# Efficient Circuit-Based PSI with Linear Communication 

Benny Pinkas ${ }^{1(\boxtimes)}$, Thomas Schneider ${ }^{2}$, Oleksandr Tkachenko ${ }^{2}$, and Avishay Yanai ${ }^{1}$<br>${ }^{1}$ Bar-Ilan University, Ramat Gan, Israel benny@pinkas.net, ay.yanay@gmail.com<br>${ }^{2}$ TU Darmstadt, Darmstadt, Germany<br>\{schneider,tkachenko\}@encrypto.cs.tu-darmstadt.de


#### Abstract

We present a new protocol for computing a circuit which implements the private set intersection functionality (PSI). Using circuits for this task is advantageous over the usage of specific protocols for PSI, since many applications of PSI do not need to compute the intersection itself but rather functions based on the items in the intersection.

Our protocol is the first circuit-based PSI protocol to achieve linear communication complexity. It is also concretely more efficient than all previous circuit-based PSI protocols. For example, for sets of size $2^{20}$ it improves the communication of the recent work of Pinkas et al. (EUROCRYPT'18) by more than 10 times, and improves the run time by a factor of 2.8 x in the LAN setting, and by a factor of 5.8 x in the WAN setting.

Our protocol is based on the usage of a protocol for computing oblivious programmable pseudo-random functions (OPPRF), and more specifically on our technique to amortize the cost of batching together multiple invocations of OPPRF.


Keywords: Private Set Intersection • Secure computation

## 1 Introduction

The functionality of Private Set Intersection (PSI) enables two parties, $P_{1}$ and $P_{2}$, with respective input sets $X$ and $Y$ to compute the intersection $X \cap Y$ without revealing any information about the items which are not in the intersection. There exist multiple constructions of secure protocols for computing PSI, which can be split into two categories: (i) constructions that output the intersection itself and (ii) constructions that output the result of a function $f$ computed on the intersection. In this work, we concentrate on the second type of constructions (see Sect. 1.2 for motivation). These constructions keep the intersection $X \cap Y$ secret from both parties and allow the function $f$ to be securely computed on top of it, namely, yielding only $f(X \cap Y)$. Formally, denote by $\mathcal{F}_{\mathrm{PSI}, f}$ the functionality $(X, Y) \mapsto(f(X \cap Y), f(X \cap Y))$.

A functionality for computing $f(X \cap Y)$ can be naively implemented using generic MPC protocols by expressing the functionality as a circuit. However, naive protocols for computing $f(X \cap Y)$ have high communication complexity, which is of paramount importance for real-world applications. The difficulty in designing a circuit for computing the intersection is in deciding which pairs of items of the two parties need to be compared. We refer here to the number of comparisons computed by the circuit as the major indicator of the overhead, since it directly affects the amount of communication in the protocol (which is proportional to the number of comparisons, times the length of the representation of the items, times the security parameter). Since the latter factors (input length and security parameter) are typically given, and since the circuit computation mostly involves symmetric key operations, the goal is to minimize the communication overhead as a function of the input size. We typically state this goal as minimizing the number of comparisons computed in the circuit. The protocol presented in this paper is the first to achieve linear communication overhead, which is optimal.

Suppose that each party has an input set of $n$ items. A naive circuit for this task compares all pairs and computes $O\left(n^{2}\right)$ comparisons. More efficient circuits are possible, assuming that the parties first order their respective inputs in specific ways. For example, if each party has sorted its input set then the intersection can be computed using a circuit which first computes, using a merge-sort network, a sorted list of the union of the two sets, and then compares adjacent items [HEK12]. This circuit computes only $O(n \log n)$ comparisons. The protocol of [PSSZ15] (denoted "Circuit-Phasing") has $P_{1}$ map its items to a table using Cuckoo hashing, and $P_{2}$ maps its items using simple hashing. The intersection is computed on top of these tables by a circuit with $O(n \log n / \log \log n)$ comparisons. This protocol is the starting point of our work.

A recent circuit-based PSI construction [PSWW18] is based on a new hashing algorithm, denoted "two-dimensional Cuckoo hashing", which uses a table of size $O(n)$ and a stash of size $\omega(1)$. Each party inserts its inputs to a separate table, and the hashing scheme assures that each value in the intersection is mapped by both parties to exactly one mutual bin. Hence, a circuit which compares the items that the two parties mapped to each bin, and also compares all stash items to all items of the other party, computes the intersection in only $\omega(n)$ comparisons (namely, the overhead is slightly more than linear, although it can be made arbitrarily close to being linear).

Our work is based on the usage of an oblivious programmable pseudo-random function (OPPRF), which is a new primitive that was introduced in $[\mathrm{KMP}+17]$. An OPRF-oblivious pseudo-random function (note, this is different than an OPPRF) - is a two-party protocol where one party has a key to a PRF $F$ and the other party can privately query $F$ at specific locations. An OPPRF is an extension of the protocol which lets the key owner "program" $F$ so that it has specific outputs for some specific input values (and is pseudo-random on all other values). The other party which evaluates the OPPRF does not learn whether it learns a "programmed" output of $F$ or just a pseudo-random value.

### 1.1 Overview of Our Protocol

The starting point for our protocols is the Circuit-Phasing PSI protocol of [PSSZ15], in which $O(n)$ bins are considered and the circuit computes $O(n \log n / \log \log n)$ comparisons. Party $P_{1}$ uses Cuckoo hashing to map at most one item to each bin, whereas party $P_{2}$ maps its items to the bins using simple hashing (two times, once with each of the two functions used in the Cuckoo hashing of the first party). Thus, $P_{2}$ maps up to $S=O(\log n / \log \log n)$ items to each bin. Since the parties have to hide the number of items that are mapped to each bin, they pad the bins with "dummy" items to the maximum bin size. That is, $P_{1}$ pads all bins so they all contain exactly one item and $P_{2}$ pads all bins so they all contain $S$ items.

Both parties use the same hash functions, and therefore for each input element $x$ that is owned by both parties there is exactly one bin to which $x$ is mapped by both parties. Thus, it is only needed to check whether the item that $P_{1}$ places in a bin is among the items that are placed in this bin by $P_{2}$. This is essentially a private set membership (PSM) problem: As input, $P_{1}$ has a single item $x$ and $P_{2}$ has a set $\Sigma$ with $|\Sigma|$ items, where $S=|\Sigma|$. As for the output, if $x \in \Sigma$ then both parties learn the same random output, otherwise they learn independent random outputs. These outputs can then be fed to a circuit, which computes the intersection. The Circuit-Phasing protocol [PSSZ15] essentially computes the PSM functionality using a sub-circuit of the overall circuit that it computes. Namely, let $S=O(\log n / \log \log n)$ be an upper bound on the number of items mapped by $P_{2}$ to a single bin. For each bin the sub-circuit receives one input from $P_{1}$ and $S$ inputs from $P_{2}$, computes $S$ comparisons, and feeds the result to the main part of the circuit which computes the intersection itself (and possibly some function on top of the intersection). Therefore the communication overhead is $O(n S)=O(n \log n / \log \log n)$. A very recent work in [CO18] uses the same hashing method and computes the PSM using a specific protocol whose output is fed to the circuit. The circuit there computes only $\omega(n)$ comparisons but the PSM protocol itself incurs a communication overhead of $O(\log n / \log \log n)$ and is run $O(n)$ times. Therefore, the communication overhead of [CO18] is also $O(n \log n / \log \log n)$.

We diverge from the protocol of [PSSZ15] in the method for comparing the items mapped to each bin. In our protocol, the parties run an oblivious programmable PRF (OPPRF) protocol for each bin $i$, such that party $P_{2}$ chooses the PRF key and the programmed values, and the first party learns the output. The function is "programmed" to give the same output $\beta_{i}$ for each of the $O(\log n / \log \log n)$ items that $P_{2}$ mapped to this bin. Therefore, if there is any match in this bin then $P_{1}$ learns the same value $\beta_{i}$. Then, the parties evaluate a circuit, where for each bin $i$ party $P_{1}$ inputs its output in the corresponding OPPRF protocol, and $P_{2}$ inputs $\beta_{i}$. This circuit therefore needs to compute only a single comparison per bin.

The communication overhead of an OPPRF is linear in the number of programmed values. Thus, a stand alone invocation of an OPPRF for every bin incurs an overall overhead of $O(n \log n / \log \log n)$. We achieve linear overhead
for comparing the items in all bins, by observing that although each bin is of maximal size $O(\log n / \log \log n)$ (and therefore naively requires to program this number of values in the OPPRF), the total number of items that need to be programmed in all bins is $O(n)$. We can amortize communication so that the total communication of computing all $O(n)$ OPPRFs is the same as the total number of items, which is $O(n)$.

In addition to comparing the items that are mapped to the hash tables, the protocol must also compare items that are mapped to the stash of the Cuckoo hashing scheme. Fixing a stash size $s=O(1)$, the probability that the stash does not overflow is $O\left(n^{-(s+1)}\right)$ [KMW09]. It was shown in [GM11] that a stash of size $O(\log n)$ ensures a negligible failure probability (namely, a probability that is asymptotically smaller than any polynomial function). Each item that $P_{1}$ places in the stash must be compared to all items of $P_{2}$, and therefore a straightforward implementation of this step requires the circuit to compute $\omega(n)$ comparisons. However, we show an advanced variant of our protocol that computes all comparisons (including elements in the stash) with only $O(n)$ comparisons.

In addition to designing a generic $O(n)$ circuit-based PSI protocol, we also investigate an important and commonly used variant of the problem where each item is associated with some value ("payload"), and it is required to compute a function of the payloads of the items in the intersection. (For example, compute the sum of financial transactions associated with these items). The challenge is that each of the $S$ items that the second party maps to a bin has a different payload and therefore it is hard to represent them using a single value. (The work in [PSSZ15, CO18], for example, did not consider payloads). We describe a variant of our PSI protocol which injects the correct payloads to the circuit while keeping the $O(n)$ overhead.

Overall, the work in this paper improves the state of the art in two dimensions:

- With regards to asymptotic performance, we show a protocol for circuitbased PSI which has only $O(n)$ communication. This cost is asymptotically smaller than that of all known circuit-based constructions of PSI, and matches the obvious lower bound on the number of comparisons that must be computed.
- With regards to concrete overhead, our most efficient protocols improve communication by a factor of 2.6 x to 12.8 x , and run faster by factor 2.8 x to 5.8 x compared to the previous best circuit-based PSI protocol of [PSWW18]. We demonstrate this both analytically and experimentally.


### 1.2 Motivation for Circuit-Based PSI

Most research on computing PSI focused on computing the intersection itself (see Sect.1.4). On the other hand, many applications of PSI are based on computing arbitrary functions of the intersection. For example, Google reported a PSI-based application for measuring the revenues from online ad viewers who later perform a related offline transaction (namely, ad conversion
rates) [Yun15, Kre17]. This computation compares the set of people who were shown an ad with the set of people who have completed a transaction. These sets are held by the advertiser, and by merchants, respectively. A typical use case is where the merchant inputs pairs of the customer-identity and the value of the transactions made by this customer, and the computation calculates the total revenue from customers who have seen an ad, namely customers in the intersection of the sets known to the advertiser and the merchant. Google reported implementing this computation using a Diffie-Hellman-based PSI cardinality protocol (for computing the cardinality of the intersection) and Paillier encryption (for computing the total revenues) [IKN+17,Kre18]. In fact, it was recently reported that Google is using such a "double-blind encryption" protocol in a beta version of their ads tool. ${ }^{1}$ However, their protocol reveals the size of the intersection, and has substantially higher runtimes than our protocol as it uses public key operations, rather than efficient symmetric cryptographic operations (cf. Sect. 7.4).

Another motivation for running circuit-based PSI is adaptability. A protocol that is specific for computing the intersection, or a specific function such as the cardinality of the intersection, cannot be easily changed to compute another function of the intersection (say, the cardinality plus some noise to preserve differential privacy). Any change to a specialized protocol will require considerable cryptographic know-how, and might not even be possible. On the other hand, the task of writing a new circuit component which computes a different function of the intersection is rather trivial.

Circuit-based protocols also benefit from the existing code base for generic secure computation. Users only need to design the circuit to be computed, and can use available libraries of optimized code for secure computation, such as [HEKM11,EFLL12,DSZ15,LWN+15].

