which can realize (pseudorandom) OT extension with minimal (logarithmic) communication.

Since the top-performing PSI protocols rely on efficient OT extension, using PCG-based techniques to improve their efficiency is a natural idea. And indeed, this was done recently for OKVS-based PSI in [27], leading to the most efficient PSI protocol known to date (OKVS stands for oblivious key-value store [13]; the use of OKVS is the leading paradigm for the design of PSI protocols). To give a single datapoint, computing the intersection between two databases of size  $n = 2^{20}$  with the protocol of [27] communicates as little as  $426n$  bits in total. In addition, some of the tools used in [27] have been significantly improved since: replacing their OKVS (which is the PaXoS OKVS of [21]) by the more recent 3H-GCT OKVS of [13], and replacing their PCG (which is the one from [32]) by the recent PCG of [11], the cost goes down to an impressive  $247n$  bits of total communication. In comparison, even the *insecure* approach of exchanging the hashes of all items in the databases already requires  $160n$  bits of communication. OKVS-based PSI protocols are now firmly established as the leading paradigm in the field, and the use of PCGs to reduce their communication overhead even more seems to further widen the gap with the other paradigms.

### 1.1 Our Contributions

We thoroughly investigate how the use pseudorandom correlation generators can reduce communication in PSI protocols. We obtain several contributions:

- A new family of semi-honest hash-based PSI protocols. Our protocols can be instantiated using several hashing techniques, and achieve very low communication, especially for databases whose entries have a small bitlength.
- New maliciously secure hash-based PSI protocols. Here, interestingly, we revive the dual execution technique, which had been used previously to design malicious PSI protocols in [26], but was considered outdated. We show that, combined with our new approach, it leads to very competitive protocols, which achieve lower communication than all known alternatives for databases with small entries.
- Eventually, we design a new maliciously secure polynomial-based PSI protocol. Our protocol enjoys several powerful features: competitive communication, security in the standard model under the ring-LPN assumption (in

contrast, other maliciously secure PSI use the ROM), and the possibility for a client to publish a single encoding of its database, and later retrieve the intersection of its database with that of multiple servers independently, with a single server-to-client message, plus minimal (database-independent) additional communication.

Below, we elaborate on each of our contributions.

**Low communication PSI for databases with small entries.** Modern PSI protocols have communication  $O(\kappa \cdot n)$ , where  $n$  is the database size, and  $\kappa$  is a computational security parameter. More precisely, the receiver-to-sender communication is  $O(\kappa n)$ , while the sender-to-receiver communication is  $O(\lambda \cdot n)$ , where  $\lambda$  is a *statistical* security parameter (typically,  $\kappa = 128$  and  $\lambda = 40$ ). We introduce a new protocol, that combines hashing techniques (e.g. Cuckoo hashing or its variants, as initially used in [18]) with a new PCG-based oblivious pseudorandom function (OPRF). In contrast to all previous works, our work avoids the  $O(\kappa \cdot n)$  overhead: it reduces the receiver-to-sender communication to be roughly  $\ell \cdot n$  (where  $\ell$  is the bitsize of the database items), leading to a significant reduction in the overall communication. To our knowledge, our protocol is the first to achieve communication independent of  $\kappa$  (up to low order terms). To give a datapoint, for  $n = 2^{20}$ , with 64-bit entries, our protocol communicates  $210n$  bits, and with 32-bit entries, it communicates only  $148n$  bits. For the same parameters, the leading OKVS-based PSI of [27] communicates  $197n$  bits, even after improving it with all relevant optimization (such as using the 3H-GCT OKVS of [13], and the recent PCG of [11]). We provide further datapoints and comparisons to the state of the art on Table 1, when instantiating our protocols with various hashing methods.

**Fast maliciously-secure PSI for small entries.** We then turn our attention to maliciously secure PSI. We provide two alternative protocols which achieve malicious security; both use standard paradigms for upgrading PSI to malicious security. The first protocol combines our new PCG-based OPRF with simple hashing, and applies the standard paradigm used in most previous OKVS-based PSI to achieve malicious security (e.g. [27]). This requires to increase the sender-to-receiver message length, from  $O(\lambda \cdot n)$  to  $O(\kappa \cdot n)$  ( $\lambda$  is a statistical security parameter,  $\kappa$  is a computational security parameter; typically,  $\lambda = 40$  and  $\kappa = 128$ ) to allow for extraction of the sender input. Along the way, we also notice a small mistake in the parameter choices of [27]: they devise a new ROM-based extraction strategy in the malicious setting, and prove that a  $Q$ -query adversary will make extraction fail with probability bounded  $Q \cdot n / 2^\kappa$  (this is the probability that one of the  $Q$  queries of the malicious receiver collides with an element of the sender set). This implies that, to target 128 bits of computational security, one must set  $\kappa = 128 + \log n$ . However, the numbers reported in [27] correspond to choosing  $\kappa = 128$  at the 128-bit security level. We took this minor inconsistency into account in our tables.

More interestingly, our second protocol applies *dual execution* [26] to our PCG-based protocol with simple hashing. We observe that, in our context, this

**Table 1.** Comparison of the communication cost of several PSI protocols in the semi-honest setting and in the malicious setting, for various choices of the database size  $n$  (we assume that both parties have a database of the same size).  $\ell$  denote the bit-length of the inputs in the database; we set the computational security parameter  $\kappa$  to 128 and the statistical security parameter  $\lambda$  to 40 (for usual applications) or 30 (which can be suitable for lower risk applications). For all protocols, we take into account the optimization of [31] which reduces the costs of sending  $n$  elements of bitlength  $\lambda + 2 \cdot \log n$  to  $n \cdot (\lambda + \log n)$ . GCH stands for Generalized Cuckoo hashing (here, with 2 hash functions and 3 items per bin), 2CH for 2-choice hashing, and SH for simple hashing ( $N$  is the number of bins).

	$n = 2^{14}$	$n = 2^{16}$	$n = 2^{20}$	$n = 2^{24}$
<b>Semi-honest setting</b>				
KKRT16 [18]	$930n$	$936n$	$948n$	$960n$
PRTY19 [20] low*	$491n$	$493n$	$493n$	$494n$
PRTY19 [20] fast*	$560n$	$571n$	$579n$	$587n$
CM20 [10]	$668n$	$662n$	$674n$	$676n$
PRTY20 [21]	$1244n$	$1192n$	$1248n$	$1278n$
RS21 [27]	$2024n$	$898n$	$406n$	$374n$
RS21 [27] enhanced**	$280n$	$260n$	$263n$	$275n$
Ours ( $\ell = 64$ , GCH)	$246n$	$220n$	$210n$	$209n$
Ours ( $\ell = 48$ , GCH)	$215n$	$189n$	$179n$	$178n$
Ours ( $\ell = 32$ , GCH)	$184n$	$158n$	$148n$	$147n$
Ours ( $\ell = 64$ , 2CH)	$214n$	$190n$	$183n$	$185n$
Ours ( $\ell = 48$ , 2CH)	$193n$	$169n$	$162n$	$164n$
Ours ( $\ell = 32$ , 2CH)	$171n$	$148n$	$141n$	$142n$
Ours ( $\ell = 64$ , SH, $N = n/10$ )	$332n$	$302n$	$284n$	$276n$
Ours ( $\ell = 48$ , SH, $N = n/10$ )	$261n$	$230n$	$209n$	$198n$
Ours ( $\ell = 32$ , SH, $N = n/10$ )	$191n$	$158n$	$133n$	$120n$
Ours ( $\ell = 64$ , SH, $N = 1$ ) ***	$154n$	$131n$	$125n$	$128n$
Ours ( $\ell = 48$ , SH, $N = 1$ ) ***	$138n$	$115n$	$109n$	$112n$
Ours ( $\ell = 32$ , SH, $N = 1$ ) ***	$122n$	$99n$	$93n$	$96n$
<b>Malicious setting</b>				
RS21 [27] enhanced**	$343n$	$320n$	$315n$	$318n$
Ours ( $\ell = 48$ , SH, $N = n/10$ )	$430n$	$393n$	$356n$	$332n$
Ours ( $\ell = 40$ , SH, $N = n/10$ )	$359n$	$321n$	$281n$	$253n$
Ours ( $\ell = 32$ , SH, $N = n/10$ )	$289n$	$249n$	$205n$	$175n$

\* PRTY19 has two variants, SpOT-low (lowest communication, higher computation) and SpOT-fast (higher communication, better computation). Both use expensive polynomial interpolation and require significantly more computation compared to all other protocols in this table.

\*\* Using the 3H-GCT OKVS of [13] instead of PaXoS, and the VOLE of [11] instead of the one from [32]. Setting  $\kappa_{RS21}$  to  $\kappa + \log n$  to achieve  $\kappa$  bits of security.

\*\*\* Using  $N = 1$  requires an expensive degree- $n$  polynomial interpolation.

allows to achieve malicious security without having to increase the length of the sender-to-receiver message, at the cost of increasing the receiver-to-sender communication by a factor 2. Since our approach makes this communication as low as  $O(\ell \cdot n)$ , this turns out to be an excellent tradeoff whenever the database entries are not too large. Therefore, our results show that the landscape of maliciously secure PSI is more subtle than previously thought: for large entries, the standard approach still dominates, but for smaller entries (e.g.  $\ell \leq 40$ ), the dual execution technique leads to better performances. This revives the dual execution technique, which was previously considered obsolete compared to the modern alternatives.

**Efficient PSI in the standard model.** Eventually, our last contribution is a new “polynomial-based” PSI protocol that does not rely on the random oracle model, following the high level structure of previous works [14, 15, 17]. To this end, we introduce the notion of PCG for the *subfield ring-OLE* correlation, and show how a simple variant of the recent PCG for ring-OLE of [7] leads to efficient instantiations of this primitive. Then, we describe a new PSI protocol built on top of this PCG, which enjoys a number of very interesting features.

*Security features.* Our PSI protocol is in the standard model: unlike our first protocol, it does not require the random oracle model, or any tailor-made correlation-robustness assumptions. We rely solely on the (relatively well-established) ring-LPN assumption over polynomial rings with irreducible polynomials. To our knowledge, our protocol is the first standard model protocol which offers competitive performances compared to protocols using the random oracle heuristic or tailored assumptions. Furthermore, our PSI protocol enjoys full malicious security (for both parties) *almost for free*. This stems from the use of PCGs, which allows to confine the “price” of achieving malicious security to the distributed seed generation only, which has logarithmic communication and computation (in the set size  $n$ ).

We note that, though malicious security comes for free communication- and computation-wise, the tweaks used to guarantee malicious security in our protocol are not straightforward. In fact, achieving malicious security efficiently in polynomial-based PSI protocols is known to be complex and error prone. For example, previous works [14] used a superficially similar approach and claimed malicious security, but their protocol was found to be insecure in a recent preprint, which described powerful concrete attacks on this proposal [1]. Leveraging the specific structure of our protocol, we manage to get around these nontrivial subtleties with careful structural checks, for a minimal cost (independent of the database size).

*Efficiency features.* Our PSI protocol enjoys a very low communication, considerably lower than all previous PSI protocols in the standard model which we are aware of (excluding iO- or FHE-based protocol, which can have very low communication but poor concrete efficiency). In fact, communication-wise, our PSI protocol is even on par with the best *ROM-based* PSI protocols of previous

works. Concretely, for sets of size  $n$  with  $\ell$ -bit entries, our protocol communicates  $(2\ell + 3\lambda + 3 \log n) \cdot n + o(n)$  bits. To give a single datapoint, for  $\ell = 32$  and  $n = 2^{20}$ , we estimate the total communication to be  $278n$  bits. This is on par with the best maliciously secure protocol [27], which communicates  $279n$  bits in the same setting, with comparable computation (it also uses polynomial interpolation), but without standard model security.

On Table 1.1, we compare our protocol to the current fastest maliciously secure PSI protocols [21, 27, 29]. As the table shows, the communication of our protocol is almost on par with that of the best protocol (the protocol of [27], enhanced with the latest VOLE protocol) for small-ish input size, and large enough set sizes. Yet, our protocol is in the standard model under the ring-LPN assumption, while [27] is only proven secure in the ROM.

**Table 2.** Comparison of the communication cost of several PSI protocols in the malicious model, for various choices of the database size  $n$  (we assume that both parties have a database of the same size) and statistical security parameter  $\lambda = 40$ , using the encoding technique of [31].  $\ell$  denote the bit-length of the inputs in the database; we set the computational security parameter  $\kappa$  to 128. For fairness of comparison, since our standard model PSI uses interpolation, we compare it to RS21 with an interpolation-based OKVS (which has better communication), and we compare our other PSIs with RS21 instantiated with (computationally) efficient OKVS.

Protocol	Communication					Hardness Assumption	Standard Model
	$n = 2^{16}$	$n = 2^{18}$	$n = 2^{20}$	$n = 2^{22}$	$n = 2^{24}$		
Our Standard PSI						Ring-LPN	✓
$\ell = 64$	$724n$	$423n$	$342n$	$324n$	$323n$	+ OT	
$\ell = 48$	$692n$	$391n$	$310n$	$292n$	$291n$		
$\ell = 32$	$660n$	$359n$	$278n$	$260n$	$259n$		
RS21 [27] enhanced*	$318n$	$286n$	$279n$	$279n$	$280n$	LPN + OT	✗
Our Direct PSI							✗
$\ell = 64$	$421n$	$385n$	$374n$	$369n$	$365n$		
$\ell = 48$	$348n$	$311n$	$298n$	$292n$	$286n$		
$\ell = 32$	$277n$	$237n$	$223n$	$215n$	$208n$	LPN + OT	
Our Dual PSI							
$\ell = 64$	$609n$	$535n$	$511n$	$499n$	$489n$		
$\ell = 48$	$465n$	$388n$	$361n$	$345n$	$333n$		
$\ell = 32$	$321n$	$240n$	$210n$	$192n$	$176n$		
PRTY20 [21]			$1766n$			OT	✗
RT21 [29]			$512n$			DH	✗
RS21 [27] enhanced**	$320n$	$315n$	$315n$	$317n$	$318n$	LPN + OT	✗

\* Using interpolation instead of PaXoS, and the VOLE of [11] instead of the one from [32]. Sets  $\kappa_{\text{RS21}}$  to  $\kappa + \log n$  to achieve  $\kappa$  bits of security.