### 1.3 Computing Symmetric Functions

We focus in this work on constructing a circuit which computes the intersection. On top of that circuit it is possible to compose a circuit for computing any function that is based on the intersection. In order to preserve privacy, that function must be a symmetric function of the items in the intersection. Namely, the output of the function must not depend on the order of its inputs.

If the function that needs to be computed is non-symmetric, then the circuit for computing the intersection must shuffle its output, in order to place each item of the intersection in a location which is independent of the other values. The result is used as the input to the function. The size of this "shuffle" step is $O(n \log n)$, as is described in [HEK12], and it dominates the $O(n)$ size of the intersection circuit. We therefore focus on the symmetric case. ${ }^{2}$

[^0]Most interesting functions of the intersection (except for the intersection itself) are symmetric. Examples of symmetric functions include:

- The size of the intersection, i.e., PSI cardinality (PSI-CA).
- A threshold function that is based on the size of the intersection. For example identifying whether the size of the intersection is greater than some threshold (PSI-CAT). An extension of PSI-CAT, where the intersection is revealed only if the size of the intersection is greater than a threshold, can be used for privacy-preserving ridesharing [HOS17]. Other public-key based protocols for this functionality appear in [ZC17,ZC18].
- A differentially private [Dwo06] value of the size of the intersection, which is computed by adding some noise to the exact count.
- The sum of values associated with the items in the intersection. This is used for measuring ad-generated revenues (cf. Sect.1.2).

The circuits for computing all these functions are of size $O(n)$. Therefore, with our new construction the total size of the circuits for applying these functions to the intersection is $O(n)$.

### 1.4 Related Work

We classify previous works into dedicated protocols for PSI, generic protocols for circuit-based PSI, and dedicated protocols for PSI cardinality.

PSI. The first PSI protocols were based on public-key cryptography, e.g., on the Diffie Hellman function (e.g. [Mea86], with an earlier mention in [Sha80]), oblivious polynomial evaluation [FNP04], or blind RSA [DT10]. More recent protocols are based on oblivious transfer (OT) which can be efficiently instantiated using symmetric key cryptography [IKNP03, ALSZ13]: these protocols use either Bloom filters [FNP04] or hashing to bins [PSZ14,PSSZ15,KKRT16,PSZ18]. All these PSI protocols have super-linear complexity and many of them were compared experimentally in [PSZ18]. PSI protocols have also been evaluated on mobile devices, e.g., in [HCE11,ADN+13, CADT14,KLS+17]. PSI protocols with input sets of different sizes were studied in [KLS+17,PSZ18,RA18].
Circuit-based PSI. These protocols use secure evaluation of circuits for PSI. A trivial circuit for PSI computes $O\left(n^{2}\right)$ comparisons which result in $O\left(\sigma n^{2}\right)$ gates, where $\sigma$ is the bit-length of the elements.

The sort-compare-shuffle (SCS) PSI circuit of [HEK12] computes $O(n \log n)$ comparisons and is of size $O(\sigma n \log n)$ gates (even without the final shuffle layer).

The Circuit-Phasing PSI circuit of [PSSZ15] uses Cuckoo hashing to $O(n)$ bins by one party and simple hashing by the other party which maps at most $O(\log n / \log \log n)$ elements per bin. Therefore, the Circuit-Phasing circuit has a size of $O(\sigma n \log n / \log \log n)$ gates.

The recent circuit-based PSI protocol of [CO18] applies a protocol based on OT extension to compute private set membership in each bin. The outputs of the invocations of this protocol are input to a comparison circuit. The circuit
itself computes a linear number of comparisons, but the total communication complexity of the private set membership protocols is of the same order as that of the Circuit-Phasing circuit [PSSZ15] with $O(\sigma n \log n / \log \log n)$ gates.

Another recent circuit-based PSI protocol of [FNO18, Sect. 8] has communication complexity $O(\sigma n \log \log n)$. It uses hashing to $O(n)$ bins where each bin has multiple buckets and then runs the SCS circuit of [HEK12] to compute the intersection of the elements in the respective bins.

The two-dimensional Cuckoo hashing circuit of [PSWW18] uses a new variant of Cuckoo hashing in two dimensions and has an almost linear complexity of $\omega(\sigma n)$ gates.

In this work, we present the first circuit-based PSI protocol with a true linear complexity of $O(\sigma n)$ gates.

PSI Cardinality. Several protocols for securely computing the cardinality of the intersection, i.e., $|X \cap Y|$, were proposed in the literature. These protocols have linear complexity and are based on public-key cryptography, namely DiffieHellman [DGT12], the Goldwasser-Micali cryptosystem [DD15], or additively homomorphic encryption [DC17]. However, these protocols reveal the cardinality of the intersection to one of the parties. In contrast, circuit-based PSI protocols can easily be adapted to efficiently compute the cardinality and even functions of it using mostly symmetric cryptography.

### 1.5 Our Contributions

In summary, in this paper we present the following contributions:

- The first circuit-based PSI protocol with linear asymptotic communication overhead. We remark that achieving a linear overhead is technically hard since hashing to a table of linear size requires a stash of super-linear size in order to guarantee a negligible failure probability. It is hard to achieve linear overhead with objects of super-linear size.
- A circuit-based PSI protocol with small constants and an improved concrete overhead over the state of the art. As a special case, we consider a very common variant of PSI, namely threshold PSI, in which the intersection is revealed only if it is bigger/smaller than some threshold. Surprisingly, our protocol is $1-2$ orders of magnitude more efficient than the state-of-the-art [ZC18] and has the same asymptotic communication complexity of $O(n)$, despite the fact that the protocol in [ZC18] is a special purpose protocol for threshold-PSI.
- Our protocol supports associating data ("payload") with each input (from both parties), and compute a function that depends on the data associated with the items in the intersection. This property was not supported by the Phasing circuit-based protocol in [PSSZ15]. It is important for applications that compute some function of data associated with the items in the intersection, e.g., aggregate revenues from common users (cf. Sect.1.2).
- On a technical level, we present a new paradigm for handling $\omega(1)$ stash sizes and obtaining an overall overhead that is linear. This is achieved by running an extremely simple dual-execution variant of the protocol.
- Finally, with regards to concrete efficiency, we introduce a circuit-based PSI protocol with linear complexity. This is achieved by using Cuckoo hashing with $K=3$ instead of $K=2$ hash functions, and no stash. This protocol substantially reduces communication (by a factor of 2.6 x to 12.8 x ) and runtime (by a factor of 2.8 x to 5.8 x ) compared to the best previous circuit-based PSI protocol of [PSWW18].


## 2 Preliminaries

### 2.1 Setting

There are two parties, which we denote as $P_{1}$ (the "receiver") and $P_{2}$ (the "sender"). They have input sets, $X$ and $Y$, respectively, each of which contains $n$ items of bitlength $\lambda$. We assume that both parties agree on a function $f$ and wish to securely compute $f(X \cap Y)$. They also agree on a circuit $C$ that receives the items in the intersection as input and computes $f$. That is, $C$ has $O(n \lambda)$ input wires if we consider a computation on the elements themselves or $O(n(\lambda+\rho))$ if we consider a computation on the elements and their associated payloads where the associated payload of each item has bitlength $\rho$. We denote the computational and statistical security parameters by $\kappa$ and $\sigma$, respectively. Denote the set $1, \ldots, c$ by $[c]$. We use the notation $X(i)$ to denote the $i$-th element in the set $X$.

### 2.2 Security Model

This work, similar to most protocols for private set intersection, operates in the semi-honest model, where adversaries may try to learn as much information as possible from a given protocol execution but are not able to deviate from the protocol steps. This is in contrast to malicious adversaries which are able to deviate arbitrarily from the protocol. PSI protocols for the malicious setting exist, but they are less efficient than protocols for the semi-honest setting, e.g., [FNP04,DSMRY09,HN10,DKT10,FHNP16,RR17a,RR17b]. The only circuit-based PSI protocol that can be easily secured against malicious adversaries is the Sort-Compare-Shuffle protocol of [HEK12]: here a circuit of size $O(n)$ can be used to check that the inputs are sorted, resulting in an overall complexity of $O(n \log n)$. For the recent circuit-based PSI protocols that rely on Cuckoo hashing, ensuring that the hashing was done correctly remains the challenge.The semi-honest adversary model is appropriate for scenarios where execution of the intended software is guaranteed via software attestation or business restrictions, and yet an untrusted third party is able to obtain the transcript of the protocol after its execution, by stealing it or by legally enforcing its disclosure.

## FUNCTIONALITY 1 (Two-Party Computation)

Parameters. The Boolean circuit $C$ to be computed, with $I_{1}, I_{2}$ inputs and $O_{1}, O_{2}$ outputs associated with $P_{1}$ and $P_{2}$ resp.
Inputs. $P_{1}$ inputs bits $x_{1}, \ldots, x_{I_{1}}$ and $P_{2}$ inputs bits $y_{1}, \ldots, y_{I_{2}}$.
Outputs. The functionality computes the circuit $C$ on the parties' inputs and returns the outputs to the parties.

### 2.3 Secure Two-Party Computation

There are two main approaches for generic secure two-party computation of Boolean circuits with security against semi-honest adversaries: (1) Yao's garbled circuit protocol [Yao86] has a constant round complexity and with today's most efficient optimizations provides free XOR gates [KS08], whereas securely evaluating an AND gate requires sending two ciphertexts [ZRE15]. (2) The GMW protocol [GMW87] also provides free XOR gates and also sends two ciphertexts per AND gate using OT extension [ALSZ13].

The main advantage of the GMW protocol is that all symmetric cryptographic operations can be pre-computed in a constant number of rounds in a setup phase, whereas the online phase is very efficient, but requires interaction for each layer of AND gates. In more detail, the setup phase is independent of the actual inputs and precomputes multiplication triples for each AND gate using OT extension in a constant number of rounds. The online phase begins when the inputs are provided and involves a communication round for each layer of AND gates. See [SZ13] for a detailed description and comparison between Yao and GMW.

In our protocol we make use of Functionality 1.

### 2.4 Cuckoo Hashing

Cuckoo hashing [PR01] uses two hash functions $h_{0}, h_{1}$ to map $n$ elements to two tables $T_{0}, T_{1}$ which each contain $(1+\varepsilon) n$ bins. (It is also possible to use a single table $T$ with $2(1+\varepsilon) n$ bins. The two versions are essentially equivalent). Each bin accommodates at most a single element. The scheme avoids collisions by relocating elements when a collision is found using the following procedure: Let $b \in\{0,1\}$. An element $x$ is inserted into a bin $h_{b}(x)$ in table $T_{b}$. If a prior item $y$ exists in that bin, it is evicted to $\operatorname{bin} h_{1-b}(y)$ in $T_{1-b}$. The pointer $b$ is then assigned the value $1-b$. The procedure is repeated until no more evictions are necessary, or until a threshold number of relocations has been reached. In the latter case, the last element is put in a special stash. It was shown in [KMW09] that for a stash of constant size $s$ the probability that the stash overflows is at most $O\left(n^{-(s+1)}\right)$. It was also shown in [GM11] that this failure probablity is negligilble when the stash is of $\operatorname{size} O(\log n)$. An observation in [KM18] shows that this is also the case when $s=O\left(\omega(1) \cdot \frac{\log n}{\log \log n}\right)$. After insertion, each item can be found in one of two locations or in the stash.

### 2.5 PSI Based on Hashing

Existing constructions for circuit-based PSI require the parties to reorder their inputs before inputting them to the circuit. In the sorting network-based circuit of [HEK12], the parties sort their inputs. In the hashing-based construction of [PSSZ15], the parties map their items to bins using a hashing scheme.

It was observed as early as [FNP04] that if the two parties agree on the same hash function and use it to assign their respective input to bins, then the items that one party maps to a specific bin only need to be compared to the items that the other party maps to the same bin. However, the parties must be careful not to reveal to each other the number of items that they mapped to each bin, since this leaks information about their input sets. Therefore, the parties agree beforehand on an upper bound $m$ for the maximum number of items that can be mapped to a bin (such upper bounds are well known for common hashing algorithms, and can also be substantiated using simulation), and pad each bin with random dummy values until it has exactly $m$ items in it. If both parties use the same hash algorithm, then this approach considerably reduces the overhead of the computation from $O\left(n^{2}\right)$ to $O\left(\beta \cdot m^{2}\right)$ where $\beta$ is the number of bins.

When using a random hash function $h$ to map $n$ items to $n$ bins such that $x$ is mapped to bin $h(x)$, the most occupied bin has at most $m=\frac{\ln n}{\ln \ln n}(1+o(1))$ items with high probability [Gon81]. For instance, for $n=2^{20}$ and a desired error probability of $2^{-40}$, a careful analysis shows that $m=20$. Cuckoo hashing is much more promising, since it maps $n$ items to $2 n(1+\varepsilon)$ bins, where each bin stores at most $m=1$ items.

It is tempting to let both parties, $P_{1}$ and $P_{2}$, map their items to bins using Cuckoo hashing, and then only compare the item that $P_{1}$ maps to a bin with the item that $P_{2}$ maps to the same bin. The problem is that $P_{1}$ might map $x$ to $h_{0}(x)$ whereas $P_{2}$ might map it to $h_{1}(x)$. Unfortunately, they cannot use a protocol where $P_{1}$ 's value in bin $h_{0}(x)$ is compared to the two bins $h_{0}(x), h_{1}(x)$ in $P_{2}$ 's input, since this reveals that $P_{1}$ has an item which is mapped to these two locations. The solution used in [FHNP16, PSZ14, PSSZ15] is to let $P_{1}$ map its items to bins using Cuckoo hashing, and $P_{2}$ map its items using simple hashing. Namely, each item of $P_{2}$ is mapped to both bins $h_{0}(x), h_{1}(x)$. Therefore, $P_{2}$ needs to pad its bins to have exactly $m=O(\log n / \log \log n)$ items in each bin, and the total number of comparisons is $O(n \log n / \log \log n)$.

## 3 OPPRF - Oblivious Programmable PRF

Our protocol builds on a (batched) oblivious programmable pseudorandom function (OPPRF). In this section we gradually present the properties required by that kind of a primitive, by first describing simpler primitives, namely, Programmable PRF (and its batched version) and Oblivious PRF.

### 3.1 Oblivious PRF

An oblivious PRF (OPRF) [FIPR05] is a two-party protocol implementing a functionality between a sender and a receiver. Let $F$ be a pseudo-random
function (PRF) such that $F:\{0,1\}^{\kappa} \times\{0,1\}^{\ell} \rightarrow\{0,1\}^{\ell}$. The sender inputs a key $k$ to $F$ and the receiver inputs $q_{1} \ldots, q_{c}$. The functionality outputs $F\left(k, q_{1}\right), \ldots, F\left(k, q_{c}\right)$ to the receiver and nothing to the sender. In another variant of oblivious PRF the sender is given a fresh random key $k$ as an output from the functionality rather than choosing it on its own. In our protocol we will make use of a "one-time" OPRF functionality in which the receiver can query $a$ single query, namely, the sender inputs nothing and the receiver inputs a query $q$; the functionality outputs to the sender a key $k$ and to the receiver the result $F_{k}(q)$. Let us denote that functionality by $\mathcal{F}_{\text {OPRF }}$.

## 3.2 (One-Time) Programmable PRF (PPRF)

A programmable PRF (PPRF) is similar to a PRF, with the additional property that on a certain "programmed" set of inputs the function outputs "programmed" values. Namely, for an arbitrary set $X$ and a "target" multi-set $T$, where $|X|=|T|$ and each $t \in T$ is uniformly distributed ${ }^{3}$, it is guaranteed that on input $X(i)$ the function outputs $T(i)$. Let $\mathcal{T}$ be a distribution of such multi-sets, which may be public to both parties.

The restriction of the PPRF to be only one-time comes from the fact that we allow the elements in $T$ to be correlated. If the elements are indeed correlated then by querying it two times (on the correlated positions) it would be easy to distinguish it from a random function.

We capture the above notion by the following formal definition:
Definition 1. An $\ell$-bits PPRF is a pair of algorithms $\hat{F}=($ Hint,$F)$ as follows:

- $\operatorname{Hint}(k, X, T) \rightarrow$ hint $_{k, X, T}$ : Given a uniformly random key $k \in\{0,1\}^{\kappa}$, the set $X$ where $|X(i)|=\ell$ for all $i \in[|X|]$ and a target multi-set $T$ with $|T|=|X|$ and all elements in $T$ are uniformly distributed (but may be correlated), output the hint hint $t_{k, X, T} \in\{0,1\}^{\kappa \cdot|X|}$.
$-F(k$, hint,$x) \rightarrow y^{\star}$. Given a key $k \in\{0,1\}^{\kappa}$, a hint hint $\in\{0,1\}^{\kappa \cdot|X|}$ and an input $x \in\{0,1\}^{\ell}$, output $y^{\star} \in\{0,1\}^{\ell}$.

We consider two properties of a PPRF, correctness and security:

- Correctness. For every $k, T$ and $X$, and for every $i \in[|X|]$ we have:

$$
F(k, \text { hint, } X(i))=T(i) .
$$

- Security. We say that an interactive machine $M$ is a PPRF oracle over $\hat{F}$ if, when interacting with a "caller" $\mathcal{A}$, it works as follows:

1. $M$ is given a set $X$ from $\mathcal{A}$.
2. $M$ samples a uniformly random $k \in\{0,1\}^{\kappa}$ and $T$ from $\mathcal{T}$, invokes hint $\leftarrow \operatorname{Hint}(k, X, T)$ and hands hint to $\mathcal{A}$.
[^1]
## CONSTRUCTION 2 (PPRF)

Let $F^{\prime}:\{0,1\}^{\kappa} \times\{0,1\}^{\ell} \rightarrow\{0,1\}^{\ell}$ be a PRF.

- $\operatorname{Hint}(k, X, T)$. Interpolate a polynomial $p$ over the points $\left\{\left(X(i), F_{k}^{\prime}(X(i)) \oplus T(i)\right)\right\}_{i \in[|X|]}$. Return $p$ as the hint.
- $F(k$, hint,$x)$. Interpret hint as a polynomial, denoted $p$. Return $F_{k}^{\prime}(x) \oplus$ $p(x)$.

3. $M$ is given an input $x \in\{0,1\}^{\ell}$ from $\mathcal{A}$ and responds with $F(k$, hint, $x)$.
4. $M$ halts (any subsequent queries will be ignored).

The scheme $\hat{F}$ is said to be secure if, for every $X$ input by $\mathcal{A}$ (i.e. the caller), the interaction of $\mathcal{A}$ with $M$ is computationally indistinguishable from the interaction with the PPRF oracle $\mathcal{S}$, where $\mathcal{S}$ outputs a uniformly random "hint" $\{0,1\}^{\kappa \cdot|X|}$ and a "PRF result" from $\{0,1\}^{\ell}$.

The definition is reminiscent of a semantically secure encryption scheme. Informally, semantic security means that whatever is efficiently computable about the cleartext given the ciphertext, is also efficiently computable without the ciphertext. Also here, whatever can be efficiently computable given $X$ is also efficiently computable given only $|X|$. That implicitely means that the interaction with a PPRF oracle $M$ over (KeyGen, $F$ ) does not leak the elements in $X$.

Our security definition diverges from that of $[K M P+17]$ in two aspects:

1. In $[\mathrm{KMP}+17], \mathcal{A}$ has many queries to $M$ in Step 3 of the interaction, whereas our definition allows only a single query. In the ( $n, t$ )-security definition in $[\mathrm{KMP}+17]$ this corresponds to setting $t=1$. Our definition is weaker in this sense, but this is sufficient for our protocol as we invoke multiple instances of the one-time PPRF.
2. The definition in $[\mathrm{KMP}+17]$ compensates for the fact that $\mathcal{A}$ has many queries, by requiring that the function $F$ outputs an independent target value for every $x \in X$. Our definition is stronger as it allows having correlated target elements in $T$. In the most extreme form of correlation all values in $T$ are equal, which makes the task of the adversary "easier". We require the security property to hold even in this case.

We present in Construction 2 a polynomial-based PPRF scheme that is based on the construction in $[\mathrm{KMP}+17]$.

Theorem 3. Construction 2 is a PPRF.
Proof. It is easy to see that this construction is correct. For every $k, X$ and $T$, let $p=\operatorname{Hint}(k, X, T)$, then for all $i \in|X|$ it holds that

$$
\begin{aligned}
F(k, p, X(i)) & =F^{\prime}(k, X(i)) \oplus p(X(i)) \\
& =F^{\prime}(k, X(i)) \oplus F^{\prime}(k, X(i)) \oplus T(i) \\
& =T(i)
\end{aligned}
$$

as required. We now reduce the security of the scheme to the security of a PRF (i.e., to the standard PRF definition, with many oracle accesses). Let $M$ be a PPRF oracle over $\hat{F}$ of Construction 2. Assume there exists a distinguisher $\mathcal{D}$ and a caller $\mathcal{A}$ such that $\mathcal{D}$ distinguishes between the output of $M$ after interacting with $\mathcal{A}$, when $\mathcal{A}$ chooses $X$ and $x$ as its inputs, and the output of $\mathcal{S}\left(1^{\kappa},|X|\right)$ (where $\mathcal{S}$ is the simulator described in Definition 1) with probability $\mu$.

We present a distinguisher $\mathcal{D}_{\tilde{\sim}}^{\prime}$ that has an oracle access to either a truly random function $R(\cdot)$ or a $\operatorname{PRF} \tilde{F}(k, \cdot)$. The distinguisher $\mathcal{D}^{\prime}$ runs as follows:

Given an oracle $\mathcal{O}$ to either $R(\cdot)$ or $\tilde{F}(k, \cdot), \mathcal{D}^{\prime}$ samples $T$ from $\mathcal{T}$, then, for every $i \in[|X|]$ it queries the oracle on $X(i)$ and obtains $\mathcal{O}(X(i))$. It interpolates the polynomial $p$ using the points $\{(X(i), \mathcal{O}(X(i)) \oplus T(i))\}_{i \in|X|}$ and provides $p$ 's coefficients to $\mathcal{D}$. For the query $x, \mathcal{D}^{\prime}$ hands $\mathcal{D}$ the value $\mathcal{O}(x) \oplus p(x)$ and outputs whatever $\mathcal{D}$ outputs.

Observe that if $\mathcal{O}$ is truly random, then the values $\{R(X(i)) \oplus T(i))\}_{i \in[|X|]}$ are uniformly random and thus the polynomial $p$ is uniformly random and independent of $T$. If $x \notin X$ then the value $R(x) \oplus p(x)$ is obviously random since $R(x)$ is independent of $p$. In addition, if $x=X(i)$ for some $i$, then the value $R(x) \oplus p(x)$ equals $T(i)$ for some $i \in[|X|]$, which is uniformly random since $T$ is sampled from $\mathcal{T}$ and every $t \in T$ is distributed uniformly. Therefore, the pair $(p, R(x) \oplus p(x))$ is distributed identically to the output of $\mathcal{S}$. On the other hand, if $\mathcal{O}$ is a pseudorandom function, then the values $\left.\left\{F_{k}(X(i)) \oplus T(i)\right)\right\}_{i \in[|X|]}$ from which the polynomial $p$ is interpolated, along with the second output $F_{k}(x) \oplus p(x)$, are distributed identically to the output of $M$ upon an interaction with $\mathcal{A}$. This leads to the same distinguishing success probability $\mu$, for both $\mathcal{D}$ and $\mathcal{D}^{\prime}$, which must be negligible.

### 3.3 Batch PPRF

Note that the size of the hint generated by algorithm KeyGen is $\kappa \cdot|X|$ (i.e., the polynomial is represented by $|X|$ coefficients, each of size $\kappa$ bits). In our setting we use an independent PPRF per bin, where each bin contains at most $O(\log n / \log \log n)$ values. Therefore the hint for one bin is of size $O(\kappa \cdot \log n / \log \log n)$, and the size of all hints is $O(\kappa \cdot n \cdot \log n / \log \log n)$. However, we know that the total number of values in all $P_{2}$ 's bins is $2 n$, since each value is stored in (at most) two locations of the table ${ }^{4}$. We next show that it is possible to combine the hints of all bins to a single hint of length $2 n$, thus reducing the total communication for all hints to $O(n)$.