\*\* Using the new OKVS of [13] instead of PaXoS, and the VOLE of [11] instead of the one from [32]. Sets  $\kappa_{\text{RS21}}$  to  $\kappa + \log n$  to achieve  $\kappa$  bits of security.

*Batch non-interactive PSI.* On top of these security and efficiency features, the structure of our protocol allows to obtain a powerful interaction pattern: it leads to a batch non-interactive PSI, where after a short interaction with each server, a client  $C$  with set  $X$  can broadcast a *single* encoding of its database, and receive afterwards at anytime a single message from each server  $S_i$  with set  $X_i$  (plus, in the malicious setting, a small database-size-independent 2-round structural check), from which they can decode  $X \cap X_i$ . To achieve this feature, we build upon the fact that the PCG for subfield ring-OLE correlations is *programmable*, which means that we can enforce that a target party will receive the same pseudorandom string across executions with many different parties. Concretely, we achieve the following form of *batch non-interactive PSI* between a client  $C$  with database  $X$  and multiple servers  $S_i$  with datasets  $X_i$  (all of size  $n$ ):

1. In a preprocessing phase,  $C$  interacts with each of the servers, using  $O(\log n)$  communication *and* computation in each interaction, in a small constant number of rounds.
2. Then,  $C$  performs a single  $\tilde{O}(n)$  cost local computation, and broadcasts a single  $2\ell n$ -size *encoding*  $E_X$  of  $X$ .
3. Each server  $S_i$  can, at any time, send a single message  $M_i = m(X_i, E_X)$ , of length  $3(\lambda + \log n)n$ , using  $\tilde{O}(n)$  computation.
4. Eventually, given  $X$  and  $M_i$ , the client  $C$  can run a  $\tilde{O}(n)$  cost decoding procedure and recover  $X \cap X_i$ , without further interaction.

When the number of servers becomes large, our batch PSI protocol leads to strong savings for the client compared to executing a PSI protocol individually with each server. Furthermore, in this setting, the amortized communication (per PSI instance) is reduced to  $(2\ell/N_S + 3\lambda + \log n) \cdot n + o(n)$ , where  $N_S$  denotes the number of servers. Even for relatively small number of servers, the amortized communication quickly outperforms that of even the best ROM-based maliciously secure PSI protocols. For example, for  $n = 2^{24}$  and  $\ell = 32$ , the amortized communication per secure set intersection approaches  $195n$  bits with our protocol, versus  $280n$  for [27].

## 1.2 Concurrent work

In a concurrent and independent work, recently accepted at CCS'22, Rindal and Raghuraman [25] introduced a new PSI protocol, using an approach similar to ours: the authors also leveraged subfield-VOLE to achieve communication independent of the computational security parameter  $\kappa$ . Our results have been obtained independently of theirs, around the same time period. Although their main result bears similarities to our first two contributions, we highlight some important distinctions between our work and theirs:

- The work of [25] uses an OKVS-based construction, and achieves a receiver-to-sender communication of  $(\lambda + 2 \log n) \cdot n$ . In contrast, we use a hash-based protocol, and achieve an  $(\ell - \log n) \cdot n$  receiver-to-sender communication. Therefore, we get smaller communication overall in the setting where the databases have small entries, but a slightly larger computation.

- For malicious security, the work of [25] only considers the standard paradigm of previous works (e.g. [27]), hence having a  $O(\kappa \cdot n)$  receiver-to-sender (and overall) communication. In contrast, we give two protocols, including one based on dual execution which achieves communication independent of  $\kappa$  (and smaller concrete communication for databases with small entries).
- Eventually, our last contribution, a “batchable” ring-OLE-based malicious PSI in the standard model with low communication, is unique to our work.

### 1.3 Structure of the Paper

We provide preliminaries in Section 2, and a detailed technical overview of our contributions in Section 3. Section 4 covers our ROM-based semi-honest and malicious protocols. Due to space limitation, our second malicious protocol, based on dual execution, is presented in the full version [9]. Section 5 covers our standard model PSI. Note that the additional preliminaries and all the missing proofs appear in full version [9].

## 2 Preliminaries

**Notation.** Throughout the paper we use the following notations: we let  $\kappa, \lambda$  denote the computational and statistical security parameters, respectively. We write  $[1, m]$  to denote a set  $\{1, 2, \dots, m\}$ . For a vector  $\mathbf{x}$  we define by  $x_i$  its  $i$ -th coordinate. Given distribution ensembles  $\{X_n\}, \{Y_n\}$ , we write  $X_n \approx Y_n$  to denote that  $X_n$  is computationally indistinguishable to  $Y_n$ .

We typically write  $\mathbb{F}_q$  to denote a field with an arbitrary subfield  $\mathbb{F}_p$ , where  $p$  is a prime power and  $q = p^t$ . We use  $\mathcal{R}_p = \mathbb{F}_p[X]/F(X)$  for the ring over the field  $\mathbb{F}_p$  where  $F(x)$  is some polynomial, and also denote  $\mathcal{R}_q = \mathbb{F}_{p^t}[X]/F(X)$ . Note that all operations in our paper are field/ring operations not modular arithmetic.

**PSI functionality.** A private set intersection (PSI) protocol allows two parties to compute the intersection of their input sets while concealing all other information. We typically denote by  $n$  the input set sizes. For completeness, the ideal functionalities for PSI (in the semi-honest and in the malicious settings) are given in Appendix of the full version [9].

**Pseudorandom correlation generators (PCG).** Pseudorandom correlations generators have been introduced in a recent line of work [3–5]. A PCG allows to compress long correlations into short, correlated seeds that can later be locally expanded into pseudorandom instances of the target correlation. Slightly more formally, a PCG for a target correlation  $C$  (which samples pairs of long correlated strings  $(y_0, y_1)$ ) is a pair  $(\text{Gen}, \text{Expand})$  of algorithms such that  $\text{Gen}(1^\lambda)$  outputs a pair of short, correlated keys  $(k_0, k_1)$  and  $\text{Expand}(\sigma, k_\sigma)$  outputs a long string  $\tilde{y}_\sigma$ . Correctness states that  $(\tilde{y}_0, \tilde{y}_1)$  are indistinguishable from a random sample from  $C$ , while security states that given  $k_{1-\sigma}$ ,  $\tilde{y}_\sigma$  looks like a random sample from  $C$  conditioned on satisfying the target correlation with  $\text{Expand}(1-\sigma, k_{1-\sigma})$ , for  $\sigma = 0, 1$ .



A PCG does not in itself provide a protocol to efficiently generate long pseudorandom correlations. To get the latter, one must combine a PCG with a *distributed key generation* protocol, which allows two parties to obviously run  $\text{Gen}(1^\lambda)$  such that each party gets one of the keys. Fortunately, for most PCGs of interest (and in particular, for all PCGs we use in this work), there exists very efficient low-communication distributed setup protocols [4, 7]. Combining a PCG with a distributed setup protocols allows to securely instantiate (with low communication) functionalities that distribute instances of the target correlation. In this work, we will directly rely in a black-box way on such functionalities, and use known protocols to instantiate them. We now expand on the two main functionalities we use in this work.

PARAMETERS:

- 2 parties, a sender and receiver, an integer  $n$ , the size of the output vector.
- A finite field  $\mathbb{F}_q$  where  $q = p^r$ ,  $p$  is a power of prime,  $r$  an integer.

FUNCTIONALITY:

- Depending on the parties:
  - If the sender is corrupted then wait for  $\mathcal{A}$  to send 2 vectors  $\mathbf{u} \in \mathbb{F}_p^n, \mathbf{v} \in \mathbb{F}_q^n$ ; samples  $\Delta \leftarrow_r \mathbb{F}_q$  and computes  $\mathbf{w} := \Delta \cdot \mathbf{u} + \mathbf{v}$ .
  - If the receiver is corrupted then wait for  $\mathcal{A}$  to send  $\mathbf{w} \in \mathbb{F}_q^n, \Delta \in \mathbb{F}_q$ ; samples  $\mathbf{u} \leftarrow_r \mathbb{F}_p^n$  and computes  $\mathbf{v} := \mathbf{w} - \Delta \cdot \mathbf{u}$ .
  - Otherwise, samples  $\mathbf{u} \in \mathbb{F}_p^n, \mathbf{v} \in \mathbb{F}_q^n, \Delta \leftarrow_r \mathbb{F}_q$  and computes  $\mathbf{w} := \Delta \cdot \mathbf{u} + \mathbf{v}$ .
- The functionality sends  $\mathbf{u} \in \mathbb{F}_p^n, \mathbf{v} \in \mathbb{F}_q^n$  to sender and  $\Delta \in \mathbb{F}_q, \mathbf{w} := \Delta \cdot \mathbf{u} + \mathbf{v}$  to receiver.

**Fig. 1.** Ideal functionality  $(n, p, q) - \mathcal{F}_{\text{svole}}$  of subfield vector-OLE

*Subfield Vector-OLE.* We described the subfield vector-OLE correlation in the technical overview of [9]. We represent on Figure 1 the ideal functionality that distributes a subfield VOLE correlation. In our concrete instantiations, we will instantiate this functionality using the efficient protocol of [4]. The latter provides a general template which can be instantiated under various flavors of the LPN assumption, and provides a conservative choice under LPN for quasi-cyclic choice. A variant of LPN that leads to a considerably more efficient protocol, when plugged in the template of [4], was recently put forth in the work [11] (we note that our communications estimate are oblivious to the underlying variant: only the computational costs depends on the LPN flavor).

*Subfield Ring-OLE.* Recently, a new PCG construction was described in [7] for the *ring-OLE* correlation. The ring-OLE correlation over a ring  $\mathcal{R}_q$  is the following correlation:  $\{((x_0, z_0), (x_1, z_1)) \mid x_0, x_1, z_0 \leftarrow_r \mathcal{R}_q, z_1 \leftarrow x_0 \cdot x_1 - z_0\}$ . In this work, we rely on a slight variant of the ring-OLE correlation, where  $x_0$  is instead sampled from a subring  $\mathcal{R}_p$  of  $\mathcal{R}_q$ . We represent the corresponding variant of the ideal functionality in the full version [9]. We note that the protocol of [7]

to instantiate the ring-OLE functionality can be adapted to handle the subfield ring-OLE functionality in a straightforward way.

### 3 Technical Overview

Our starting point is the classical KKRT protocol [18], which combines Cuckoo hashing with a batch related-key oblivious pseudorandom function (BaRK-OPRF). We assume some familiarity with the KKRT protocol in this technical overview. For completeness, we provide a high level overview of KKRT, the notion of BaRK-OPRF (batch related-key oblivious pseudorandom function), and its communication costs in Appendix of full version [9]. Our construction will also rely on a functionality that distributes *subfield vector-OLE* correlation (the sVOLE functionality): Alice gets  $(\mathbf{u}, \mathbf{v})$ , and Bob gets  $(\Delta, \mathbf{w} = \Delta \mathbf{u} + \mathbf{v})$ . Such correlation can be distributed with very low communication using pseudorandom correlation generators.

#### 3.1 A New sVOLE-Based PSI for Databases with Small Entries

Subfield-VOLE leads to a simple and natural construction of BaRK-OPRF. Let  $\ell$  be the bitlength of Alice's inputs, and let  $\mathbf{x} = (x_1, \dots, x_n)$  be the inputs of Alice, viewed as elements of  $\mathbb{F}_{2^\ell}$ . We assume for simplicity that  $\ell$  divides  $\kappa$ , the computational security parameter. Alice and Bob use an sVOLE protocol (e.g. [11]) over the field  $\mathbb{F}_{2^\kappa}$ , with subfield  $\mathbb{F}_{2^\ell}$ ; let  $(\mathbf{u}, \mathbf{v})$  be the output of Alice, and  $(\Delta, \mathbf{w})$  be the output of Bob. Recall that  $\mathbf{w} = \Delta \cdot \mathbf{u} + \mathbf{v}$ . Alice sends  $\mathbf{z} = \mathbf{x} - \mathbf{u}$  to Bob, who defines the BaRK-OPRF keys to be  $\Delta$  and  $(K_1, \dots, K_n) = \Delta \cdot \mathbf{z} + \mathbf{w}$ . The BaRK-OPRF is defined as follows:  $F_{\Delta, K_i}(y) = H(i, K_i - \Delta \cdot y)$  (all operations are over  $\mathbb{F}_{2^\kappa}$ ). Eventually, Alice outputs  $(H(i, v_i))_{i \leq n}$ . Observe that

$$\begin{aligned} H(i, v_i) &= H(i, w_i - \Delta u_i) = H(i, K_i - \Delta(z_i + u_i)) \\ &= H(i, K_i - \Delta \cdot x_i) = F_{\Delta, K_i}(x_i) \end{aligned}$$

The use of sVOLE, rather than OT extension as in the original KKRT BaRK-OPRF, has two main advantages: first, the bitwise AND is now replaced by a field multiplication. In particular, this means that we do not need anymore to use error-correcting codes, and that  $y \cdot \Delta$  retains the entire entropy of  $\Delta$ . In other words, it suffices for  $\Delta$  to be  $\kappa$ -bit long to achieve  $\kappa$  bits of security for the construction (in contrast, KKRT had to use around  $5\kappa$  bits). Second, and most importantly, the use of *subfield* VOLE allows us to completely decorrelate the size of  $\mathbf{u}$  from that of  $\Delta$ , something which can fundamentally not be achieved with the INKP OT extension. Concretely, this means that  $\mathbf{u}$  only needs to mask the input vector  $\mathbf{x}$  of Alice. If  $\mathbf{x} \in \mathbb{F}_{2^\ell}^n$ , then so do  $\mathbf{u}$  and  $\mathbf{z}$ : the communication now depends solely on the input size.