We first present a formal definition of the notion of batch PPRF.
Definition 2. An $\ell$-bits, $\beta$-bins PPRF (or $(\ell, \beta)$-PPRF) is a pair of algorithms $\hat{F}=($ KeyGen,$F)$ as follows:

[^2]
## CONSTRUCTION 4 (Batched PPRF)

Let $F^{\prime}:\{0,1\}^{\kappa} \times\{0,1\}^{\ell} \rightarrow\{0,1\}^{\ell}$ be a PRF.
$-\operatorname{Hint}(k, X, T)$.
Given the keys $k=k_{1}, \ldots, k_{\beta}$, the sets $X=X_{1}, \ldots, X_{\beta}$ and the target multi-sets $T=T_{1}, \ldots, T_{\beta}$, interpolate the polynomial $p$ using the points $\left\{\left(X_{j}(i), F^{\prime}\left(k_{j}, X_{j}(i)\right) \oplus T_{j}(i)\right)\right\}_{j \in \beta ; i \in\left[\left|X_{j}\right|\right]}$. Return $p$ as the hint.

- $F(k$, hint, $x)$.

Interpret hint as a polynomial, denoted $p$. Return $F^{\prime}(k, x) \oplus p(x)$. (Same as in Construction 2.)

- $\operatorname{Hint}(k, X, T) \rightarrow \operatorname{hint}_{k, X, T}$. Given a set of uniformly random and independent keys $k=k_{1}, \ldots, k_{\beta} \in\{0,1\}^{\kappa}$, the sets $X=X_{1}, \ldots, X_{\beta}$ where $\left|X_{j}(i)\right|=\ell$ for all $j \in[\beta]$ and $i \in[|X|]$ and a target multi-sets $T=T_{1}, \ldots, T_{\beta}$ where for every $j \in[\beta]$ it holds that $\left|T_{j}\right|=\left|X_{j}\right|$ and all elements in $T_{j}$ are uniformly distributed (but, again, may be correlated), output the hint hint ${ }_{k, X, T} \in\{0,1\}^{\kappa \cdot N}$ where $N=\sum_{j=1}^{\beta}\left|X_{j}\right|$.
- $F(k$, hint, $x) \rightarrow y^{\star}$. Given a key $k \in\{0,1\}^{\kappa}$, a hint hint $\in\{0,1\}^{\kappa \cdot N}$ and an input $x \in\{0,1\}^{\ell}$, output $y^{\star} \in\{0,1\}^{\ell}$.

As before, we want a batched PPRF to have the following properties:

- Correctness. For every $k=k_{1}, \ldots, k_{\beta}, T=T_{1}, \ldots, T_{\beta}$ and $X=X_{1}, \ldots, X_{\beta}$ as above, we have

$$
F\left(k_{j}, \text { hint, } X_{j}(i)\right)=T_{j}(i)
$$

for every $j \in[\beta]$ and $i \in\left[\left|X_{j}\right|\right]$.

- Security. We say that an interactive machine $M$ is a batched PPRF oracle over $\hat{F}$ if, when interacting with a "caller" $\mathcal{A}$, it works as follows:

1. $M$ is given $X=X_{1}, \ldots, X_{\beta}$ from $\mathcal{A}$.
2. $M$ samples uniformly random keys $k=k_{1}, \ldots, k_{\beta}$ and target multi-sets $T=T_{1}, \ldots, T_{\beta}$ from $\mathcal{T}$, and invokes hint $\leftarrow \operatorname{Hint}(k, X, T)$ hands hint to $\mathcal{A}$.
3. $M$ is given $\beta$ queries $x_{1}, \ldots, x_{\beta}$ from $\mathcal{A}$ and responds with $y_{1}^{\star}, \ldots, y_{\beta}^{\star}$ where $y_{j}^{\star}=F\left(k_{j}\right.$, hint, $\left.x_{j}\right)$.
4. $M$ halts.

The scheme $\hat{F}$ is said to be secure if for every disjoint sets $X_{1}, \ldots, X_{\beta}$ (where $\left.N=\sum_{j \in[\beta]}\left|X_{j}\right|\right)$ input by a PPT machine $\mathcal{A}$, the output of $M$ is computationally indistinguishable from the output of $\mathcal{S}\left(1^{\kappa}, N\right)$, such that $\mathcal{S}$ outputs a uniformly random hint $\in\{0,1\}^{\kappa \cdot N}$ and a set of $\beta$ uniformly random values from $\{0,1\}^{\ell}$.

Construction 4 is a batched version of Construction 2.
Theorem 5. Construction 4 is a secure $(\ell, \beta)-P P R F$.

Proof. For correctness, note that for every $j \in[\beta]$ and $i \in\left[\left|X_{j}\right|\right]$ it holds that

$$
\begin{aligned}
F\left(k_{j}, p, X_{j}(i)\right) & =F^{\prime}\left(k_{j}, X_{j}(i)\right) \oplus p\left(X_{j}(i)\right) \\
& =F^{\prime}\left(k_{i}, X_{j}(i)\right) \oplus F^{\prime}\left(k_{i}, X_{j}(i)\right) \oplus T_{j}(i) \\
& =T_{j}(i)
\end{aligned}
$$

The security of the scheme is reduced to the security of a batch PRF $\tilde{F}$. Informally, a batch PRF works as follows: Sample uniform keys $k_{1}, \ldots, k_{\beta} \in$ $\{0,1\}^{\kappa}$ and for a query $(j, x)$ respond with $\tilde{F}\left(k_{j}, x\right)$. One can easily show that a batch PRF is indistinguishable from a set of $\beta$ truly random functions $R_{1}, \ldots, R_{\beta}$ where on query $(j, x)$ the output is $R_{j}(x)$.

Let $M$ be a batched PPRF oracle over $\hat{F}$ of Construction 4. Assume there exists a distinguisher $\mathcal{D}$ and a caller $\mathcal{A}$ such that $\mathcal{D}$ distinguishes between the output of $M$ after interacting with $\mathcal{A}$, when $\mathcal{A}$ chooses $X_{1}, \ldots, X_{\beta}$ and $x_{1}, \ldots, x_{\beta}$ as its inputs, and the outputs of $\mathcal{S}\left(1^{\kappa}, N\right)$, where $\mathcal{S}$ is the simulator described in Definition 2.

We present a distinguisher $\mathcal{D}^{\prime}$ that has an oracle access $\mathcal{O}$, to either a batch $\operatorname{PRF} \tilde{F}\left(k_{j}, \cdot\right)$ or a set of truly random functions $R_{j}(\cdot)$ (where $j \in[\beta]$ ). The distinguisher $\mathcal{D}^{\prime}$ works as follows: sample $T_{1}, \ldots, T_{\beta}$ from $\mathcal{T}$, interpolate a polynomial $p$ with the points $\left\{\left(X_{j}(i), \mathcal{O}\left(j, X_{j}(i)\right) \oplus T_{j}(i)\right)\right\}_{j \in[\beta] ; i \in\left[\left|X_{j}\right|\right]}$ and hand $p$ 's coefficients to $\mathcal{D}$ as the hint. Then, for query $x_{j}$ of $\mathcal{D}$, respond with $y_{j}^{\star}=\mathcal{O}\left(x_{j}\right) \oplus p\left(x_{j}\right)$. Finally, $\mathcal{D}^{\prime}$ outputs whatever $\mathcal{D}$ outputs.

First note that if $\mathcal{O}$ is a set of truly random functions then the polynomial $p$ is uniformly random and independent of $y_{1}^{\star}, \ldots, y_{\beta}^{\star}$ because all interpolation points are uniformly random. Now, if $x_{j} \notin X_{j}$ then the result is obviously uniformly random. Otherwise, if $x_{j}=X_{j}(i)$ for some $i$ then note that the result is $T_{j}(i)$ which is uniformly random as well, since the other elements in $T_{j}$ are unknown. Thus, this is distributed identically to the output of $\mathcal{S}\left(1^{\kappa}, N\right)$. On the other hand, if $\mathcal{O}$ is a batch PRF then the interpolation points $\left\{\left(X_{j}(i), \mathcal{O}\left(j, X_{j}(i)\right) \oplus\right.\right.$ $\left.\left.T_{j}(i)\right)\right\}_{j \in[\beta] ; i \in\left[\left|X_{j}\right|\right]}$ along with $y_{1}^{\star}, \ldots, y_{\beta}^{\star}$ are distributed identically to the output of $M$ upon an interaction with $\mathcal{A}$. This leads to the same distinguishing success probability for both $\mathcal{D}$ and $\mathcal{D}^{\prime}$, which must be negligible.

### 3.4 Batch Oblivious Programmable Pseudorandom Functions

In this section we define a two-party functionality for batched oblivious programmable pseudorandom function (Functionality 6), which is the main building block in our PSI protocols. The functionality is parametrised by a $(\ell, \beta)$-PPRF $\hat{F}=($ Hint,$F)$ and interacts with a sender, who programs $\hat{F}$ with $\beta$ sets, and a receiver who queries $\hat{F}$ with $\beta$ queries. The functionality guarantees that the sender does not learn what are the receiver's queries and the receiver does not learn what are the programmed points.

Given a protocol that realizes $\mathcal{F}_{\text {OPRF }}$ and a secure $(\ell, \beta)$-PPRF, the realization of Functionality 6 is simple and described in Protocol 7.

Theorem 8. Given an $(\ell, \beta)-P P R F$, Protocol 7 securely realizes Functionality 6 in the $\mathcal{F}_{\text {OPRF-hybrid }}$ model.

## FUNCTIONALITY 6 (Batch Oblivious PPRF)

Parameters. A $(\ell, \beta)$-PPRF $\hat{F}=($ Hint,$F)$.
Sender's inputs. These are the following values:

- Disjoint sets $X=X_{1}, \ldots, X_{\beta}$ where $\left|X_{j}(i)\right| \in\{0,1\}^{\ell}$ for every $j \in[\beta]$ and $i \in\left[\left|X_{j}\right|\right]$. Let the total number of elements in all sets be $N=\sum_{j}\left|X_{j}\right|$.
- The sets $T=T_{1}, \ldots, T_{\beta}$ sampled independently from $\mathcal{T}$.

Receiver's inputs. The queries $x_{1}, \ldots, x_{\beta} \in\{0,1\}^{\ell}$.
The functionality works as follows:

1. Sample uniformly random and independent keys $k=k_{1}, \ldots, k_{\beta}$.
2. Invoke $\operatorname{Hint}(k, X, T) \rightarrow$ hint.
3. Output hint to $P_{1}\left(P_{2}\right.$ can compute it on its own from $\left.k, X, T\right)$.
4. For every $j \in[\beta]$ output $F\left(k_{j}\right.$, hint, $\left.x_{j}\right)$ to the receiver.

## PROTOCOL 7 (Batch Oblivious PPRF)

The protocol is defined in the $\mathcal{F}_{\text {OPRF-hybrid model and receives an }(\ell, \beta) \text {-PPRF }}$ $\hat{F}=($ Hint,$F)$ as a parameter. The underlying PRF in both $\mathcal{F}_{\text {OPRF }}$ and $\hat{F}$ is the same and denoted $F^{\prime}$. The protocol proceeds as follows:

1. The parties invoke $\beta$ instances of $\mathcal{F}_{\text {OPRF }}$. In the $j \in[\beta]$ instance, $P_{2}$ inputs nothing and receives the key $k_{j}$, and $P_{1}$ inputs $x_{j}$ and receives $F^{\prime}\left(k_{j}, x_{j}\right)$.
2. Party $P_{2}$ invokes $p \leftarrow \operatorname{Hint}(k, X, T)$ and sends $p$ to $P_{1}$.
3. For every $j \in[\beta]$, party $P_{1}$ outputs $F^{\prime}\left(k_{j}, x_{j}\right) \oplus p\left(x_{j}\right)$.

Proof. Note that party $P_{2}$ receives nothing in the functionality but receives $k_{1}, \ldots, k_{\beta}$ in the real execution as output from $\mathcal{F}_{\text {OPRF }}$. Therefore, $P_{2}$ 's view can be easily simulated with the simulator of $\mathcal{F}_{\text {OPRF }}$.

As for the view of $P_{1}$, from the security of the PPRF it follows that it is indistinguishable from the output of $\mathcal{S}\left(1^{\kappa}, N\right)$ where $\mathcal{S}$ is the simulator from Definition 2.

## 4 A Super-Linear Communication Protocol

### 4.1 The Basic Construction

Let $C_{a, b}$ be a Boolean circuit that has $2 \cdot a \cdot(b+\lambda)$ input wires, divided to $a$ sections of $2 b+\lambda$ inputs wires each. For each section, the first (resp. second) $\beta$ input wires are associated with $P_{1}$ (resp. $P_{2}$ ). The last $\lambda$ input wires are associated with $P_{1}$ as well. Denote the first (resp. second) $\beta$ bits input to the $i$ th section by $u_{i, 1}$ (resp. $v_{i, 2}$ ) and the last $\lambda$ bits by $z_{i}$. The circuit first compares $u_{i, 1}$ to $v_{i, 2}$ for every $i \in[\alpha]$ and produces $w_{i}=1$ if $u_{i, 1}=v_{i, 2}$ and 0 otherwise.

## PROTOCOL 9 (Private Set Intersection)

Inputs. $P_{1}$ has $X=\left\{x_{1}, \ldots, x_{n}\right\}$ and $P_{2}$ has $Y=\left\{y_{1}, \ldots, y_{n}\right\}$.
Protocol. The protocol proceeds in 3 steps as follows:

1. Hashing. The parties agree on hash functions $H_{1}, H_{2}:\{0,1\}^{\ell} \rightarrow[\beta]$, which are used as follows:

- $P_{1}$ uses $H_{1}, H_{2}$ in a Cuckoo hashing construction that maps $x_{1}, \ldots, x_{n}$ to a table Table ${ }_{1}$ of $\beta=2(1+\varepsilon) n$ entries, where input $x_{i}$ is mapped to either entry Table $\left[H_{1}\left(x_{i}\right)\right]$ or Table ${ }_{1}\left[H_{2}\left(x_{i}\right)\right]$ or the stash Stash (which is of size $s)^{a}$. Since $\beta>n, P_{1}$ fills the empty entries in Table ${ }_{1}$ with a uniformly random value.
- $P_{2}$ maps $y_{1}, \ldots, y_{n}$ to Table ${ }_{2}$ of $\beta$ entries using both $H_{1}$ and $H_{2}$. That is, $y_{i}$ is placed in both $\operatorname{Table}_{2}\left[H_{1}\left(y_{i}\right)\right]$ and $\operatorname{Table}_{2}\left[H_{2}\left(y_{i}\right)\right]$. (Obviously, some bins will have multiple items mapped to them. This is not an issue, and there is even no need to use a probabilistic upper bound on the occupancy of the bin.)

2. Computing batch OPPRF. $P_{2}$ samples uniformly random and independent target values $t_{1}, \ldots, t_{\beta} \in\{0,1\}^{\kappa}$. The parties invoke an $(\lambda, \beta)$ OPPRF (Functionality 6; recall that $\lambda$ is the bit-length of the items). $P_{2}$ inputs $Y_{1}, \ldots, Y_{\beta}$ and $T_{1}, \ldots, T_{\beta}$ where $Y_{j}=$ Table $_{2}[j]=\{y \| j \mid y \in$ $\left.Y \wedge j \in\left\{H_{1}(y), H_{2}(y)\right\}\right\}$ and $T_{j}$ has $\left|Y_{j}\right|$ elements, all equal to $t_{j}$. If, $j=H_{1}(y)=H_{2}(y)$ for some $y \in Y$ then $P_{2}$ adds a uniformly random element to Table $_{2}[j] . P_{1}$ inputs Table ${ }_{1}[1], \ldots$, Table $_{1}[\beta]$ and receives $y_{1}^{\star}, \ldots, y_{\beta}^{\star}$. According to the definition of the OPPRF, if Table ${ }_{1}[j] \in$ Table $_{2}[j]$ then $y_{j}^{\star}=t_{j}$.
3. Computing the circuit. The parties use a two-party computation (Functionality 1) with the circuit $C_{\beta+s \cdot n, \gamma} \cdot{ }^{b}$ For section $j \in[\beta]$ of the circuit, party $P_{1}$ inputs the first $\gamma$ bits of $y_{j}^{\star}$ and Table $_{1}[j]$, and $P_{2}$ inputs the first $\gamma$ bits $t_{j}$; for the $\beta+j$-th section $P_{1}$ inputs $\operatorname{Stash}[\lceil j / n\rceil+1]$ and $P_{2}$ inputs Table $[(j \bmod n)+1]$.
${ }^{a}$ We discuss the value of $s$ in Sect. 4.2 and the value of $\varepsilon$ in Sect.7.1.
${ }^{b}$ We discuss the value of $\gamma$ in Sect.4.2.

Then, the circuit computes and outputs $f(Z)$ where $Z=\left\{z_{i} \mid w_{i}=1\right\}_{i \in[a]}$ and $f$ is the function required to be computed in the $\mathcal{F}_{\mathrm{PSI}, f}$ functionality.

Correctness. If $z \in X \cap Y$ then $z$ is mapped to both $\operatorname{Table}_{2}\left[H_{1}(z)\right]$ and Table $_{2}\left[H_{2}(z)\right]$ by $P_{2}$. There are two cases: (1) $z$ is mapped to $\operatorname{Table}_{1}\left[H_{b}(z)\right]$ by $P_{1}$ for $b \in\{1,2\}$. (2) $z$ is mapped to Stash by $P_{1}$. In the first case the match is found in section $H_{b}(z)$ of the circuit; in the second case the match is certainly found since every item in the Stash is compared to every item in $Y$.

Two items $x \in X$ and $y \in Y$ where $x \neq y$ will not be matched, since by the properties of the PPRF $P_{1}$ receives a pseudorandom output. Since the parties only input the first $\gamma$ bits of the PPRF results, those values will be matched with probability $2^{-\gamma}$. See Sect. 4.2 for a discussion on limiting the failure probability.

Security. The security of the protocol follows immediately from the security of the OPPRF and the two-party computation functionalities.

### 4.2 Limiting the Failure Probability

Protocol 9 might fail due to two reasons:

- Stash size. For an actual implementation, one needs to fix $s$ and $\varepsilon$ so that the stash failure probability will be smaller than $2^{-\sigma}$. If the stash is overflowed (i.e., more than $s$ items are mapped to it) then the protocol fails. ${ }^{5}$ As discussed in Sect. 2, setting $s=O(\log n / \log \log n)$ makes the failure probability negligible.
- Input encoding. The circuit compares the first $\gamma$ bits of $y_{j}^{\star}$ of $P_{1}$ to the first $\gamma$ bits of $t_{j}$ of $P_{2}$. Thus, the false positive error probability in each comparison equals $2^{-\gamma}$ (due to $F(x)$, for $x \notin Y$, being equal to the programmed output), and therefore the overall probability of a false positive is at most $\beta \cdot 2^{-\gamma}=$ $2(1+\varepsilon) n \cdot 2^{-\gamma}$.


### 4.3 Reducing Computation

A major computation task of the protocol is interpolating the polynomial which encodes the hint. If we use Cuckoo hashing with $K=O(1)$ hash functions then the polynomial encodes $O(n)$ items and is of degree $O(n)$. This section describes how to reduce the asymptotic overhead of computing the polynomial and therefore we will use asymptotic notation. The concrete overhead is discussed in Sect. 7.2.

The overhead of interpolating a polynomial of degree $O(n)$ over arbitrary points is $O\left(n^{2}\right)$ operations using Lagrange interpolation, or $O\left(n \log ^{2} n\right)$ operations using FFT. The overhead can be reduced by dividing the polynomial to several lower-degree polynomials. In particular, let us divide the $\beta=O(n)$ bins to $B$ "mega-bins", each encompassing $\beta / B$ bins. Suppose that we have an upper bound such that the number of items in a mega-bin is at most $m$, except with negligible probability. Then the protocol can invoke a batch OPPRF for each mega-bin, using a different hint polynomial. Each such polynomial is of degree $m$. Therefore the computation overhead is $O\left(B \cdot m \log ^{2} m\right)$. Ideally, the upper bound on the number of items in a mega-bin, $m$, is of the same order as the expected number of items in a mega-bin, $O(n / B)$. In this case the computation overhead is $O\left(n / B \cdot B \cdot \log ^{2}(n / B)\right)=O\left(n \log ^{2}(n / B)\right)$ and will be minimized when the number of mega-bins $B$ is maximal.

It is known that when mapping $O(n)$ items to $B=n / \log n$ (mega-)bins, then with high probability the most occupied bin has less than $m=O(n / B)=$ $O(\log n)$ items. When interested in concrete efficiency we can use the analysis

[^3]in [PSZ18] to find the exact number of mega-bins to make the failure probability sufficiently small (see Sect.7.2). When interested in asymptotic analysis, it is easy to deduce from the analysis in [PSZ18] that with $B=n / \log n$ mega-bins, the probability of having more than $\omega(\log n)$ items in a megabin is negligible. Therefore when using this number of mega-bins, the computation overhead is only $\omega\left((n / \log n) \cdot \log ^{2}(n)\right)=\omega(n \log n)$ using Lagrange interpolation. Using FFT interpolation, the asymptotic overhead is reduced to $\omega\left((n / \log n) \log n(\log \log n)^{2}\right)=\omega\left(n \cdot(\log \log n)^{2}\right)$. But since we map relatively few items to each mega-bin the gain in practice of using FFT is marginal.

## 5 A Linear Communication Protocol

We describe here a protocol in which the circuit computes only $O(n)$ comparisons. This protocol outperforms the protocols in Sect. 4.1 or in [PSWW18, CO18] which have a circuit that computes $\omega(n)$ comparisons. A careful analysis reveals that those protocols require $O(n)$ comparisons to process all items that were mapped to the Cuckoo hash table, and an additional $s \cdot n$ comparisons to process the $s=\omega(1)$ items that were mapped to the stash. We note that the concurrent and independent work of [FNO18] proposes to use a PSI protocol for unbalanced set sizes, such as in the work of [KLS +17$]$, to reduce the complexity of handling the stash from $\omega(n)$ to $O(n)$ in PSI protocols. However, their idea can only be applied when the output is the intersection itself. When the output is a function of the intersection then their protocol has communication complexity $O(n \log \log n)$, cf. Sect.1.4). In contrast, we achieve $O(n)$ communication even when the output is a function of the intersection.

We present two different techniques to achieve a linear communication protocol with failure probability that is negligible in the statistical security parameter $\sigma$. The first technique (see Sect.5.1) is implied by a mathematical analysis of the failure probability (as argued in Sect.1.4). The second technique (see Sect.5.2) is implied by the empirical analysis presented in [PSZ18].

### 5.1 Linear Communication via Dual Execution

We overcome the difficulty of handling the stash by running a modified version of the protocol in three phases. The first phase is similar to the basic protocol, but ignores the items that $P_{1}$ maps to the stash. Therefore this phase inputs to the circuit the $O(n)$ results of comparing $P_{1}$ 's input items (except those mapped to the stash) with all of $P_{2}$ 's items. The second phase reverses the roles of the parties, and in addition now $P_{1}$ inputs only the items that it previously mapped to the stash. In this phase $P_{2}$ uses Cuckoo hashing and might map some items to the stash. The last phase only compares the items that $P_{1}$ mapped to the stash in the first phase, to the items that $P_{2}$ mapped to the stash in the second phase, and therefore only needs to handle very few items. Below, we describe our protocol in more detail: In Protocol 10, we describe our protocol in more detail.