In total, our BaRK-OPRF communicates  $\ell \cdot n$  bits, plus the cost of distributing the seeds for the sVOLE generator. Using the protocol of [4] to distribute the

seeds<sup>3</sup>, the cost is logarithmic in  $n$ , hence its effect on the overall communication vanishes for large enough  $n$ .

**Combining the new OPRF with permutation-based hashing.** Plugging our new BaRK-OPRF into KKRT, and using the same parameters for Cuckoo hashing, leads to a protocol with total communication  $(1.3 \cdot \ell + 3 \cdot (\lambda + 2 \log n))n + o(n)$  bits (where the  $o(n)$  terms capture the costs of distributing the PCG seeds). Concretely, for  $n = 2^{20}$  and  $\ell = 32$  (resp. 64), this already brings the cost down, from  $1008n$  bits to  $282n$  bits (resp.  $324n$  bits). However, this can be further improved using the well-established notion of *permutation-based* hashing [22]. Concretely, in *permutation-based* hashing, an item  $x$  is written as  $x_L || x_R$ , where  $x_L$  is  $\log(1.3n)$ -bit long. The item  $x$  is inserted by mapping  $x_R$  to the bin  $x_L \oplus f(x_R)$ , where  $f$  is a  $k$ -wise independent hash function, for some large enough  $k$ . This guarantees that no collision occurs, because if two items  $x, x'$  end up mapping the same value to the same bin, this means that  $x_R = x'_R$  and  $x_L \oplus f(x_R) = x'_L \oplus f(x'_R)$ , hence  $x = x'$ . When multiple hash functions are used, as in Cuckoo hashing, the index of the hash function must be appended to  $x_R$ .

Interestingly, our use of sVOLE is crucial to enabling a permutation-hashing-based optimization: the latter only provides savings when the communication involves a  $O(\ell \cdot n)$  component (which neither KKRT nor any modern OKVS-based PSI has). In our protocol, however, it further reduces the communication to  $(1.3 \cdot (\ell - \log(1.3n) + 1) + 3 \cdot (\lambda + 2 \log n))n + o(n)$  bits, which gives  $275n$  bits for  $n = 2^{20}$  and 32-bit items, or  $317n$  bits for 64-bit items. In itself, this is a really small communication improvement. However, it has an important consequence: it implies that the Alice-to-Bob communication is now completely dominated by the Bob-to-Alice communication. Concretely, this means that we can easily afford to use a much higher number of bins (which is  $1.3n$  currently) if it can allow us to reduce the number of hash functions (which is 3). This brings us to our last optimization.

**Packing multiple items per bin with generalized Cuckoo hashing.** In this last optimization, our goal is to reduce the number of hash functions used in the Cuckoo hashing protocol, from 3 to 2, by increasing the number of bins to compensate. Unfortunately, this does not work directly with standard cuckoo hashing even while using a reasonably small stash since the cost of handling the stash is high, and nullifies all communication benefits of using two hash functions in the first place. Instead, we use a different approach: we add one degree of freedom to the Cuckoo hashing parameters, *by allowing bins to contain multiple items*. This generalization of Cuckoo hashing is not new: it has been studied in details in several works [12, 33], because it comes with a much nicer cache-friendliness than standard Cuckoo hashing.

In  $(d, k)$ -Cuckoo hashing,  $n$  items are mapped to  $(1 + \varepsilon) \cdot n$  bins using  $k$  hash functions, and each bin is allowed to contain up to  $d$  items. Allowing more items per bins significantly improves the efficiency; for example,  $(3, 2)$ -Cuckoo hashing is known to perform strictly better than standard  $(1, 3)$ -Cuckoo hashing

<sup>3</sup> This protocol uses a length- $t$  reverse VOLE protocol as a blackbox, which we instantiate with the construction of [2].

in terms of occupancy (i.e., the total number of slots  $N = d \cdot (1 + \varepsilon) \cdot n$  which must be used to guarantee a  $o(1)$  failure probability). Based on existing analysis of this variant [33], it seems reasonable to expect that (3,2)-Cuckoo hashing already achieves a strictly smaller failure probability compared to (1,3)-Cuckoo hashing, with a smaller number of bins.

We relied on extensive computer simulations on small values of  $n$  (from 256 to 2048) to select parameters, and extrapolated from these results parameters for larger values of  $n$ . More precisely, we ran  $10^7$  experiments with (3,2)-Cuckoo hashing for  $n \in \{2^8, 2^9, 2^{10}\}$  (we also experimented with  $2^{11}$ , but with a smaller number of experiments) with  $c \cdot n$  bins for various values of  $c$ . Even for a value as low as  $c = 0.65$  and values of  $n$  as low as  $2^9$ , our experiments never reported any insertion failure, indicating that the empirical failure probability should already be way below  $2^{-20}$ . Since the theoretical failure probability is known to scale as  $O(1/n^\delta)$  for some constant  $\delta$  with reasonably small constant factors, we extrapolate that for large enough values of  $n$ , e.g.  $n \geq 2^{18}$ , the failure probability should be well below  $2^{-40}$ .

**Alternative hashing variants.** Alternatively, when allowing multiple items per bins, we can consider other hashing variants. Two natural choices are two-choice hashing [20], where each bin can have up to two  $d$  items and each item is placed in the least-full of two bins, and simple hashing, where a single hash function is used to map the items to bins (standard results show that, when hashing  $n$  items to  $O(n)$  bins this way, the maximum load will be of the order of  $\log n / \log \log n$  with high probability). As we will see, these choices of hashing lead to various communication versus computation tradeoffs in our protocols, and the optimal choice also depends on the database size.

**A membership BaRK-OPRF.** There remains a non-trivial task: to use some of the above hashing variants, we need a protocol to handle hashing with up to  $d$  items per bins. Intuitively, denoting  $\mathbf{x}_i = (x_i^{(1)}, \dots, x_i^{(d)})$  the  $d$  entries of the bin  $i$ , we want to construct a new kind of *membership* OPRF (similar in spirit to the notion of multi-point OPRF in the literature), where Bob obtains  $F_{\Delta, K_i}(y)$  and Alice obtains the set  $F_{\Delta, K_i}(\mathbf{x}_i) = \{F_{\Delta, K_i}(x_i^{(j)})\}_{j \leq d}$ . This implies that  $F_{\Delta, K_i}(y) \in F_{\Delta, K_i}(\mathbf{x}_i)$  if and only if  $y$  is equal to any entry of  $\mathbf{x}_i$ , and  $F_{\Delta, K_i}(y)$  looks pseudorandom to Alice otherwise.

Going back to the BaRK-OPRF, recall that for a bin  $i$  where Alice placed  $x_i$  and Bob placed  $y_i$ , Alice computes  $H(i, v_i)$  and Bob computes  $H(i, K_i - \Delta y_i) = H(i, \Delta \cdot (x_i - y_i) + v_i)$ . Here, we view the  $x_i - y_i$  term as  $P_{x_i}(y_i)$ , where  $P_{x_i} = X - x_i$  is a degree-1 polynomial with root  $x_i$ . This view suggests a natural generalization of this approach, where the  $P_{x_i}$  polynomials are replaced by higher degree polynomials. Define  $P_{\mathbf{x}_i}$  to be the polynomial  $\prod_{j=1}^d (X - x_i^{(j)})$ , and let  $(c_{j,i})_{0 \leq j \leq d-1}$  denote its coefficients:  $P_{\mathbf{x}_i}(X) = X^d + \sum_{j=0}^{d-1} c_{j,i} \cdot X^j$ . Our new *membership* BaRK-OPRF is a direct generalization of the BaRK-OPRF from Section 3.1, which we sketch below.

**Our construction.** Let  $m$  be the bitlength of Alice's inputs inside the bins, and let  $(\mathbf{x}_1, \dots, \mathbf{x}_N)$  be the inputs of Alice in each of the  $N$  bins, where the

inputs in each bin are viewed as length- $d$  vectors of elements of  $\mathbb{F}_{2^m}$ . We assume for simplicity that  $m$  divides  $\kappa$ , the computational security parameter. Alice and Bob use  $d$  sVOLE protocol (e.g. [11]) over the field  $\mathbb{F}_{2^\kappa}$ , with subfield  $\mathbb{F}_{2^m}$ , *with the same value  $\Delta$* .<sup>4</sup> Let  $(\mathbf{u}_j, \mathbf{v}_j)_{j \leq d}$  be the outputs of Alice, and  $(\Delta, (\mathbf{w}_j)_{j \leq d})$  be the output of Bob. Recall that  $\mathbf{w}_j = \Delta \cdot \mathbf{u}_j + \mathbf{v}_j$ .

For each  $\mathbf{x}_i$ , let  $(c_{0,i}, \dots, c_{d-1,i})$  be the coefficients of the polynomial  $P_{\mathbf{x}_i}$  (omitting the coefficient of  $X^d$ , which is always 1). Let  $\mathbf{c}_j$  denote the vector  $(c_{j,i})_{i \leq N}$  for  $j = 0$  to  $d-1$ . Alice sends  $\mathbf{z}_j = \mathbf{c}_j - \mathbf{u}_j$  for  $j = 0$  to  $d-1$  to Bob, who defines the membership BaRK-OPRF keys to be  $\Delta$  and  $K_i = (k_{j,i})_{0 \leq j \leq d-1} = (\Delta \cdot z_{j,i} + w_{j,i})_{0 \leq j \leq d-1}$  for  $i = 1$  to  $N$ . Define the following degree- $d$  polynomial  $P_{\Delta, K_i}$  over  $\mathbb{F}_q$ :  $P_{\Delta, K_i}(X) = \Delta \cdot X^d + \sum_{j=0}^{d-1} k_{j,i} \cdot X^j$ . The OPRF is defined as follows:  $F_{\Delta, K_i}(y) = H(i, P_{\Delta, K_i}(y))$  (all operations are over  $\mathbb{F}_{2^\kappa}$ ). Eventually, for each bin  $i$ , Alice sets her  $d$  tuple of outputs to be  $F_{\Delta, K_i}(\mathbf{x}_i) = \{H(i, \sum_{j=0}^{d-1} v_{j,i} \cdot (x_i^{(k)})^j)\}_{k \leq d}$ . Observe that, since  $k_{j,i} = \Delta z_{j,i} + w_{j,i} = \Delta c_{j,i} + v_{j,i}$  for all  $i, j$ , we have  $H(i, P_{\Delta, K_i}(y)) = H\left(i, \Delta \cdot \left(y^d + \sum_{j=0}^{d-1} c_{j,i} y^j\right) + \sum_{j=0}^{d-1} v_{j,i} y^j\right)$ , which is equal to  $H\left(i, \Delta \cdot P_{\mathbf{x}_i}(y) + \sum_{j=0}^{d-1} v_{j,i} y^j\right)$ . Therefore, if there exists  $k \in \{1, \dots, d\}$  such that  $y = x_i^{(k)}$ , we have  $P_{\mathbf{x}_i}(y) = 0$ , and  $H(i, P_{\Delta, K_i}(y)) = H(i, \sum_{j=0}^{d-1} v_{j,i} \cdot (x_i^{(k)})^j) \in F_{\Delta, K_i}(\mathbf{x}_i)$ . On the other hand, whenever  $P_{\mathbf{x}_i}(y) \neq 0$ , then the  $\Delta \cdot P_{\mathbf{x}_i}(y)$  term in the hash makes the output pseudorandom from the viewpoint of Alice, under the correlation robustness of the hash function.

**Tying up loose ends.** Using the new construction from the previous Section, together with (3, 2)-Cuckoo hashing, leads to a total communication of  $(0.65 \cdot 3(\ell - \log(0.65n) + 1) + 2 \cdot (\lambda + 2 \log n))n + o(n)$  bits, where the  $o(n)$  corresponds to the cost of setting up the PCG seeds. For  $n = 2^{20}$  and 32 bits items, this gives 148n bits of communication. We mention a few remaining details. First, in the construction of membership BaRK-OPRF, Alice and Bob need to invoke  $d = 3$  length- $N$  sVOLE. In fact, it suffices to invoke a single length- $3N$  sVOLE, and to cut the output in three equal length parts, to obtain the necessary correlation. This means that the concrete cost of distributing the sVOLE seeds remains that of generating a single sVOLE (e.g.  $\approx 0.7n$  bits for  $n = 2^{20}$ ).

Second, in the above, we overlooked an important subtlety: a bin can possibly contain less than  $d$  items. In KKRT, this was handled by adding dummy items to empty bins. We use instead a more efficient approach with a negligible extra cost called a *variant* of our OPRF (details in section 4).