## PROTOCOL 10 (PSI with Linear Communication)

Inputs. $P_{1}$ has $X=\left\{x_{1}, \ldots, x_{n}\right\}$ and $P_{2}$ has $Y=\left\{y_{1}, \ldots, y_{n}\right\}$.
Protocol. The protocol proceeds in 3 phases as follows:

1. Run steps 1-2 of Protocol 9. Denote the items mapped to $P_{1}$ 's table by $X_{T}$ (i.e., excluding the items mapped to the stash). In the end of this phase, for every $j \in[\beta], P_{1}$ holds the OPPRF result $y_{j}^{\star}$ and $P_{2}$ holds the target value $t_{j}$.
2. Reverse the roles of $P_{1}$ and $P_{2}$ and run steps 1-2 of Protocol 9 again, where $P_{1}$ inputs $X_{S}=X \backslash X_{T}$ (i.e., only the items that were previously mapped to the stash) and $P_{2}$ inputs $Y$. Since the roles are reversed then $P_{1}$ maps $X_{S}$ using simple hashing and $P_{2}$ maps $Y$ using Cuckoo hashing. Denote the items mapped to the table and stash of $P_{2}$ by $Y_{T}$ and $Y_{S}$, respectively. In the end of this phase, for every $j \in[\beta], P_{1}$ has the target value $\tilde{t}_{j}$ and $P_{2}$ has the OPPRF result $\tilde{y}_{j}^{\star}$.
3. The parties use secure two-party computation (Functionality 1) with the circuit $C_{2 \beta+s^{2}, \gamma}$ (where $s$ is the stash size). For section $j \in[\beta]$ of the circuit, $P_{1}$ and $P_{2}$ input the first $\gamma$ bits of $y_{j}^{\star}$ and $t_{j}$ resp. For section $j \in\{\beta+1, \ldots, 2 \beta\}$ of the circuit $P_{1}$ and $P_{2}$ input the first $\gamma$ bits of $\tilde{t}_{j}$ and $\tilde{y}_{j}^{\star}$, respectively. Finally, for the rest $s^{2}$ sections of the circuit, the parties input every combination of $X_{S} \times Y_{S}$ (padded with uniformly random items so that $\left.\left|X_{S}\right|=\left|Y_{S}\right|=s\right)$.

Correctness \& Efficiency. The protocol compares every pair in $X \times Y$ and therefore every item in the intersection is input to the circuit exactly once: Sections $1, \ldots, \beta$ of the circuit cover all pairs in $X_{T} \times Y$, sections $\beta+1, \ldots, 2 \beta$ cover all pairs in $X_{S} \times Y_{T}$ and sections $2 \beta+1, \ldots, 2 \beta+s^{2}$ covers all pairs in $X_{S} \times Y_{S}$. This implies that the result of the three-phase construction is exactly the intersection $X \cap Y$. The communication complexity in the first two steps of the protocol is $O(n \cdot \kappa)$ as they involve the execution of a OPPRF with at most $O(n)$ items to the parties. The communication complexity of the third step is $O(n \cdot \gamma)$ since it involves $2 n+s^{2}$ comparisons of $\gamma$-bit elements. Since the stash size is $s=O(\log n)$, overall there are $O(n)$ comparisons.

Security. As in the basic protocol (see Sect.4.1), the security of this protocol is implied by the security of the OPPRF and secure two-party computation.

### 5.2 Linear Communication via Stash-Less Cuckoo Hashing

The largest communication cost factor in our protocols is the secure evaluation of the circuit. The asymptotically efficient Protocol 10 requires computing at least two copies of the basic circuit (for Phases 1 and 2), and it is therefore preferable to implement a protocol which has better concrete efficiency. We design a protocol that requires no stash (while achieving a small failure probability of less than $2^{-40}$ ), and hence uses no dual execution.

In order to be able to not use the stash, hashing is done with $K>2$ hash functions. We take into account the results of [PSZ18], which ran an empirical evaluation for the failure probability of Cuckoo hashing (failure is defined as the event where an item cannot be stored in the table and must be stored in the stash). They run experiments for a failure probability of $2^{-30}$ with $K=3,4$ and 5 hash functions, and extrapolated the results to yield the minimum number of bins for achieving a failure probability of less than $2^{-40}$. The results showed that $\beta=1.27 n, 1.09 n$, and $1.05 n$ bins are required for $K=3,4$, and 5 , respectively.

The main obstacle in using more than two hash functions in previous works on PSI was that the communication was still linear in $O\left(\max _{b} \cdot \beta\right)$, where $\max _{b}$ is the maximal number of elements in a bin of the simple hash table. The value of $\max _{b}$ increases with $K$ since each item is stored $K$ times in the simple hash table. In our protocol the communication for the circuit is independent of $\max _{b}$, as it only depends on the number of bins $\beta$. The communication for sending the polynomials, whose size is $O(K \cdot n \cdot \kappa)$, is just a small fraction of the overall communication and was in our experiments always smaller than $3 \%$. In this paper, we therefore use $K=3$ hash functions for our stash-less protocol.

## 6 PSI with Associated Payload

In many cases, each input item of the parties has some "payload" data associated with it. For example, an input item might include an id which is a credit card number, and a payload which is a transaction that was paid using this credit card. The parties might wish to compute some function of the payloads of the items in the intersection (for example, the sum of the financial transactions associated with these items). However, a straightforward application of our techniques does not seem to support this type of computation: Recall that $P_{2}$ might map multiple items to each bin. The OPPRF associates a single output $\beta$ to all these items, and this value is compared in the circuit with the output $\alpha$ of $P_{1}$. But if $P_{2}$ inserts a single item to the circuit, it seems that this item cannot encode the payloads of all items mapped to this bin.

The 2D Cuckoo hashing circuit-based PSI protocol of [PSWW18] handles payloads well, since each comparison involves only a single item from each party. While our basic protocol cannot handle payloads, we show here how it can be adapted to efficiently encode payloads in the input to the circuit.

Let Table ${ }_{1}$ and Stash be $P_{1}$ 's table and stash after mapping its items using Cuckoo hashing and let Table $2_{2}$ be $P_{2}$ 's table after mapping its items using simple hashing. In addition, denote by $U(x)$ and $V(y)$ the payloads associated with $x \in X$ and $y \in Y$ respectively and assume that all payloads have the same length $\delta$. The parties invoke two instances of batch OPPRF as follows:

1. A batch OPPRF where $P_{1}$ inputs Table $_{1}[1], \ldots, \operatorname{Table}_{1}[\beta]$ and $P_{2}$ inputs Table $_{2}[1], \ldots$, Table $_{2}[\beta]$ and $T_{1}, \ldots, T_{\beta}$ where $T_{j}$ has $\mid$ Table $_{2}[j] \mid$ elements, all equal to a uniformly random and independent value $t_{j} \in\{0,1\}^{\lambda}$. This is the same invocation of a batch OPPRF as in Protocol 9. At the end, $P_{1}$ has the OPPRF results $y_{1}^{\star}, \ldots, y_{\beta}^{\star}$ and $P_{2}$ has the target values $t_{1}, \ldots, t_{j}$.
2. In the second batch OPPRF, $P_{2}$ chooses the target values such that the elements in the set $T_{j}$ are not equal. Specifically, $P_{1}$ inputs Table $1[1], \ldots$, Table $_{1}[\beta]$ and $P_{2}$ samples $\tilde{t}_{1}, \ldots, \tilde{t}_{\beta}$ uniformly, and inputs $\operatorname{Table}_{2}[1], \ldots$, Table $_{2}[\beta]$ and $T_{1}, \ldots, T_{\beta}$ where $T_{j}(i)=\tilde{t}_{j} \oplus V\left(\operatorname{Table}_{2}[j](i)\right)$. Denote the OPPRF results that $P_{1}$ obtains by $\tilde{y}_{1}^{\star}, \ldots, \tilde{y}_{\beta}^{\star}$.

Then, the circuit operates in the following way: For the $j$-th section, $P_{1}$ inputs Table $_{1}[j], y_{j}^{\star}, \tilde{y}_{j}^{\star}$ and $U\left(\right.$ Table $\left._{1}[j]\right)$, and $P_{2}$ inputs $t_{j}$ and $\tilde{t}_{j}$. The circuit compares $y_{j}^{\star}$ to $t_{j}$. If they are equal then it forwards to the sub-circuit that computes $f$ the item Table $_{1}[j]$ itself, $P_{1}$ 's payload $U\left(\right.$ Table $\left._{1}[j]\right)$ and $P_{2}$ 's payload $\tilde{y}_{j}^{\star} \oplus \tilde{t}_{j}$. This holds since if Table $_{1}[j]$ is the $i$-th item in $P_{2}$ 's table, namely, $\operatorname{Table}_{2}[j](i)$, then the value $\tilde{y}_{j}^{\star}$ received by $P_{1}$ is $\tilde{y}_{j}^{\star}=\tilde{t}_{j} \oplus V\left(\operatorname{Table}_{2}[j](i)\right)$. Thus, $\tilde{y}_{j}^{\star} \oplus \tilde{t}_{j}=$ $V\left(\right.$ Table $\left._{2}[j](i)\right)$ as required.

Efficiency. The resulting protocol has the same asymptotic complexity as our initial protocols without payloads. The number of comparisons in the circuit is the same as in the basic circuit.

Table 1. The results of [PSZ18] for the required stash sizes $s$ for $K=2$ hash functions and $\beta=2.4 n$ bins, and the minimum OPPRF output bitlength $\gamma$ to achieve failure probability $<2^{-40}$ when mapping $n$ elements into $\beta$ bins with Cuckoo hashing. For $K>2$ hash functions we choose a large enough number of bins $\beta$ to achieve stash failure probability $<2^{-40}$.

| \# Elements $n$ | $2^{8}$ | $2^{12}$ | $2^{16}$ | $2^{20}$ | $2^{24}$ |  |
| :--- | :--- | :--- | :--- | :--- | :--- | :--- |
| Stash size $s$ for $K=2$ | 12 | 6 | 4 | 3 | 2 |  |
| OPPRF output length $\gamma$ | $K=2, \beta=2.4 n$ | 50 | 54 | 58 | 62 | 66 |
|  | $K=3, \beta=1.27 n, s=0$ | 49 | 53 | 57 | 61 | 65 |
|  | $K=4, \beta=1.09 n, s=0$ | 49 | 53 | 57 | 61 | 65 |
|  | $K=5, \beta=1.05 n, s=0$ | 49 | 53 | 57 | 61 | 65 |

## 7 Concrete Costs

In this section we evaluate the concrete costs of our protocol for concrete values of the security parameters. We set the computational security parameter to $\kappa=128$, and the statistical security parameter to $\sigma=40$.

### 7.1 Parameter Choices for Sufficiently Small Failure Probability

For $K=2$ hash functions, following previous works on PSI (e.g., [PSSZ15, PSWW18]), we set the table size parameter for Cuckoo hashing to $\epsilon=0.2$, and use a Cuckoo table with $\beta=2 n(1+\epsilon)=2.4 n$ bins. The resulting stash sizes for mapping $n$ elements into $\beta=2.4 n$ bins, as determined by the experiments in
[PSZ18], are summarized in Table 1. Note that we use here concrete values for the stash size, and are aiming for a failure probability smaller than $2^{-40}$. This can either be achieved using the basic protocol of Sect. 4.1 with the right choice of the stash size, or by running the three rounds $O(n)$ complexity protocol of Sect. 5.

Another option is described in Sect. 5.2, where we use more than two hash functions (specifically, use $K=3,4$, or 5 functions), with the hash table being of size $\beta=1.27 n, 1.09 n$, or $1.05 n$, respectively. These parameters achieve a failure probability smaller than $2^{-40}$ according to the experimental analysis in [PSZ18].

As described in Sect. 4.2, even if there are no stash failures, the scheme can fail due to collisions in the output of the PRF, with probability $\beta \cdot 2^{-\gamma}$, where $\gamma$ is the output bitlength of the OPPRF. To make this failure probability smaller than the statistical security parameter (which we set to 40 ), the output bitlength of the OPPRF must be $\gamma=40+\log _{2} \beta$ bits.

### 7.2 Computing Polynomial Interpolation

We implemented interpolation of polynomials of degree $d$ using an $O\left(d^{2}\right)$ algorithm based on Lagrange interpolation in a prime field where the prime is the Mersenne prime $2^{61}-1$. The runtime for interpolating a polynomial of degree $d=1024$ was 7 ms , measured on an Intel Core i7-4770K CPU with a single thread. The runtimefor different values of $d$ behaved (very accurately) as a quadratic function of $d$. The actual algorithmsare those implemented in NTL v10.0 with field arithmetics replaced with our customized arithmetic operationsover the Mersenne prime $2^{61}-1$. Most importantly, this field enables an order of magnitude faster multiplication of field elements: multiplying $x \cdot y$ with $|x|,|y| \leq 61$ is implemented by multiplying $x$ and $y$ over $\mathbb{Z}$ to obtain $z=x y$ with $|z| \leq 122$. Then the result is the sum of the element represented by the lower 61 bits of $z$ with the element represented by the higher 61 bits of $z$ (and therefore no expensive modular reduction is required). The Mersenne prime $2^{61}-1$ allows the use of at least 40-bit statistical security for up to $n=2^{20}$ elements for all our algorithms using permutation-based hashing (cf. [PSSZ15]). To use larger sets, we see two possible solutions: (i) using a larger Mersenne prime or (ii) reducing the statistical security parameter $\sigma$ (e.g., using $\sigma=38$ for achieving less than $2^{-\sigma}$ failure probability for $n=2^{22}$ elements, $K=3$ hash functions, and $\beta=1.27 n$ bins). The required minimum bit-length of the elements using permutation-based hashing with failure probability $2^{-\sigma}$ is computed as $\ell=\sigma+2 \log _{2} n-\log _{2} \beta$. The OPPRF output is also $\leq 61$ bits in most cases as shown in Table 1 .

For reducing the computation complexity of our protocol, we use the approach described in Sect. 4.3, where instead of interpolating a polynomial of degree $K \cdot n$, where $K$ is the number of hash functions and $n$ is the number of elements for PSI, we interpolate multiple smaller polynomials of degree at most $d=1024$. We therefore have to determine the minimum number of mega-bins $B$ such that when mapping $N=K \cdot n$ elements to $B$ bins, the probability of having a bin with more than $\max _{b}=1024$ elements is smaller than $2^{-40}$. As in the analysis for simple hashing in [PSZ18], we use the formula from [MR95]:

$$
\begin{aligned}
P\left(" \exists \text { bin with } \geq \max _{b} \text { elements" }\right) & \leq \sum_{i=1}^{B} P\left(\text { "bin } i \text { has } \geq \max _{b} \text { elements" }\right) \\
& =B \cdot \sum_{i=\max _{b}}^{N}\binom{N}{i} \cdot\left(\frac{1}{B}\right)^{i} \cdot\left(1-\frac{1}{B}\right)^{N-i} .
\end{aligned}
$$

We depict the corresponding numbers in Table 2. With these numbers and our experiments for polynomial interpolation described above, the estimated runtimes for the polynomial interpolation are $B \cdot 7 \mathrm{~ms}$. The hints (polynomials) that need to be sent have size $B \cdot \max _{b} \cdot \gamma$ bits which is only slightly larger than the ideal communication of $K \cdot n \cdot \gamma$ bits when using one large polynomial as shown in Table 2.

Note that in contrast to many PSI solutions whose main run-time bottleneck is already network bandwidth (which cannot be easily improved in many settings such as over the Internet), the run-time of our protocols can be improved by using multiple threads instead of one thread. Since the interpolation of polynomials for different mega-bins is independent of each other, the computation scales linearly in the number of physical cores and thus can be efficiently parallelized.

Table 2. Parameters for mapping $N=K \cdot n$ elements to $B$ mega-bins s.t. each megabin has at most $\max _{b} \leq 1024$ elements with probability smaller than $2^{-40}$. The lower half of the table contains the expected costs for the polynomial interpolations.

| \# hash functions | $K=2$ |  |  | $K=3$ |  |  |
| :--- | ---: | ---: | ---: | ---: | ---: | ---: |
| Set size | $n=2^{12}$ | $n=2^{16}$ | $n=2^{20}$ | $n=2^{12}$ | $n=2^{16}$ | $n=2^{20}$ |
| \# mega-bins $B$ | 11 | 165 | 2663 | 16 | 248 | 4002 |
| Maximum number of <br> elements max | 944 | 1021 | 1024 | 975 | 1021 | 1024 |
| Polynomial interpola- <br> tion [in milliseconds] | 126 | 1815 | 29293 | 183 | 2809 | 45335 |
| Size of hints [in bits] | 560736 | 9770970 | 169068544 | 826800 | 14432856 | 249980928 |
| Ideal size of hints for <br> one polynomial [in <br> bits] | 436330 | 7505580 | 128477895 | 651264 | 11206656 | 191889408 |

### 7.3 Communication and Depth Comparison

We first compute the communication complexity of our basic construction from Sect.4.1. The communication is composed of (a) the OPRF evaluations for each of the $B$ bins, (b) the hints consisting of the polynomials, (c) the circuit for comparing the outputs of the OPPRFs in each bin, and (d) the circuit for comparing the $s$ elements on the stash with the $n$ elements of $P_{2}$.

With regards to (a), the OPRF protocol of [KKRT16], which was also used in $[\mathrm{KMP}+17]$, has an amortized communication of at most 450 bits per OPRF
evaluation for set sizes up to $n=2^{24}$ elements (cf. [KKRT16, Table 1]). This amounts to $B \cdot 450$ bits of communication.

With regards to (b), for the size of the hints in the OPPRF construction we use the values given in Table 2. These numbers represent the communication when using mega-bins, and are slightly larger than the ideal communication of $K \cdot n$ coefficients of size $\gamma$ bits each, that would have been achieved by using a single polynomial for all values. However, it is preferable to use mega-bins since their usage substantially improves the computation complexity as described in Sect.4.3, while the total communication for the hints is at most $3 \%$ of the total communication. (This also shows that any improvements of the size of the hints will have only a negligible effect on the total communication).

With regards to (c), the circuit compares $B$ elements of bitlength $\gamma$, and hence requires $B \cdot(\gamma-1)$ AND gates. With 256 bits per AND gate [ALSZ13, ZRE15] this yields $B \cdot(\gamma-1) \cdot 256$ bits of communication.

With regards to (d), the final circuit consists of $s \cdot n$ comparisons of bitlength $\sigma$. This requires $\operatorname{sn}(\sigma-1) \cdot 256$ bits of communication.

We now analyze the communication complexity of our $O(n)$ protocol described in Sect.5. The main difference compared to the basic protocol analyzed above is that a different method is used for comparing the elements of the stash, i.e., replacing step (d) above. The new method replaces this step by letting $P_{2}$ use Cuckoo hashing of its $n$ elements into $B$ bins and then evaluating OPRF for each of these bins. This requires $B \cdot 450$ bits of communication plus $B$ comparisons of $\gamma$ bit values. Overall, this amounts to $B \cdot(450+(\gamma-1) \cdot 256)$ bits of communication. For simplicity, we omit the communication for comparing the elements for phase 3 which compares the elements on the two stashes, as it is negligible.

Comparison to Previous Work. In Table 3, we compare the resulting communication of our protocols to those of previous circuit-based PSI protocols of [HEK12,PSSZ15,PSWW18,FNO18]. As can be seen from this table, our protocols improve communication by an integer factor, where the main advantage of our protocols is that their communication complexity is independent of the bitlength of the input elements. Namely, for arbitrary input bitlengths, our nostash protocol improves the communication over the previous best protocol of [PSWW18] by a factor of $12.8 x$ for $n=2^{12}$ to a factor of $10.1 x$ for $n=2^{20}$. For fixed bitlength of $\sigma=32$ bits, our no-stash protocol improves communication over [PSWW18] by a factor of $5.7 x$ for $n=2^{12}$ to a factor of 2.6x for $n=2^{20}$.

Circuit Depth. For some secure circuit evaluation protocols like GMW [GMW87] the round complexity depends on the depth of the circuit. In Table 4, we depict the circuit depths for concrete parameters of our protocols and previous work, and show that our circuits have about the same low depth as the best previous works [PSSZ15,PSWW18]. In more detail, the Sort-CompareShuffle (SCS) circuit of [HEK12] has depth $\log _{2} \sigma \cdot \log _{2} n$ when using depthoptimized comparison circuits. The protocols of [PSSZ15,PSWW18] have depth

Table 3. Communication in MB for circuit-based PSI on $n$ elements of fixed bitlength $\sigma=32$ (left) and arbitrary bitlength hashed to $\sigma=40+2 \log _{2}(n)-1$ bits (right). The numbers for previous protocols are based on the circuit sizes given in [PSWW18, Table 3] with 256 bit communication per AND gate. The best values are marked in bold.

| Protocol | $\sigma=32$ |  |  | Arbitrary $\sigma$ |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: |
|  | $2^{12}$ | $2^{16}$ | $2^{20}$ | $2^{12}$ | $2^{16}$ | $2^{20}$ |
| SCS [HEK12] | 104 | 2174 | 42976 | 205 | 4826 | 106144 |
| Circuit-Phasing [PSSZ15] | 130 | 1683 | 21004 | 320 | 5552 | 97708 |
| Hashing + SCS [FNO18] | - | 1537 | 21207 | - | 3998 | 72140 |
| 2D CH [PSWW18] | 51 | 612 | 6582 | 115 | 1751 | 25532 |
| Ours Basic Sect. 4.1 | 41 | 550 | 8123 | 65 | 870 | 12731 |
| Ours Advanced Sect. 5 | 35 | 604 | 10277 | 35 | 604 | 10277 |
| Ours No-Stash Sect. 5.2 | 9 | 149 | 2540 | 9 | 149 | 2540 |
| Breakdown: |  |  |  |  |  |  |
| OPRF | 0.3 (3\%) | 5 (3\%) | 72 (3\%) | 0.3 (3\%) | 5 (3\%) | 72 (3\%) |
| Sending polynomials | 0.1 (1\%) | 2 (1\%) | 30 (1\%) | 0.1 (1\%) | 2 (1\%) | 30 (1\%) |
| Circuit | 9 (96\%) | $142(96 \%)$ | 2438 (96\%) | 9 (96\%) | $142(96 \%)$ | 2438 (96\%) |
| Improvement factor | 5.7 x | 4.1 x | 2.6 x | 12.8 x | 11.8 x | 10.1x |

$\log _{2} \sigma$. A depth-optimized SCS circuit for the construction in [FNO18] has depth $\log _{2}\left(\sigma-\log _{2}(n / b)\right) \cdot \log _{2}((1+\delta) b)$, where concrete parameters for $n, \delta, b$ are given in [FNO18, Table 1]. Our protocols consist of circuits for comparing the elements on the stash of bitlength $\sigma$ and the outputs of the OPPRFs of length $\gamma$ and therefore have depth $\max (\log \sigma, \log \gamma)=\max \left(\log \sigma, \log _{2}\left(40+2 \log _{2}(n)-1\right)\right)$.

Table 4. Circuit depth for circuit-based PSI on $n$ elements of fixed bitlength $\sigma=$ 32 (left) and arbitrary bitlength hashed to $\sigma=40+2 \log _{2}(n)-1$ bits (right).

|  | $\sigma=32$ |  |  | Arbitrary $\sigma$ |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| Protocol $n=$ |  | $2^{16}$ | $2^{20}$ | $2^{12}$ | $2^{16}$ | $2^{20}$ |
| SCS [HEK12] | 60 | 80 | 100 | 72 | 98 | 126 |
| Circuit-Phasing [PSSZ15] | 5 | 5 | 5 | 6 | 7 |  |
| Hashing + SCS [FNO18] | - | 42 | 36 |  | 54 | 51 |
| 2D CH [PSWW18] | 5 | 5 | 5 | 6 | 7 |  |
| Our Protocols | 6 | 6 | 6 | 6 | 7 |  |

### 7.4 Runtime Comparison

In this section we compare the runtimes of different PSI protocols. In Sect. 7.2 we conducted experiments for polynomial interpolation, the main new part of our protocol, and we show below that this step takes only a small fraction of the total runtime. We also implemented our most efficient protocol (see Sect. 5.2). ${ }^{6}$ In addition, we estimate the runtime of our less efficient basic protocol (see

[^4]Sect.4.1) and the protocol with linear communication overhead (see Sect.5) based on the experiments of the interpolation procedure and rigorous estimations from previous works.

Previous Work. As we have seen in the analysis of the communication overhead in Sect. 7.3, our protocols provide better improvements to performance in the case of arbitrary bitlengths. The previous work of [PSWW18] gave runtimes only for fixed bitlength of 32 bits in [PSWW18, Table 4]. Therefore, we extrapolate the runtimes of the previous protocols from fixed bitlength to arbitrary bitlength based on the circuit sizes given in [PSWW18, Table 3]. The estimated runtimes are given in Table 5. The LAN setting is a 1 Gbit/s network with roundtrip time of 1 ms and the WAN setting is a $100 \mathrm{Mbit} / \mathrm{s}$ network with round-trip time of 100 ms . Runtimes were not presented in [FNO18], but since their circuit sizes and depths are substantially larger than those of [PSWW18] (cf. Tables 3 and 4), their runtimes will also be substantially higher than those of [PSWW18].