### 3.2 Malicious Security

We then turn our attention to maliciously secure PSI. Here, it is well known that Cuckoo hashing and two-choice hashing are not usable. Consequently, we focus on simple hashing as our choice of the underlying hash technique. Using maliciously secure subfield-VOLE, which can be implemented very efficiently [4,

<sup>4</sup> Note that all known sVOLE protocols allow Bob to choose the value of  $\Delta$ , hence Bob can enforce the use of the same  $\Delta$  across all instances.

11], we enhance our membership BaRK-OPRF to the malicious setting, with a minimal overhead. Then, we apply two standard methods to achieve security against malicious adversaries in our PSI protocol:

*First method: direct approach.* The first method increases the PRF output length to  $\kappa$ . Using the analysis of [27], this suffices to allow for extracting the input of a malicious sender. However, this makes the communication depend linearly on  $\kappa$ , which severely harms communication complexity.

*Second method: dual execution.* To recover a  $\kappa$ -independent communication complexity, we then turn our attention to the dual execution technique [26]. Here, the idea is simple: the parties will invoke the malicious BaRK-OPRF twice, exchanging their roles. Then, the sender sends, for each entry  $x$  of his database, a value of the form  $\text{PRF}_A(x) \oplus \text{PRF}_B(x)$ , where  $\text{PRF}_A(x)$  is obtained by the sender when invoking the BaRK-OPRF functionality as sender, and  $\text{PRF}_B(x)$  is the PRF output obtained when invoking the functionality as receiver. Here, it becomes possible to extract the input set of each party simply from its call as receiver to the BaRK-OPRF functionality, which does not require to increase the output length of the OPRF. The price to pay is that the protocol now uses two calls to the BaRK-OPRF. Concretely, the total communication becomes  $(2 \cdot N \cdot d(\ell - \log(N)) + (\lambda + \log n)n + o(n))$ , where  $N$  is the number of bins,  $d$  the maximum load of a bin, and  $\ell$  the input size (e.g. for  $n = 2^{20}$ , one can choose  $N = n/10$  and  $d = 47$ , see [26, Figure 5]). For small database entries, this outperforms all known malicious PSI protocols.

### 3.3 An Efficient PSI in the Standard Model

In our last construction, we use a different functionality: we rely on the subfield ring-OLE functionality (given on Appendix of full version [9]), that generates a subfield ring-OLE correlation over the rings  $\mathcal{R}_p = \mathbb{F}_p[X]/F(X)$ ,  $\mathcal{R}_q = \mathbb{F}_q[X]/F(X)$ , and  $F(X)$  is some polynomial of degree  $2n + 1$  (more generally, when the two parties have sets of different size  $n$  and  $m$ ,  $F$  will be of degree  $n + m + 1$ ). At a high level, the functionality  $\mathcal{F}_{\text{sole}}$  distributes to Alice  $(a, s_A) \in \mathcal{R}_p \times \mathcal{R}_q$  and  $(b, s_B) \in (\mathcal{R}_q)^2$  to Bob such that  $ab = s_A + s_B$ . Our protocol makes a single black-box call to this functionality. Consider two parties, a sender Alice and a receiver Bob, where Alice has a set  $A = \{x_1, x_2, \dots, x_n\} \in \mathbb{F}_p^n$  and Bob has a set  $B = \{y_1, y_2, \dots, y_n\} \in \mathbb{F}_p^n$ . Define  $p_A := \prod_{i=1}^n (X - x_i) \in \mathcal{R}_p$  and  $p_B := \prod_{i=1}^n (X - y_i) \in \mathcal{R}_p$ . Let  $I := A \cap B$  denote the target output. The protocol computes the common roots of  $p_A$  and  $p_B$ , i.e.,  $\text{gcd}(p_A, p_B)$ .

By revealing appropriate linear combination of their shares and their input polynomials, Alice and Bob will “derandomize” this correlation, allowing Alice to learn the polynomial  $u = p_A b_0 + p_B b'_0$ , where  $b_0, b'_0$  are two uniformly random degree- $n$  polynomials known by Bob (this also requires revealing the high-order coefficients of  $b$ , to reduce the degree- $2n$  random polynomial  $b$  to a degree- $n$  random polynomial  $b_0$ ). Using some standard lemmas about polynomials, the polynomial  $u$  can be factored as  $\text{gcd}(p_A, p_B) \cdot p_R$ , where with high probability,  $p_R$  has no common root with  $p_A$ . This allows Alice to compute the intersection  $I = A \cap B$  as  $I = \{x_i \in A : u(x_i) = 0\}$ . Concretely:

- Alice computes and sends  $t_A = a - p_A$  to Bob.
- Bob sets  $s'_B \leftarrow s_B - t_A b$ . Then, Bob decomposes  $b$  as  $b = b_0 + b_1 \cdot X^n$  (where  $b_0, b_1$  are degree- $n$  polynomials), sets  $s'_B \leftarrow s_B - t_A b$ , and picks a random degree- $n$  polynomial  $b'_0$  over  $\mathcal{R}_q$ . He sends  $b_1$  and  $t_B \leftarrow s'_B + p_B b'_0$  to Alice.
- **(Output)** Alice sets  $u \leftarrow t_B - p_A b_1 \cdot X^n + s_A$ ; note that  $u = p_A b_0 + p_B b'_0$ . Alice outputs the set  $I = \{x \in A \mid u(x) = 0\}$ .

We prove that this construction achieves “augmented semi-honest security”, a strengthening of honest-but-curious corruption where the adversary is allowed to change the corrupted parties’ inputs. Furthermore, we securely realize the functionality  $\mathcal{F}_{\text{sole}}$  using the PCG-based protocol of [7], which is secure under the ring-LPN assumption. Instantiating the subfield ring OLE this way allows to import a powerful feature of the PCG of [7], which is its *programmability*: when generating a ring-OLE correlation, the receiver can ensure that her output  $a$  remains *identical* across multiple instances of the protocol with different parties. Using this programmability feature, we show that our protocol can be *batched*: a single  $O(\ell \cdot n)$ -size client message encoding her database  $A$  can be reused with  $N$  different servers with databases  $B_i$ , allowing her to learn  $A \cap B_i$  using a single message from each server afterwards.

**Achieving malicious security.** We then turn our attention to security against malicious adversaries. Our upgrade introduces only a minimal communication overhead to the protocol, independent of the set sizes  $n$ . At a high level, the main issues that can occur in the malicious setting is when Alice sets  $p_A = 0$ , or when Bob sets  $p_B b'_0 = 0$ . Indeed, since Alice gets  $u = p_A b_0 + p_B b'_0$ , if  $p_A = 0$ , she can learn Bob’s entire input set  $p_B$ . On the other hand, if  $p_B b'_0 = 0$ , Bob forces the output to be  $A$ .

We handle both issues separately. The second issue is intuitively simpler to handle, since when Bob carries out this attack, Alice will notice that her output is exactly her set  $A$ . This suggests a simple way around: if Alice notice at the end of the protocol that the output is equal to  $A$ , she aborts the protocol. Of course, a honest Bob could have an input  $B$  with  $A \subseteq B$ , in which case this modification would harm correctness. But there is a simple way around: prior to the protocol, Alice and Bob can just agree on a reserved dummy item  $d$  (we will pick  $d = 1$  in the protocol, but this choice is arbitrary), which is guaranteed to be in neither databases. If database entries are elements of a field  $\mathbb{F}_{p'}$ , this can simply be done by choosing any slightly larger field  $\mathbb{F}_p$  of size  $|\mathbb{F}_p| \geq |\mathbb{F}_{p'}| + 1$ , reserving one element of  $\mathbb{F}_p$  to encode  $d$ , and mapping the elements of  $\mathbb{F}_{p'}$  to the remaining elements. Then, Alice and Bob execute the protocol on inputs  $A \cup \{1\}$  and  $B$ , which guarantees that  $B$  does not contain  $A$ .

For the first issue, Bob must check before sending  $t_B = s'_B + p_B b'_0$  that Alice did not set  $p_A$  to be 0 when computing  $t_A = a - p_A$ . Intuitively, this will be done by letting Bob check that  $p_A(x) \neq 0$ , for an appropriate input  $x$ . This, however, must be done with some care, since learning  $p_A(x)$  could leak information to a corrupted Bob. We handle this issue by reserving a second element of  $\mathbb{F}_p$  (hence we now need  $|\mathbb{F}_p| \geq |\mathbb{F}_{p'}| + 2$ ), which we assume w.l.o.g. to be 0, which should again be in neither set. Then, Alice will define the encoding of her set to be

the degree- $n$  polynomial  $p_A$  such that  $p_A(\text{map}(a)) = 0$  for every  $a \in A$ , and  $p_A(0) = 1$ . Then, we let Bob first send  $b_1$ , without sending  $t_B$ . Afterwards, Bob computes  $s'_B \leftarrow s_B - t_A b$  and Alice computes  $s'_A \leftarrow s_A - p_A b_1 \cdot X^n$ . Observe that if both parties behave honestly,  $s'_a + s'_b = ab - t_A b - p_A b_1 \cdot X^n = ab - ab + p_A b - p_A b_1 \cdot X^n = p_A b_0$ . To enforce  $p_A \neq 0$ , we will check that the above equation holds for some nonzero  $p_A$ . Crucially, since both  $p_A$  and  $b_0$  have degree at most  $n$ , no reduction modulo  $F(X)$  occurs in the right hand side of the equation. This implies that we can simply check that the equation holds for the reserved input  $x = 0$  (since a honest  $p_A$  is guaranteed to satisfy  $p_A(0) = 1 \neq 0$ ). To check this, we let Alice send  $s'_A(0)$  to Bob, who checks that  $s'_A(0) = b_0(0) - s'_B(0)$ ; if the check fails, Bob aborts the protocol.

## 4 PSI from Subfield-VOLE

### 4.1 A new membership batched OPRF

Our BaRK-OPRF allows the sender to hold a set of keys  $(\mathbf{k}_i)_{i \leq N}$  such that each key is assigned with a tuple of  $d$  input elements of the receiver and then the receiver learns a PRF output on each element in this tuple corresponding with the same key. More formally, denoting  $\mathbf{x}_i = (x_i^{(1)}, \dots, x_i^{(d)})$  consisting of  $d$  entries, the sender gets  $F(i, y)$  and the receiver obtains a set  $\{F(i, x_i^{(j)})\}_{j \leq d}$  such that  $F(i, y) \in \{F(i, x_i^{(j)})\}_{j \leq d}$  if and only if  $y$  is equal to any entry of  $\mathbf{x}_i$ , and  $F(i, y)$  looks pseudorandom to the receiver otherwise.

#### PARAMETERS:

$\mathbb{F}_p$  is a finite field. There are 2 parties, a sender and a receiver with input set  $X = \{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_N\} \subseteq \mathbb{F}_p$  where  $\mathbf{x}_i = (x_i^{(1)}, \dots, x_i^{(d)})$ .

#### FUNCTIONALITY:

- Wait for input (sender, id) from the sender and (receiver, id, X) from the receiver. The functionality samples a PRF  $F$  then  $\forall x \in \mathbf{x}_i$  outputs  $F(i, x)$  to the receiver for  $i \in [1, N]$ .
- When the sender inputs any  $(i, y) \in [1, N] \times \mathbb{F}_p$ , functionality gives  $F(i, y)$  to the sender.

**Fig. 2.** Ideal functionality  $\mathcal{F}_{\text{oprf}}$

**Main construction.** Assume that the receiver inputs the set of  $n = Nd$  elements:  $X = \{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_N\} \subseteq \mathbb{F}_p$  where  $\mathbf{x}_i = (x_i^{(1)}, \dots, x_i^{(d)})$ . First, the sender and the receiver invoke the  $\mathcal{F}_{\text{s vole}}$  protocol of dimension  $n$ , with their roles reversed, to get a random sVOLE correlation. Specifically, the receiver learns a pair of vectors  $(\mathbf{u}, \mathbf{v})$  where  $\mathbf{u} \in \mathbb{F}_p^n$ ,  $\mathbf{v} \in \mathbb{F}_q^n$ , the sender gets  $\Delta \in \mathbb{F}_q$  and  $\mathbf{w} := \Delta \cdot \mathbf{u} + \mathbf{v}$ . Denoting  $\mathbf{u} = (\mathbf{u}_1, \mathbf{u}_2, \dots, \mathbf{u}_N)$  where  $(u_{j,i})_{1 \leq j \leq d}$  are  $d$  entries of vector  $\mathbf{u}_i$ . This notation is the same for  $\mathbf{v}, \mathbf{w}$ . Consider  $\mathbf{x}_i$  and its associated polynomial as  $P_{\mathbf{x}_i}(X) = \prod_{j=1}^d (X - x_i^{(j)}) = X^d + \sum_{j=1}^d c_{j,i} \cdot X^{j-1}$  where  $c_{j,i} \in \mathbb{F}_p$  for  $i \in [1, N]$ ,  $j \in [1, d]$ .



Now, the receiver defines  $\mathbf{c}_i := (c_{j,i})_{j \leq d}$ ,  $\mathbf{c} := (\mathbf{c}_1, \mathbf{c}_2, \dots, \mathbf{c}_N)$ , and then  $\forall i \in [1, N]$  sends to the sender  $\mathbf{z}_i := \mathbf{c}_i - \mathbf{u}_i \in \mathbb{F}_p^d$ . Above, the  $\mathbf{u}_i$  are masks for the coefficients  $\mathbf{c}_i$  of (the polynomial associated)  $\mathbf{x}_i$ . Indeed,  $\mathbf{u}_i$  are distributed uniformly at random in the subfield  $\mathbb{F}_p$ , then the vector  $\mathbf{z}_i$  is a uniformly random over  $\mathbb{F}_p^n$  from the viewpoint of the sender. The two parties will run a coin flipping protocol to get a random value  $t \leftarrow \mathbb{F}_q$ . For  $i \in [1, N]$ , the receiver defines the PRF output on each input  $x \in \mathbf{x}_i$  as  $F(i, x) = H\left(i|t|x, \sum_{j=1}^d v_{j,i} \cdot x^{j-1}\right)$ .

On the other hand, after receiving the vectors  $\mathbf{z}_i$ , for  $i \in [1, N]$ , the sender defines the vector  $\mathbf{k}_i := \mathbf{w}_i + \Delta \cdot \mathbf{z}_i$ . As a consequence, for any input  $(i, y) \in [1, N] \times \mathbb{F}_p$ , its PRF output is computed as:  $F(i, y) = H\left(i|t|y, \Delta \cdot y^d + \sum_{j=1}^d k_{j,i} \cdot y^{j-1}\right)$ .

**Correctness and Security.** To see why PRF output is defined as above. Observe that  $\mathbf{k}_i := \mathbf{w}_i + \Delta \cdot \mathbf{z}_i = \mathbf{v}_i + \Delta \cdot \mathbf{c}_i$ . Then, we have

$$\begin{aligned} \Delta \cdot y^d + \sum_{j=1}^d k_{j,i} \cdot y^{j-1} &= \Delta \cdot y^d + \sum_{j=1}^d (v_{j,i} + \Delta \cdot c_{j,i}) \cdot y^{j-1} \\ &= \Delta \cdot (y^d + \sum_{j=1}^d c_{j,i} \cdot y^{j-1}) + \sum_{j=1}^d v_{j,i} \cdot y^{j-1} = \Delta \cdot P_{\mathbf{x}_i}(y) + \sum_{j=1}^d v_{j,i} \cdot y^{j-1} \end{aligned}$$

so if  $y \in \mathbf{x}_i$  then  $P_{\mathbf{x}_i}(y) = 0$  which leads to  $F(i, y) \in \{F(i, x_i^{(j)})\}_{j \leq d}$ .

**Theorem 1.** *The protocol  $\Pi_{\text{opr}}f$  (Figure 3) instantiated with random oracles  $H, H'$ , securely realizes the ideal functionality of  $\mathcal{F}_{\text{opr}}f$  (Figure 2) against a malicious setting in the  $\mathcal{F}_{\text{svole}}$  hybrid model.*

Note that the output  $v$  of  $H$  is chosen depending on the concrete structure of PSI and the target setting (semi-honest or malicious). This parameter is detailed in the section 4.2 for a semi-honest setting and the section 4.3 for a malicious setting.

## 4.2 A new semi-honest PSI from mOPRF

**A variant of BaRK-OPRF.** We now propose a variant of our BaRK-OPRF to deal with the case when the size of each tuple input is not necessarily equal to  $d$ . This means that the receiver now can divide the input set to  $N$  tuples  $\mathbf{x}_i$  and each tuple has less than or equal to  $d$  items. Meanwhile, the sender is not allowed to learn about how many exactly items are in each tuple. This functionality can be obtained from our BaRK-OPRF plus a small extra cost, i.e, a *subfield* VOLE of length  $N$  over the subfield  $\mathbb{F}_2$ .

The idea is as follows. The receiver's input set  $X = \{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_N\} \subseteq \mathbb{F}_p$  where  $\mathbf{x}_i = (x_i^{(1)}, \dots, x_i^{(j_i)})$ ,  $j_i \leq d$ . The polynomial associated to  $\{\mathbf{x}_i\}_{i \leq N}$  will be expressed as a polynomial of degree  $d$ :  $P_{\mathbf{x}_i}(X) = \prod_{j=1}^{j_i} (X - x_i^{(j)}) = \sum_{j=1}^{d+1} c_{j,i} \cdot X^{j-1}$  where  $c_{j,i} \in \mathbb{F}_p$ .

As a result, the set of the coefficients of  $P_{\mathbf{x}_i}(X) = (c_{1,i}, c_{2,i}, \dots, c_{d+1,i})$ . We remark that, compared to the associated polynomial in our original BaRK-OPRF which has a constant coefficient of degree  $d$  of 1, in our variant version

## PARAMETERS:

- Given  $\mathbb{F}_p \subseteq \mathbb{F}_q$  where  $\mathbb{F}_q \approx O(2^\kappa)$ ,  $H : \{0, 1\}^* \times \mathbb{F}_q \rightarrow \{0, 1\}^v$  and  $H' : \mathbb{F}_q \rightarrow \mathbb{F}_q$  are random oracles.
- The sender has no input and the receiver inputs a set  $X = \{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_N\} \subseteq \mathbb{F}_p$  where  $\mathbf{x}_i = (x_i^{(1)}, \dots, x_i^{(d)})$  and  $n = Nd$ .

## PROTOCOL:

1. The sender and the receiver invoke to the  $\mathcal{F}_{\text{vole}}$  of dimension  $n$  in the  $\mathbb{F}_q$  over the  $\mathbb{F}_p$  with the inverse role. The receiver gets two random vectors  $\mathbf{u} \in \mathbb{F}_p^n, \mathbf{v} \in \mathbb{F}_q^n$  and the sender receives  $\Delta \in \mathbb{F}_q, \mathbf{w} := \Delta \mathbf{u} + \mathbf{v} \in \mathbb{F}_q^n$ . Denoting  $\mathbf{u} = (\mathbf{u}_1, \mathbf{u}_2, \dots, \mathbf{u}_N)$  where  $\mathbf{u}_i = (c_{j,i})_{1 \leq j \leq d}$ . This denotation is the same for  $\mathbf{v}, \mathbf{w}$ .
2. The receiver samples  $t_r \leftarrow \mathbb{F}_q$  and sends  $h_r := H'(t_r)$  to the sender.
3. The sender samples  $t_s \leftarrow \mathbb{F}_q$  and sends  $h_s := H'(t_s)$  to the receiver.
4. The receiver determines the associated polynomial for each  $\mathbf{x}_i$  as

$$P_{\mathbf{x}_i}(X) = \prod_{j=1}^d (X - x_i^{(j)}) = X^d + \sum_{j=1}^d c_{j,i} \cdot X^{j-1}$$

where  $c_{j,i} \in \mathbb{F}_p$  for  $i \in [1, N], j \in [1, d]$ .

5. Denoting  $\mathbf{c}_i := (c_{j,i})_{1 \leq j \leq d}; \mathbf{c} := (\mathbf{c}_1, \mathbf{c}_2, \dots, \mathbf{c}_N)$ , the receiver computes  $\mathbf{z}_i := \mathbf{c}_i - \mathbf{u}_i \in \mathbb{F}_p^d$ , and then sends  $\mathbf{z}_i$  and  $t_r$  to the sender.
6. The sender aborts if  $H'(t_r) \neq h_r$ .
7. The sender sends  $t_s$  to the receiver, the receiver aborts if  $H'(t_s) \neq h_s$  and both parties define  $t = t_s \oplus t_r$ .
8. The receiver outputs the PRF values on the input  $x \in \mathbf{x}_i$  for  $i \in [1, N]$  as

$$F(i, x) = H\left(i|t|x, \sum_{j=1}^d v_{j,i} \cdot x^{j-1}\right)$$

9. For  $i \in [1, N]$ , the sender defines  $\mathbf{k}_i = \mathbf{w}_i + \Delta \mathbf{z}_i$ . For any input  $(i, y) \in [1, N] \times \mathbb{F}_p$ , the sender computes the PRF output by below formula

$$F(i, y) = H\left(i|t|y, \Delta \cdot y^d + \sum_{j=1}^d k_{j,i} \cdot y^{j-1}\right)$$

**Fig. 3.** Our batch BaRK-OPRF  $\Pi_{\text{oprf}}$  based on subVOLE

this coefficient will equal 0 or 1 since the degree of  $P_{\mathbf{x}_i}(X)$  is *less* than or equal to  $d$ . So, it requires  $(d+1)$  masks for this polynomial instead of  $d$ , but the mask for the coefficient of degree  $d$  only needs to be in  $\mathbb{F}_2$ . For each tuple, we require an additional value  $u_i \in \mathbb{F}_2$ , so in total we need an additional subfield VOLE of length  $N$  over the subfield  $\mathbb{F}_2$ .

More formally, the sender and receiver invoke a subfield VOLE of length  $n$  over the subfield  $\mathbb{F}_p$  as before (all the notations in figure 3 are reused), and additionally invoke another subfield VOLE instance over the subfield  $\mathbb{F}_2$  of length  $N$  with an inverse role, while the receiver gets  $\mathbf{u}' \in \mathbb{F}_2^N$ , and  $\mathbf{v}' \in \mathbb{F}_q^N$  the sender holds  $\Delta \in \mathbb{F}_q$  ( $\Delta$  is the same for each time invoking *subfield* VOLE) and

$\mathbf{w}' := \Delta \cdot \mathbf{u}' + \mathbf{v}'$ . The receiver sends to the sender vectors  $\mathbf{z}_i$  as before, and an extra vector  $\mathbf{z}'$  defined as  $z'_i := c_{d+1,i} - u'_i$  for  $i \in [1, N]$ . The receiver outputs on input  $x \in \mathbf{x}_i$  are computed as  $F(i, x) = H(i|t|x, v'_i \cdot x^d + \sum_{j=1}^d v_{j,i} \cdot x^{j-1})$ . On the other hand, the sender defines their PRF values on input  $(i, y)$  where  $i \in [1, N]$ ,  $y \in \mathbb{F}_p$  as  $F(i, y) = H(i|t|y, (w'_i + \Delta z'_i) \cdot y^d + \sum_{j=1}^d k_{j,i} \cdot y^{j-1})$ .

**Main construction of a new PSI.** The sender and the receiver have two input sets  $X = \{x_1, x_2, \dots, x_n\}$  and  $Y = \{y_1, y_2, \dots, y_n\}$ . Assume that all of these elements have the bit-length  $\ell$ . Intuitively, our BaRK-OPRF is constructed from subVOLE to handle the case when having multiple items per bin. Then this specialized BaRK-OPRF can combine with some hashing techniques to form an efficient PSI protocol. In the next part 4.2, we discuss these types of hashing. Our PSI protocol is described in Figure 4; it builds upon the protocol of [18] using GCH and BaRK-OPRF. For simplicity, we describe our protocol directly with generalized Cuckoo hashing; adapting the protocol to other variants is immediate. We elaborate on our protocol below. In our protocol, the receiver first uses  $(d, k)$ -Cuckoo hashing to map his input set  $Y$  to a table with  $N$  bins, note that the bit-length of the values stored in a bin is  $\ell - \log N$  instead of  $\ell$ . Depending on the size of  $n$ , we use one of two approaches to handle the bins which are not full (the threshold was chosen empirically to optimize communication).

- If  $n \geq 2^{20}$ , the variant of our BaRK-OPRF (using an additional subfield VOLE over  $\mathbb{F}_2$ ) is used; for such sizes, the concrete cost of implementing the additional sVOLE vanishes.
- Otherwise, when  $n < 2^{20}$ , the receiver adds dummy items to bins such that each bin contains *exactly*  $d$  items. To avoid collisions between the dummy items and the elements in the same bin of the sender, we pad an extra bit to all items in the following way:  $i|x|b$  where  $i$  is the index of hash function corresponding with the stored value  $x$  while  $b = 1$  if  $x$  is a dummy item added and  $b = 0$  otherwise.

In both case, the sender computes  $k \cdot n$  PRF evaluations and sends (shuffled) to the receiver, who compares them with his OPRF outputs, and outputs the intersection set. To reduce the computational cost in this step, the sender can send separately each set  $H_i$  ( $i \in [1, k]$ ) which contains the PRF outputs of each  $x \in X$  with the related bin  $h_i(x)$ . Then for each element, the receiver only needs to search for one set (among  $k$  sets  $H_i$ ) of  $n$  items instead of  $k \cdot n$ .

**Alternative hashing methods.** There are two hashing schemes that can be fit into our PSI structure.

*2-choice hashing* [20] is a variant of Cuckoo hashing where one item  $x$  is assigned to one of two bins  $h_1(x)$  or  $h_2(x)$ . However, there is no restriction on the number of items per bin and an item is put in a bin which already has fewer items. [20] proposes both theoretical references and heuristic parameters for 2-choice hashing, which require only a small number of dummy items. Let us assume we have  $n$  items and 2 hash functions; using 2-choice hashing allows to map  $n$  items to  $N$  bins in time  $O(n \log n)$  where each bin contains at most  $L = \lceil n/N \rceil + 1$  items with a probability  $1 - O(1/N)^{L-1}$ .

*Simple hashing* uses one hash function  $h$  to map an item  $x$  to bin  $h(x)$ . For

## PARAMETERS:

- The sender and the receiver have respectively input sets  $X = \{x_1, x_2, \dots, x_n\}$  and  $Y = \{y_1, y_2, \dots, y_n\}$ , all elements of bit-length  $\ell$ .
- A  $(d, k)$ -generalized Cuckoo hashing (GCH) scheme mapping  $n$  items to  $N$  bins by  $k$  hash functions  $h_1, h_2, \dots, h_k : \{0, 1\}^* \rightarrow [N]$  where  $Nd > n$  and  $d = O(1)$  (see Section 4.2).

## PROTOCOL:

1. The receiver uses  $(d, k)$ -Cuckoo hashing with  $k$  hash functions to map the elements in  $Y$  to the table  $\mathcal{B}$  consisting of  $N$  bins, where each bin  $i$  has  $j_i \leq d$  items. Denote  $y_{j,i}$  is an element in  $Y$  assigned to position  $j$  of bin  $i$  and its stored value in table  $\mathcal{B}$  is  $y'_{j,i}$ .
2. Depending on the size of  $n$ , there are two alternatives:
  - (a)  $n \geq 2^{20}$ , the sender and receiver invoke our variant of  $\Pi_{\text{opr}}^{\text{f}}$  where the receiver uses the input set  $Y_{\mathcal{B}} = \{\mathbf{y}_1, \mathbf{y}_2, \dots, \mathbf{y}_N\}$  defined as follows:
    - $\mathbf{y}_i = \{r_{1,i}, r_{2,i}, \dots, r_{j_i,i}\}$ .
    - $r_{j,i} = t \parallel y'_{j,i}$  where  $t$  is index of a hash function such that  $h_t(y_{j,i}) = i$ .
  - (b)  $n < 2^{20}$ , the sender and receiver directly invoke the  $\Pi_{\text{opr}}^{\text{f}}$  where the receiver uses the input set  $Y_{\mathcal{B}} = \{\mathbf{y}_1, \mathbf{y}_2, \dots, \mathbf{y}_N\}$  defined as follows:
    - $\mathbf{y}_i = \{r_{1,i}, r_{2,i}, \dots, r_{d,i}\}$ .
      - For  $j \leq j_i$ :  $r_{j,i} = t \parallel y'_{j,i} \parallel 1$  where  $t$  is index of hash function such that  $h_t(y_{j,i}) = i$ .
      - Otherwise,  $r_{j,i} = t \parallel \text{dummy value} \parallel 0$  where  $t \leftarrow_r [1, k]$ .
3. The receiver obtains  $n$  instances OPRF:

$$Y' = \{\text{PRF}(i, r_{i,j}) \mid i \in [1, N], j \leq j_i\}$$

4. The sender uses the  $k$  hash functions to map the  $n$  element in  $X$  to the  $N$  bins. Let  $x_t$  denote the value stored at bin  $h_t(x)$  when mapping  $x$  for  $t \in [1, k]$ .
5. The sender computes the sets of  $k \cdot n$  PRF outputs:
  - (a) For  $n \geq 2^{20}$ :  $H_t = \{\text{PRF}(h_t(x), t \parallel x_t) \mid x \in X\}$  for  $t \in [1, k]$ .
  - (b) For  $n < 2^{20}$ :  $H_t = \{\text{PRF}(h_t(x), t \parallel x_t \parallel 1) \mid x \in X\}$  for  $t \in [1, k]$ .
 Then the sender randomly permutes and sends each set to the receiver.
6. The receiver finds the intersection:
  - if  $y \in Y$  is mapped to the position  $j$  of bin  $i$  by function  $h_t$  then check whether  $\text{PRF}(i, r_{i,j}) \in H_t$  ( $r_{i,j}$  is defined depending on  $n$ ).
  - Outputs the intersection set.

**Fig. 4.** Our new semi-honest PSI protocol from BaRK-OPRF

security, the number of items per bin can leak some information then it requires padding each bin with dummy items until having an equal number of items per bin. With very high probability, for  $N = O(n \log n)$  bins, the maximum possible items per bin is  $O(\log n)$ . The percentage of the occupation of dummy items is higher than others. However, simple hashing avoids ambiguities about where an item can be placed, a property which is crucial in the malicious setting.

**Parameters.** In this section, we discuss concrete parameters used in our new PSI semi-honest protocol. We use  $\kappa = 128$  and  $\lambda = 40$ . The protocol contains several parameters:

*The length of OPRF output.* The output domain of PRF would be  $\{0, 1\}^v$  where  $v = \lambda + 2 \log_2(n)$  guarantees a  $2^{-\lambda}$  bound on the collision probability of PRF outputs among the two size- $n$  sets. Furthermore, communicating the hashes can be reduced to communicating only  $\approx \lambda + \log n$  bits per hash, using a heuristic technique of [31] that directly leads to an optimization of our PSI protocol.

*The size of  $\mathbb{F}_p$  and  $\mathbb{F}_q$  in BaRK-OPRF.* After using permutation-based hashing, each element is mapped to a bin with a stored value in this bin, the bit-length reduces from  $\ell$  to  $\ell - \log N$ . The input set of BaRK-OPRF in PSI protocol constructs from stored values concatenating with some extra bits. Then the bit-length of an input element of BaRK-OPRF is computed as  $\ell - \log N + 1$  if  $n \geq 2^{20}$  or  $\ell - \log N + 2$  otherwise, i.e, the size of  $q = 2^{\ell - \log N + 1}$  or  $q = 2^{\ell - \log N + 2}$  respectively.

*Generalized Cuckoo hashing.* We use a  $(d, k)$ -general cuckoo hashing scheme without stash. The parameters are chosen such that the failure probability is  $2^{-\lambda}$ . When  $d = 1, k = 3$  these parameters are identical with KKRT except for the number of bins increases slightly to  $N = 1.3n$  which is a trade-off to obtain no stash. Even with the higher number of bins, our PSI protocol significantly outperforms KKRT.

To minimize the overall communication, we set  $k = 2$  to reduce the cost of sending  $k \cdot n$  PRF outputs. We used a Python script to simulate randomly assigning  $n$  values to  $N = c \cdot n$  bins using  $(d, 2)$ -Cuckoo hashing, for several values of  $d$  and  $c$ , and for  $n = 2^9, 2^{10}, 2^{11}, 2^{12}$ . For a value of  $c$  as low as 0.65, we never observed any insertion failure over  $10^7$  trials for each values of  $n$  (for  $n = 2^{12}$ , we could only do  $10^6$  trials), when using  $d = 3$  items per bins. For  $d = 2$ , the failure probability became noticeable already for  $c \approx 1$ . Based on known theoretical analysis of  $(d, k)$ -Cuckoo hashing, the failure probability is known to scale inverse polynomially with  $n$ . Therefore, we expect that for reasonably large values of  $n$  (e.g.  $n \geq 2^{18}$ ), our parameters should guarantee a failure probability significantly below  $2^{-40}$ .

*2-choice hashing.* Following the analysis of [20], we set the number  $N$  of bins to  $n/3$ , and the maximum load  $d = L + 1$  to 4. This guarantees a failure probability which we empirically estimate to be  $1/N^{L-1}$ , which is below  $2^{-40}$  for all values of  $n$  above  $2^{14}$ .

*Simple hashing.* Eventually, for simple hashing, we set arbitrarily the number of bins  $N$  to  $n/10$ , and derive the corresponding value of  $d$  from Figure 5 in [26]. We note that the parameters for simple hashing are much less heuristic than the other two, in that concrete bound can actually be achieved which are relatively close to the heuristic (computer-estimated) bounds. For example, [20] experimentally observes that for a  $2^{-40}$  failure probability, setting  $d = 47$  suffices when using  $N = n/10$  bins. Using a standard Chernoff bound, it is in fact straightforward to prove formally that  $d = 49$  already suffices to reach this failure probability, which is very close to the experimental bound. In contrast, experimental bounds in more complex hashing variants are typically much more distant from provable bounds. The choice of  $N = n/10$  is entirely arbitrary: any smaller  $N$  leads to

better communication, but requires using higher values of  $d$ , leading to worse computation (due to the need to perform  $N$  polynomial interpolations with degree- $d$  polynomial). This allows for a smooth tradeoff between communication and computation, where better computational power can be used to further reduce the communication. At the extreme end of the spectrum, using  $N = 1$  and  $d = n$  requires one expensive degree- $n$  polynomial interpolation, but can achieve extremely low communications, e.g.  $93n$  bits of communication for  $\ell = 32$  and  $n = 2^{20}$ .

**Efficiency.** We compare the communication of our protocols, using three hashing methods, on Table 1. Regarding computation, we provide a breakdown of the computation costs of our protocols in the Appendix of full version [9]. Briefly, though, compared to the protocol of [27], and when using a standard choice of parameters for our protocol (e.g.  $n = 2^{20}$ , and using generalized Cuckoo hashing with  $d = 3$  and  $N = 0.65n$ ), our protocol requires essentially a length- $1.9n$  VOLE (with a small subfield),  $0.65n$  degree-3 polynomial interpolations (roughly  $3n$  multiplications over a small field), and computing  $n$  hashes. In contrast, the enhanced version of [27] (using the OKVS of [13] and the VOLE of [11]) will require solving a linear system to set up an OKVS (this requires on the order of  $(1.3 \log n + \lambda)^3$  multiplications over  $\mathbb{F}_{2^{128}}$ , plus  $O(\lambda n)$  operations), computing a length- $1.3n$  VOLE (over  $\mathbb{F}_{2^{128}}$ ), and computing  $2n$  hashes. The cost of the VOLE dominates that of performing  $n$  hashes, so for sufficiently large set sizes ( $n \gg 2^{20}$ ), the protocol of [27] should become roughly 30% more efficient than our protocol computation-wise. For smaller sets (e.g.  $n \approx 2^{16}$ ), the cost of setting up the OKVS becomes more significant, requiring around  $20n$  field multiplications over  $\mathbb{F}_{2^{128}}$ , hence the computational efficiency of our protocol becomes roughly on par with that of [27]. Of course, real runtimes can vary due to e.g. cache misses, so these estimations should only be viewed as a first order approximation indicating that the computational efficiency of our protocols is close to that of [27] (but likely slightly larger).

In terms of computation, the main computational overhead comes from performing  $N$  polynomial interpolations of *only degree- $d$*  polynomials. Based on our analysis, to achieve  $2^{-\lambda} = 2^{-40}$  probability of insertion failure, the following parameters can be chosen:

- $N = 0.65n$  and  $d = 3$  for generalized Cuckoo hashing (GCH),
- $N = 0.33n$  and  $d = 4$  for two-choice hashing,
- $N = n/10$  and  $d \approx 46$  for simple hashing.

As the above illustrates, the cost of performing  $N$  polynomial interpolations will be very small for GCH, two-choice hashing, but becomes higher for simple hashing (though performing  $n/10$  degree-46 interpolations remains reasonably fast).

### 4.3 A malicious PSI from mOPRF

In this section, we propose a maliciously secure PSI protocol based on our BaRK-OPRF (section 4.1) and simple hashing combining a permutation-based hash

function. The PSI protocol is shown in figure 5 and its security against a corrupted adversary is proven in theorem 2. The estimated overhead communication cost of this PSI is  $Nd(\ell - \log N) + (\kappa + \log n)n + o(n)$ . Observe that the PSI protocol in section 4.2 is insecure against malicious settings since the general hashing scheme does not allow the simulation in ideal world. To handle this we use simple hashing schemes with only one permutation-based hash function. This protocol is constructed from the natural approach used recently in [10,20,21,27], i.e, Alice (a sender) and Bob (a receiver) invoke the  $\mathcal{F}_{\text{opr}}^{\text{prf}}$  then Bob gets the PRF values on his input and Alice enables to compute the PRF on any input so Alice computes on her input after that she sends these PRF values to Bob; Bob compares and outputs the intersection.

PARAMETERS:

- Alice (sender) and Bob (receiver) have respectively input set  $X = \{x_1, x_2, \dots, x_n\} \in \mathbb{F}_p$  and  $Y = \{y_1, y_2, \dots, y_n\} \in \mathbb{F}_p$ , all elements of bit-length  $\ell$ .
- A random hash functions  $h : \{0, 1\}^* \rightarrow [N]$ .
- A Permutation-based hashing  $\text{Per}_{h,X}$  maps a set  $X$  to table  $\mathcal{B}_X$  consisting of  $N$  bins such that each bin has  $d$  slots where  $Nd > |X|$ , and  $d = O(1)$ . Denote  $\text{Per}(x) := (i, x')$  where  $x'$  is the stored value of  $x$  in bin  $i$  which defined by  $h$  and  $x$  then  $\text{Per}^{-1}(i, x') = x$ .

PROTOCOL:

1. Bob uses  $\text{Per}$  to map  $Y$  to  $\mathcal{B}_Y$ , for each empty slot in each bin  $\mathcal{B}_Y[i]$ , put here a dummy item of length  $\ell - \log N$ .
2. Alice sends  $(\text{sender}, \text{id})$  and Bob sends  $(\text{receiver}, \text{id}, \mathcal{B}_Y)$  to  $\mathcal{F}_{\text{opr}}^{\text{prf}}$  then
  - Bob receives the  $Y' = \{F(i, y') \mid y' \in \mathcal{B}_Y[i]\}_{i \leq N}$ .
3. For each  $x \in X$ , Alice queries  $x$  to  $\mathcal{F}_{\text{opr}}^{\text{prf}}$  with corresponding input  $(i, x')$  such that  $\text{Per}(x) = (i, x')$ , then Alice gets  $F(i, x')$ . Alice sends to Bob

$$U = \{F(i, x') \mid x \in X \wedge \text{Per}(x) = (i, x')\}$$

4. Now for each  $y \in Y$ ,  $\text{Per}(y) = (i, y')$ , if  $F(i, y') \in U$  then Bob outputs  $y$  as an element in the intersection.

**Fig. 5.** Our malicious PSI protocol based on  $\mathcal{F}_{\text{opr}}^{\text{prf}}$

Intuitively, in a malicious setting, when the sender is corrupted, the simulation needs to extract the sender's input set  $X$  from the queries to  $\mathcal{F}_{\text{opr}}^{\text{prf}}$  and the set  $U$ . Denote  $F(y) := F(i, y')$  where  $\text{Per}(y) = (i, y')$  and the set of all elements queried to  $\mathcal{F}_{\text{opr}}^{\text{prf}}$  is  $X'$  where  $n' = |X'|$ . The extraction procedure is that  $X = \{x \in \mathbb{F}_p \mid x \in X' \wedge F(x) \in U\}$ . Observe that if there exist two distinct elements  $x_1, x_2 \in X'$  such that  $F(x_1) = F(x_2) \in U$  then more than one element is extracted to  $X$ . The probability of existing collision is  $2^{-v+2 \log n'}$  then one approach to avoid collision is choosing  $v = 2\kappa$ . However, when  $v = 2\kappa$ , the overhead communication cost significantly increases.

Therefore, another approach is that Sim only extracts elements  $x \in X'$  if its

PRF is distinct and appears in  $U$ , i.e,  $x \in X'$  such that  $F(x) \in U$  and  $\nexists x' \in X'$  where  $F(x) = F(x')$ . [27] proposed this simulation and claimed that if the output domain of PRF  $v = \kappa$  then this simulation is correct and can not be distinguishable from the real protocol. We point out the proof of [27] has a gap and show that the output of PRF should be  $\kappa + \log n$ .

Indeed, if there exist some  $x_1, x_2 \in X'$  such that  $F(x_1) = F(x_2)$  then Sim only needs to extract  $x_1, x_2$  when one of them is in  $Y$ . Let assume  $x_1 \in Y$ , the probability of  $F(x_2) = F(y)$  for some  $y \in Y$  is  $2^{-v+\log(n_Y)}$  since  $Y$  is first fixed before the function  $F$  is sampled. [27] shows  $n_Y = O(\kappa)$  then the security can hold if  $v = \kappa$ . However, this should be  $v = \kappa + \log n_Y$  since  $n_Y = O(\text{poly}(\kappa))$  instead of  $O(\kappa)$ . In particular, PSI protocols in [27] are targeted on large input set because of the usage of vector OLE.

**Theorem 2.** *The PSI protocol on Figure 5 securely realizes the ideal functionality  $\mathcal{F}_{\text{psi}}$  over the field  $\mathbb{F}_p$  for set size  $n$  and malicious set size  $n_X = n$ ,  $n_Y = Nd$  with statistical security against malicious adversaries in  $\mathcal{F}_{\text{opr}}^{\text{hybrid}}$  model.*

In general, the malicious PSI (figure 5) has a communication cost that depends on the security parameter  $\kappa$  and is dominated by  $\kappa n$ . We now present a new PSI protocol that is secure in malicious setting via a dual execution while its communication cost only depends on the statistic parameter  $\lambda$  and the set size  $n$ . The idea of using a dual execution has been used in [26] but when combining this with our BaRK-OPRF it achieves efficient results, i.e, the total communication cost is only  $2Nd(\ell - \log N) + n(\lambda + \log n) + o(n)$ . The detailed construction of dual PSI is shown in the Appendix of full version [9].

## 5 A standard PSI from subfield-ring OLE

In this section, we describe a new PSI protocol, which builds upon a (simple variant of) a pseudorandom correlation generator for the ring-OLE correlation [7]. Our protocol enjoys a number of important features: it is in the *standard model*, achieves *malicious security* at essentially no cost, has *low communication* (competitive even with the best maliciously secure PSI protocols in the random oracle model), and reasonable computation (albeit superlinear in  $n$ ). Our protocol can also be generalized to a powerful notion of *batch non-interactive PSI*, where (after a small logarithmic-cost preprocessing step with each server) a client can broadcast a single encoding of his database, and then obtain the intersection with any of the server databases at any time after receiving a single message from this server. We believe that this functionality itself is of independent interest.

### 5.1 Semi-Honest Batch Non-Interactive PSI from Subfield Ring-OLE

We describe a new PSI scheme in the semi-honest model. Our protocol enjoys two interesting features: (1) it is in the standard model, and (2) it is a *batch non-interactive* protocol, a useful communication pattern which we describe afterwards. The full construction is represented on Figure 5.1.



**Theorem 3.** *The PSI protocol on Figure 5.1 securely realizes the ideal functionality  $\mathcal{F}_{\text{psi}}$  over the field  $\mathbb{F}_p$  with set size  $n$  and malicious set size  $n' = n_X = n_Y = 2n$ , with statistical security against augmented semi-honest adversaries in the  $\mathcal{F}_{\text{sole}}$  hybrid model.*

Above, “augmented semi-honest security” refers to a strengthening of honest-but-curious corruption where the adversary is allowed to change the corrupted parties’ inputs. This is a standard strengthening of semi-honest security, which has been argued to better capture real-world security [16]. It will also facilitate upgrading security to the malicious setting later on.

**Batch non-interactivity.** To securely realize the functionality  $\mathcal{F}_{\text{sole}}$ , we rely on the PCG-based protocol of [7] (using a straightforward adaptation to the subfield setting), which is secure under the ring-LPN assumption. Interestingly, instantiating the subfield ring OLE this way allows to import a powerful feature of the PCG of [7], which is its *programmability*: when generating a ring-OLE correlation, the receiver can ensure that her output  $a$  remains *identical* across multiple instances of the protocol with different parties.

This feature enables the following communication structure: after a short (logarithmic-communication) interaction with  $N$  servers, a client, playing the role of Alice with input set  $A$ , can broadcast a single compact encoding of her dataset to all the servers (with input sets  $B_1 \cdots B_N$ ). Afterwards, each server  $B_i$  can at any time send a single message  $m_i$  to Alice, from which she can recover  $A \cap B_i$  without further interaction. To our knowledge, this batch non-interactive communication pattern was never achieved by any prior proposal; we believe that it can make our protocol appealing in realistic scenarios.

More concretely, after a logarithmic-communication preprocessing phase where Alice sets up PCG seeds with each of servers, Alice broadcasts the value  $t_A = a - p_A$  to everyone, which communicates  $2n \log p \approx 2\ell n$  bits. This message can be seen as a compact public encoding of her dataset (it is only twice as large as Alice’s set). Afterwards, each server can complete the protocol of Figure 5.1 by sending a single message  $(b_1, t_B)$  to the receiver, of length  $3n \log q \approx 3(\lambda + 2 \log n)n$ , from which the receiver can locally recover  $X \cap X_i$ . Furthermore, using the encoding technique of [31], the  $\lambda + 2 \log n$  term can be reduced to  $\lambda + \log n$  (the improvement is based on the observation that for an appropriate ordering,  $n$  random elements of a set of size  $2^{\lambda + 2 \log n}$  are on average at distance  $2^{\lambda + \log n}$  for each other, hence the cost of transmitting them can be reduced to essentially  $\lambda + \log n$  per element by sending the distance between consecutive elements instead).

**Efficiency.** The communication cost of protocol (Figure 5.1) is  $n \cdot (2 \log p + 3 \log q) + o(n)$  bits of communication. Here, the size of the subfield  $\mathbb{F}_p$  depends only on the bitsize  $\ell$  of the items in the sets  $A$  and  $B$ , hence we can set  $\log p = \ell$ . As we will see in the analysis,  $\log q$  must be set to  $\log q \approx \lambda + 2 \log n$  to guarantee  $\lambda$  bits of statistical security. This leads to a total communication of  $n \cdot (2\ell + 3\lambda + 6 \log n) + o(n)$  bits, which is reduced to  $n \cdot (2\ell + 3\lambda + 3 \log n) + o(n)$  with the encoding of [31]. The  $o(n)$  term above captures the cost of distributing

the PCG seeds of the subfield ring-OLE (we discuss the concrete value of  $o(n)$  later on, for our maliciously secure version of the protocol).

Regarding computation, the computational cost scales as  $O(n \log^2 n)$  due to the fast polynomial interpolations, or as  $O(n \log n)$  when using cyclotomic rings. We provide a concrete analysis of the computational cost of the maliciously secure version of our protocol in Section 5.2.

PARAMETERS:

- Two rings  $\mathcal{R}_p = \mathbb{F}_p[X]/F(X) \subseteq \mathcal{R}_q = \mathbb{F}_{p^t}[X]/F(X)$ , where  $F(X)$  has degree  $2n + 1$ .
- The sender (Alice) and receiver (Bob) have respective input sets  $A = \{a_1, a_2, \dots, a_n\} \subset \mathbb{F}_p$  and  $B = \{b_1, b_2, \dots, b_n\} \subset \mathbb{F}_p$ .
- A subfield ring-OLE in the ring  $\mathcal{R}_q$  over the subring  $\mathcal{R}_p$ .

PROTOCOL:

1. **(Setting up the correlation)** Alice and Bob encode their sets to  $p_A = \prod_{i=1}^n (X - a_i)$ ,  $p_B = \prod_{i=1}^n (X - b_i)$  respectively, and invoke  $\mathcal{F}_{\text{sole}}$  to generate a subfield ring-OLE correlation over  $\mathcal{R}_p, \mathcal{R}_q$ : Alice receives  $(a, s_A) \in \mathcal{R}_p \times \mathcal{R}_q$  and Bob receives  $(b, s_B) \in \mathcal{R}_q^2$  such that  $s_A + s_B = ab$ .
2. **(Broadcasting the client set encoding)** Alice computes and sends  $t_A = a - p_A$  to Bob.
3. **(Server-to-client message)** Bob sets  $s'_B \leftarrow s_B - t_A b$ . Then, Bob decomposes  $b$  as  $b = b_0 + b_1 \cdot X^n$  (where  $b_0, b_1$  are degree- $n$  polynomials), sets  $s'_B \leftarrow s_B - t_A b$ , and picks a random degree- $n$  polynomial  $b'_0$  over  $\mathcal{R}_q$ . He sends  $b_1$  and  $t_B \leftarrow s'_B + p_B b'_0$  to Alice.
4. **(Output)** Alice sets  $u \leftarrow t_B - p_A b_1 \cdot X^n + s_A$ ; note that  $u = p_A b_0 + p_B b'_0$ . Alice outputs the set  $I = \{x \in A \mid u(x) = 0\}$ .

**Fig. 6.** Augmented semi-honest PSI protocol based on ring-OLE

## 5.2 Maliciously Secure PSI in the Standard Model

In this section, we upgrade the security of our protocol to the malicious setting. Our upgrade introduces only a minimal communication overhead to the protocol, independent of the set sizes  $n$ . The full protocol is represented on Figure 5.2.

**Theorem 4.** *The PSI protocol on Figure 5.2 securely realizes the ideal functionality  $\mathcal{F}_{\text{psi}}$  over the field  $\mathbb{F}_p$  with set size  $n$  and malicious set size  $n' = n_X = n_Y = 2n$ , with statistical security against malicious adversaries in the  $\mathcal{F}_{\text{sole}}$ -hybrid model.*

**Efficiency.** Our malicious protocol has minimal communication overhead over our augmented semi-honest protocol. The main overhead stems from starting from a slightly larger field in which two elements can be “reserved elements”. If  $p'$  is a prime power and  $\ell \approx \log p'$ , the price to pay is therefore increasing  $\ell$  to  $\log p$  where  $p$  is the smallest prime power above  $p' + 2$ . While an exact expression would be rather tedious, for any reasonable input size this cost should be negligible (the

## PARAMETERS:

- A field  $\mathbb{F}_{p'}$  and two rings  $\mathcal{R}_p = \mathbb{F}_p[X]/F(X) \subseteq \mathcal{R}_q = \mathbb{F}_{p'}[X]/F(X)$ , where  $F(X)$  has degree  $2n + 1$  and  $|\mathbb{F}_{p'}| \leq |\mathbb{F}_p| - 2$ .  $\mathbf{map}$  is an efficient (and efficiently invertible) injective mapping, with  $\mathbf{map}(\mathbb{F}_{p'}) \subseteq \mathbb{F}_p \setminus \{0, 1\}$ .
- The sender (Alice) and receiver (Bob) have respective input sets  $A = \{a_1, a_2, \dots, a_n\} \subset \mathbb{F}_{p'}$  and  $B = \{b_1, b_2, \dots, b_n\} \subset \mathbb{F}_{p'}$ .
- A subfield ring-OLE in the ring  $\mathcal{R}_q$  over the subring  $\mathcal{R}_p$ .

## PROTOCOL:

1. **(Setting up the correlation)** Alice and Bob encode their sets to  $p_A = c \cdot (X - 1) \cdot \prod_{i=1}^n (X - \mathbf{map}(a_i))$  with  $c = -(\prod_{i=1}^n (-\mathbf{map}(a_i)))^{-1}$  (note that this guarantees  $p_A(0) = 1$  and  $p_A(1) = 0$ ) and  $p_B = \prod_{i=1}^n (X - \mathbf{map}(b_i))$  respectively. Alice and Bob invoke  $\mathcal{F}_{\text{sole}}$  to generate a subfield ring-OLE correlation over  $\mathcal{R}_p, \mathcal{R}_q$ : Alice receives  $(a, s_A) \in \mathcal{R}_p \times \mathcal{R}_q$  and Bob receives  $(b, s_B) \in \mathcal{R}_q^2$  such that  $s_A + s_B = ab$ .
2. **(Broadcasting the client set encoding)** Alice computes and sends  $t_A = a - p_A$  to Bob.
3. **(Server-to-client message)** Bob sets  $s'_B \leftarrow s_B - t_A b$ . Then, Bob decomposes  $b$  as  $b = b_0 + b_1 \cdot X^n$  (where  $b_0, b_1$  are degree- $n$  polynomials), and sets  $s'_B \leftarrow s_B - t_A b$ . He sends  $b_1$  to Alice.
4. **(Checking  $p_A$ )** Alice computes  $s'_A \leftarrow s_A - p_A b_1 \cdot X^n$ . Alice sends  $y \leftarrow s'_A(0)$  to Bob. If  $y \neq b_0(0) - s'_B(0)$ , Bob aborts. Else, Bob picks a random degree- $n$  polynomial  $b'_0$  over  $\mathcal{R}_q$  and sends  $t_B \leftarrow s'_B + p_B b'_0$  to Alice.
5. **(Output)** Alice sets  $u \leftarrow t_B - p_A b_1 \cdot X^n + s_A$ ; note that  $u = p_A b_0 + p_B b'_0$ . If  $u(1) = 0$ , Alice aborts; otherwise, Alice computes the set  $I = \{x \in A \mid u(\mathbf{map}(x)) = 0\}$  and outputs  $I$ .

Fig. 7. Maliciously secure PSI protocol in the  $\mathcal{F}_{\text{sole}}$ -hybrid model

simplest strategy is to pick  $p' = 2^\ell$  and  $p = 2^{\ell+1}$ , in which case  $\ell$  is increased by one bit, but much better encoding methods exist). Therefore, the communication remains  $n \cdot (2\ell + 3\lambda + 6 \log n) + o(n)$  bits, or  $n \cdot (2\ell + 3\lambda + 3 \log n) + o(n)$  with the encoding of [31]. We provide a more concrete analysis of the  $o(n)$  term (setting up the ring-OLE) in the malicious setting in the Appendix of full version [9].

*Computation cost.* Note that our standard model protocol shares with our other protocols the feature of having a communication independent of  $\kappa$ . Our protocol requires more computation compared to the best ROM-based protocols, due to its use of polynomial interpolation. However, it still allows for very fast PSI computation (we estimate a few seconds to compute the intersection between databases of size  $2^{20}$ , on one core of a standard laptop). Concretely, the protocol requires only

- a single degree- $n$  polynomial interpolation, one FFT over a polynomial ring with degree- $2n$  polynomials, and 3 multiplications of degree- $n$  polynomials for the receiver, and
- a single degree- $n$  polynomial interpolation, one FFT as above, 2 multiplications of degree- $n$  polynomials, and a single  $n$ -multipoint polynomial evaluation for the sender.

Furthermore, both polynomial interpolations only have to be performed over a field  $\mathbb{F}$ , of size  $|\mathbb{F}| \approx 2^\ell$  where  $\ell$  is the bit size of the set items (e.g. 32 or 64 bits), and the multipoint evaluation is over a field of size  $\lambda + 2 \log n$  bits. This stands in stark contrasts with previous state of the art protocols [20] that relied on polynomial interpolation (*on top* of using the ROM), where the interpolations and multipoint evaluations had to be performed over a very large field  $\mathbb{F}$  of size  $|\mathbb{F}| \approx 2^{400}$ . By using a cyclotomic ring, the FFTs and polynomial multiplications are much faster than the interpolations. On Table 1.1, we compare our protocol to the current fastest maliciously secure PSI protocols [21, 27, 29].

**On the attacks of [1].** We note that constructing maliciously secure PSI protocols using an algebraic approach, along the lines of our protocol, is known to be non-trivial and error prone. Indeed, previous works [14] used a similar approach based on polynomial manipulation, OLEs, and the lemmas about the polynomial (appear in the Appendix of full version [9]), to build a malicious PSI protocol. However, their protocol was found to be insecure in a recent preprint, which described powerful concrete attacks on this proposal [1]. Intuitively, the key technical difficulties revolve in both cases around how to handle null polynomials ( $p_A = 0$  or  $p_B = 0$ ). In our specific context, it turns out that our direct use of ring-OLE enables relatively elegant and simple (in hindsight) strategies to enforce nonzero polynomials. Our modification has almost no impact on the communication or the computation of our protocol, essentially giving us malicious security for free (though we note that we still require an additional round of communication). It is not, however, completely clear how to adapt our strategy to the setting of OLE-based algebraic PSI in [14]. We believe that this provides further support for the intuition that ring-OLE is the right primitive to build PSI protocols using this algebraic approach (beyond its direct advantage in terms of communication efficiency).

## Acknowledgement

We thank all reviewers for their helpful comments. The first author is supported by Dim Math Innov funding from the Paris Mathematical Sciences Foundation (FSMP) funded by the Paris Ile-de-France Region, and the second author acknowledges the support of the French Agence Nationale de la Recherche (ANR), under grant ANR-20-CE39-0001 (project SCENE). This work was also supported by the France 2030 ANR Project ANR-22-PECY-003 SecureCompute.

## References

1. Abadi, A., Murdoch, S.J., Zacharias, T.: Polynomial representation is tricky: Maliciously secure private set intersection revisited. Cryptology ePrint Archive, Report 2021/1009 (2021), <https://ia.cr/2021/1009>
2. Applebaum, B., Damgård, I., Ishai, Y., Nielsen, M., Zichron, L.: Secure arithmetic computation with constant computational overhead. In: Katz, J., Shacham, H. (eds.) CRYPTO 2017, Part I. LNCS, vol. 10401, pp. 223–254. Springer, Heidelberg (Aug 2017)

3. Boyle, E., Couteau, G., Gilboa, N., Ishai, Y.: Compressing vector OLE. In: Lie, D., Mannan, M., Backes, M., Wang, X. (eds.) ACM CCS 2018. pp. 896–912. ACM Press (Oct 2018)
4. Boyle, E., Couteau, G., Gilboa, N., Ishai, Y., Kohl, L., Rindal, P., Scholl, P.: Efficient two-round OT extension and silent non-interactive secure computation. In: Cavallaro, L., Kinder, J., Wang, X., Katz, J. (eds.) ACM CCS 2019. pp. 291–308. ACM Press (Nov 2019)
5. Boyle, E., Couteau, G., Gilboa, N., Ishai, Y., Kohl, L., Scholl, P.: Efficient pseudorandom correlation generators: Silent OT extension and more. In: Boldyreva, A., Micciancio, D. (eds.) CRYPTO 2019, Part III. LNCS, vol. 11694, pp. 489–518. Springer, Heidelberg (Aug 2019)
6. Boyle, E., Couteau, G., Gilboa, N., Ishai, Y., Kohl, L., Scholl, P.: Correlated pseudorandom functions from variable-density LPN. In: 61st FOCS. pp. 1069–1080. IEEE Computer Society Press (Nov 2020)
7. Boyle, E., Couteau, G., Gilboa, N., Ishai, Y., Kohl, L., Scholl, P.: Efficient pseudorandom correlation generators from ring-LPN. In: Micciancio, D., Ristenpart, T. (eds.) CRYPTO 2020, Part II. LNCS, vol. 12171, pp. 387–416. Springer, Heidelberg (Aug 2020)
8. Boyle, E., Couteau, G., Gilboa, N., Ishai, Y., Orrù, M.: Homomorphic secret sharing: Optimizations and applications. In: Thuraisingham, B.M., Evans, D., Malkin, T., Xu, D. (eds.) ACM CCS 2017. pp. 2105–2122. ACM Press (Oct / Nov 2017)
9. Bui, D., Couteau, G.: Improved private set intersection for sets with small entries. Cryptology ePrint Archive, Paper 2022/334 (2022), <https://eprint.iacr.org/2022/334>, <https://eprint.iacr.org/2022/334>
10. Chase, M., Miao, P.: Private set intersection in the internet setting from lightweight oblivious PRF. In: Micciancio, D., Ristenpart, T. (eds.) CRYPTO 2020, Part III. LNCS, vol. 12172, pp. 34–63. Springer, Heidelberg (Aug 2020)
11. Couteau, G., Rindal, P., Raghuraman, S.: Silver: Silent VOLE and oblivious transfer from hardness of decoding structured LDPC codes. pp. 502–534. LNCS, Springer, Heidelberg (2021)
12. Dietzfelbinger, M., Weidling, C.: Balanced allocation and dictionaries with tightly packed constant size bins. Theoretical Computer Science 380(1-2), 47–68 (2007)
13. Garimella, G., Pinkas, B., Rosulek, M., Trieu, N., Yanai, A.: Oblivious key-value stores and amplification for private set intersection. pp. 395–425. LNCS, Springer, Heidelberg (2021)
14. Ghosh, S., Nilges, T.: An algebraic approach to maliciously secure private set intersection. In: Ishai, Y., Rijmen, V. (eds.) EUROCRYPT 2019, Part III. LNCS, vol. 11478, pp. 154–185. Springer, Heidelberg (May 2019)
15. Ghosh, S., Simkin, M.: The communication complexity of threshold private set intersection. In: Boldyreva, A., Micciancio, D. (eds.) CRYPTO 2019, Part II. LNCS, vol. 11693, pp. 3–29. Springer, Heidelberg (Aug 2019)
16. Hazay, C., Lindell, Y.: A note on the relation between the definitions of security for semi-honest and malicious adversaries. Cryptology ePrint Archive, Report 2010/551 (2010), <https://eprint.iacr.org/2010/551>
17. Kissner, L., Song, D.X.: Privacy-preserving set operations. In: Shoup, V. (ed.) CRYPTO 2005. LNCS, vol. 3621, pp. 241–257. Springer, Heidelberg (Aug 2005)
18. Kolesnikov, V., Kumaresan, R., Rosulek, M., Trieu, N.: Efficient batched oblivious PRF with applications to private set intersection. In: Weippl, E.R., Katzenbeisser, S., Kruegel, C., Myers, A.C., Halevi, S. (eds.) ACM CCS 2016. pp. 818–829. ACM Press (Oct 2016)

19. Kolesnikov, V., Rosulek, M., Trieu, N., Wang, X.: Scalable private set union from symmetric-key techniques. In: Galbraith, S.D., Moriai, S. (eds.) ASIACRYPT 2019, Part II. LNCS, vol. 11922, pp. 636–666. Springer, Heidelberg (Dec 2019)
20. Pinkas, B., Rosulek, M., Trieu, N., Yanai, A.: SpOT-light: Lightweight private set intersection from sparse OT extension. In: Boldyreva, A., Micciancio, D. (eds.) CRYPTO 2019, Part III. LNCS, vol. 11694, pp. 401–431. Springer, Heidelberg (Aug 2019)
21. Pinkas, B., Rosulek, M., Trieu, N., Yanai, A.: PSI from PaXoS: Fast, malicious private set intersection. In: Canteaut, A., Ishai, Y. (eds.) EUROCRYPT 2020, Part II. LNCS, vol. 12106, pp. 739–767. Springer, Heidelberg (May 2020)
22. Pinkas, B., Schneider, T., Segev, G., Zohner, M.: Phasing: Private set intersection using permutation-based hashing. In: Jung, J., Holz, T. (eds.) USENIX Security 2015. pp. 515–530. USENIX Association (Aug 2015)
23. Pinkas, B., Schneider, T., Weinert, C., Wieder, U.: Efficient circuit-based PSI via cuckoo hashing. In: Nielsen, J.B., Rijmen, V. (eds.) EUROCRYPT 2018, Part III. LNCS, vol. 10822, pp. 125–157. Springer, Heidelberg (Apr / May 2018)
24. Pinkas, B., Schneider, T., Zohner, M.: Faster private set intersection based on OT extension. In: Fu, K., Jung, J. (eds.) USENIX Security 2014. pp. 797–812. USENIX Association (Aug 2014)
25. Rindal, P., Raghuraman, S.: Blazing fast PSI from improved OKVS and subfield VOLE. IACR Cryptol. ePrint Arch. p. 320 (2022), <https://eprint.iacr.org/2022/320>
26. Rindal, P., Rosulek, M.: Malicious-secure private set intersection via dual execution. In: Thuraisingham, B.M., Evans, D., Malkin, T., Xu, D. (eds.) ACM CCS 2017. pp. 1229–1242. ACM Press (Oct / Nov 2017)
27. Rindal, P., Schoppmann, P.: VOLE-PSI: Fast OPRF and circuit-PSI from vector-OLE. pp. 901–930. LNCS, Springer, Heidelberg (2021)
28. Rosulek, M., Trieu, N.: Compact and malicious private set intersection for small sets. Cryptology ePrint Archive, Report 2021/1159 (2021), <https://eprint.iacr.org/2021/1159>
29. Rosulek, M., Trieu, N.: Compact and malicious private set intersection for small sets. Cryptology ePrint Archive, Report 2021/1159 (2021), <https://ia.cr/2021/1159>
30. Schoppmann, P., Gascón, A., Reichert, L., Raykova, M.: Distributed vector-OLE: Improved constructions and implementation. In: Cavallaro, L., Kinder, J., Wang, X., Katz, J. (eds.) ACM CCS 2019. pp. 1055–1072. ACM Press (Nov 2019)
31. Tamrakar, S., Liu, J., Paverd, A., Ekberg, J.E., Pinkas, B., Asokan, N.: The circle game: Scalable private membership test using trusted hardware. In: Karri, R., Sinanoglu, O., Sadeghi, A.R., Yi, X. (eds.) ASIACCS 17. pp. 31–44. ACM Press (Apr 2017)
32. Weng, C., Yang, K., Katz, J., Wang, X.: Wolverine: fast, scalable, and communication-efficient zero-knowledge proofs for boolean and arithmetic circuits. In: 2021 IEEE Symposium on Security and Privacy (SP). pp. 1074–1091. IEEE (2021)
33. Wieder, U., et al.: Hashing, load balancing and multiple choice. Foundations and Trends® in Theoretical Computer Science 12(3–4), 275–379 (2017)
34. Yang, K., Weng, C., Lan, X., Zhang, J., Wang, X.: Ferret: Fast extension for correlated OT with small communication. In: Ligatti, J., Ou, X., Katz, J., Vigna, G. (eds.) ACM CCS 20. pp. 1607–1626. ACM Press (Nov 2020)