Our Implementation. We implemented and benchmarked our most efficient no-stash OPPRF-based PSI protocol (see Sect.5.2) on two commodity PCs with an Intel Core i7-4770K CPU. We instantiated our protocol with security parameter $\kappa=128$ bits, $K=3$ hash functions, $B=1.27 n$ bins, and no stash (see Sect.5.2). Our OPPRF implementation is based on the OPRF protocol of [PSZ18]. ${ }^{7}$ For the secure circuit evaluation, we used the ABY framework [DSZ15]. The run-times are averaged over 50 executions. The results are described in Table 5.

Comparison with PSI Protocols. As a baseline, we compare our performance with specific protocols for computing the intersection itself. (However, as is detailed in Sect. 1.2, our protocol is circuit-based and therefore has multiple advantages compared to specific PSI protocols). Our best protocol is slower by a factor of 41x than today's fastest PSI protocol of [KKRT16] for $n=2^{20}$ elements in the WAN setting (cf. Table 5).

Comparison with Public Key-Based PSI Variant Protocols. Our circuitbased protocol is substantially faster than previous public key-based protocols for computing variants of PSI, although they have similar asymptotic linear complexity. As an example, consider comparing whether the size of the intersection is greater than a threshold (PSI-CAT). In our protocol, we can compute the PSI-CAT functionality by extending the PSI circuit of Table 5 with a Hamming distance circuit (which, using the size-optimal construction of [BP06], adds less than $n$ AND gates). The final comparison with the threshold adds another $\log _{2} n$

[^5]Table 5. Total run-times in ms for PSI variant protocols on $n$ elements of arbitrary bitlength using GMW [GMW87] for secure circuit evaluation and one thread. Numbers for all but our protocols are based on [PSWW18]. The best values for generic circuitbased PSI protocols are marked in bold.

| Network | LAN |  |  | WAN |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| Protocol $n=$ | $2^{12}$ | $2^{16}$ | $2^{20}$ | $2^{12}$ | $2^{16}$ | $2^{20}$ |
| Special-purpose PSI protocols (as baseline) |  |  |  |  |  |  |
| DH/ECC PSI <br> [Sha80, Mea86, DGT12] <br> BaRK-OPRF [KKRT16] | $\begin{gathered} 3296 \\ 113 \end{gathered}$ | $\begin{array}{r} \hline 49010 \\ 295 \\ \hline \end{array}$ | $\begin{array}{\|r\|r\|} \hline 7904054 \\ 3882 \\ \hline \end{array}$ | $\begin{array}{r} 4082 \\ 540 \end{array}$ | $\begin{gathered} 51866 \\ 1247 \end{gathered}$ | 8008771 14604 |
| Generic circuit-based PSI protocols |  |  |  |  |  |  |
| Circuit-Phasing [PSSZ15] <br> 2D CH [PSWW18] | $\begin{aligned} & 7825 \\ & 5031 \end{aligned}$ | 67292 <br> 25960 | 1126848 <br> 336134 | 37380 22796 | $\begin{aligned} & 327976 \\ & 129436 \end{aligned}$ | $\begin{aligned} & 4850571 \\ & 1512505 \end{aligned}$ |
| Ours Basic Sect. 4.1 (estimated) Ours Advanced Sect. 5 (estimated) | 2908 1674 | 13767 <br> 9763 | 182204 <br> 148436 | 12934 7372 | 63861 43675 | 752695 597885 |
| Ours No-Stash Sect. 5.2, Total | 1199 | -8486 | 120731 | 5910 | 22134 | 261481 |
| Breakdown: |  |  |  |  |  |  |
| OPRF | 724 (60\%) | 1097 (13\%) | 5844 (5\%) | 2867 (49\%) | 4164 (19\%) | 26121 (10\%) |
| Polynomial interpolation | 183 (15\%) | 2809 (33\%) | 45335 (38\%) | 183 (3\%) | 2809 (13\%) | 45335 (17\%) |
| Polynomial transmission | 16 (1\%) | $145 \quad(2 \%)$ | 667 (0\%) | 816 (13\%) | 1079 (5\%) | 4012 (2\%) |
| Polynomial evaluation | 58 (5\%) | 1344 (16\%) | 21768 (18\%) | 58 (1\%) | 1344 (6\%) | 21768 (8\%) |
| Circuit | 218 (18\%) | 3091 (36\%) | 47117 (39\%) | 1986 (34\%) | $12738(57 \%)$ | 164245 (63\%) |
| Improvement over [PSWW18] | 4.2 x | 3.1x | 2.8 x | 3.9 x | 5.8 x | 5.8 x |

AND gates [BPP00] which are negligible as well. For the PSI-CAT functionality, [ZC17] report runtimes of 779 s for $n=2^{11}$ elements, [HOS17] report runtimes of 728 s for $n=2^{11}$ elements, and [ZC18] report runtimes of at least 138 s for $n=100$ elements, whereas our protocol requires 0.52 s for $n=2^{11}$ elements and 0.34 s for $n=100$ elements. Hence, we improve over [ZC17] by a factor of 1498 x , over [HOS17] by a factor of 1400 x , and over [ZC18] by a factor of 405 x . As an example for computing PSI-CAT with larger set sizes, our protocol requires 124 s for $n=2^{20}$ elements.

The protocol described by Google for computing ad revenues [Yun15, Kre17] (see Sect. 1.2) is based on the DH-based PSI protocol which is already 65 x slower than our protocol for $n=2^{20}$ elements over a LAN (cf. Table 5) and leaks the intersection cardinality as an intermediate result. Here, too, our circuit would be only slightly larger than the PSI circuit of Table 5.

Comparison with Circuit-Based PSI Protocols. As can be seen from Table 5, our no-stash protocol from Sect. 5.2 is substantially more efficient than our basic protocol and our linear asymptotic overhead protocol from Sect.4.1 and Sect.5, respectively. It improves over the best previous circuit-based PSI protocol from [PSWW18] by factors of 4.2 x to 2.8 x in the LAN setting, and by factors of 5.8 x to 3.9 x in the WAN setting. From the micro-benchmarks in Table 5, we also observe that the runtimes for the polynomial interpolation are a significant fraction of the total runtimes of our protocols ( $3 \%$ to $33 \%$ for the interpolation and $1 \%$ to $18 \%$ for the evaluation). Since polynomials are independent of each other, the interpolation and evaluation can be trivially parallelized
for running with multiple threads, which would give this part of the computation a speed-up that is linear in the number of physical cores of the processor.

Acknowledgements. We thank Ben Riva and Udi Wieder for valuable discussions about this work. This work has been co-funded by the DFG within project E4 of the CRC CROSSING and by the BMBF and the HMWK within CRISP, by the BIU Center for Research in Applied Cryptography and Cyber Security in conjunction with the Israel National Cyber Bureau in the Prime Ministers Office, and by a grant from the Israel Science Foundation.

## References

[ADN+13] Asokan, N., et al.: CrowdShare: secure mobile resource sharing. In: Jacobson, M., Locasto, M., Mohassel, P., Safavi-Naini, R. (eds.) ACNS 2013. LNCS, vol. 7954, pp. 432-440. Springer, Heidelberg (2013). https://doi. org/10.1007/978-3-642-38980-1_27
[ALSZ13] Asharov, G., Lindell, Y., Schneider, T., Zohner, M.: More efficient oblivious transfer and extensions for faster secure computation. In: CCS (2013)
[BP06] Boyar, J., Peralta, R.: Concrete multiplicative complexity of symmetric functions. In: Královič, R., Urzyczyn, P. (eds.) MFCS 2006. LNCS, vol. 4162, pp. 179-189. Springer, Heidelberg (2006). https://doi.org/10.1007/ 11821069_16
[BPP00] Boyar, J., Peralta, R., Pochuev, D.: On the multiplicative complexity of Boolean functions over the basis $(\wedge, \oplus, 1)$. TCS 235(1), 43-57 (2000)
[CADT14] H. Carter, C. Amrutkar, I. Dacosta, and P. Traynor. For your phone only: custom protocols for efficient secure functionevaluation on mobile devices. Secur. Commun. Netw. 7(7) (2014)
[CO18] Ciampi, M., Orlandi, C.: Combining private set-intersection with secure two-party computation. In: Catalano, D., De Prisco, R. (eds.) SCN 2018. LNCS, vol. 11035, pp. 464-482. Springer, Cham (2018). https://doi.org/ 10.1007/978-3-319-98113-0_25
[DC17] Davidson, A., Cid, C.: An efficient toolkit for computing private set operations. In: Pieprzyk, J., Suriadi, S. (eds.) ACISP 2017. LNCS, vol. 10343, pp. 261-278. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-59870-3_15
[DD15] Debnath, S.K., Dutta, R.: Secure and efficient private set intersection cardinality using bloom filter. In: Lopez, J., Mitchell, C.J. (eds.) ISC 2015. LNCS, vol. 9290, pp. 209-226. Springer, Cham (2015). https://doi. org/10.1007/978-3-319-23318-5_12
[DGT12] De Cristofaro, E., Gasti, P., Tsudik, G.: Fast and private computation of cardinality of set intersection and union. In: Pieprzyk, J., Sadeghi, A.-R., Manulis, M. (eds.) CANS 2012. LNCS, vol. 7712, pp. 218-231. Springer, Heidelberg (2012). https://doi.org/10.1007/978-3-642-35404-5_17
[DKT10] De Cristofaro, E., Kim, J., Tsudik, G.: Linear-complexity private set intersection protocols secure in malicious model. In: Abe, M. (ed.) ASIACRYPT 2010. LNCS, vol. 6477, pp. 213-231. Springer, Heidelberg (2010). https://doi.org/10.1007/978-3-642-17373-8_13
[DSMRY09] Dachman-Soled, D., Malkin, T., Raykova, M., Yung, M.: Efficient robust private set intersection. In: Abdalla, M., Pointcheval, D., Fouque, P.-A., Vergnaud, D. (eds.) ACNS 2009. LNCS, vol. 5536, pp. 125-142. Springer, Heidelberg (2009). https://doi.org/10.1007/978-3-642-01957-9_8
[DSZ15] Demmler, D., Schneider, T., Zohner, M.: ABY - a framework for efficient mixed-protocol secure two-party computation. In: NDSS (2015)
[DT10] De Cristofaro, E., Tsudik, G.: Practical private set intersection protocols with linear complexity. In: Sion, R. (ed.) FC 2010. LNCS, vol. 6052, pp. 143-159. Springer, Heidelberg (2010). https://doi.org/10.1007/978-3-642-14577-3_13
[Dwo06] Dwork, C.: Differential privacy. In: ICALP (2006)
[EFLL12] Ejgenberg, Y., Farbstein, M., Levy, M., Lindell, Y.: SCAPI: the secure computation application programming interface. Cryptology ePrint Archive, Report 2012/629 (2012)
[FHNP16] Freedman, M.J., Hazay, C., Nissim, K., Pinkas, B.: Efficient set intersection with simulation-based security. J. Cryptol. 29(1), 115-155 (2016)
[FIPR05] Freedman, M.J., Ishai, Y., Pinkas, B., Reingold, O.: Keyword search and oblivious pseudorandom functions. In: Kilian, J. (ed.) TCC 2005. LNCS, vol. 3378, pp. 303-324. Springer, Heidelberg (2005). https://doi.org/10. 1007/978-3-540-30576-7_17
[FNO18] Falk, B.H., Noble, D., Ostrovsky, R.: Private set intersection with linear communication from general assumptions. Cryptology ePrint Archive, Report 2018/238 (2018)
[FNP04] Freedman, M.J., Nissim, K., Pinkas, B.: Efficient private matching and set intersection. In: Cachin, C., Camenisch, J.L. (eds.) EUROCRYPT 2004. LNCS, vol. 3027, pp. 1-19. Springer, Heidelberg (2004). https://doi.org/ 10.1007/978-3-540-24676-3_1
[GM11] Goodrich, M.T., Mitzenmacher, M.: Privacy-preserving access of outsourced data via oblivious RAM simulation. In: Aceto, L., Henzinger, M., Sgall, J. (eds.) ICALP 2011. LNCS, vol. 6756, pp. 576-587. Springer, Heidelberg (2011). https://doi.org/10.1007/978-3-642-22012-8_46
[GMW87] Goldreich, O., Micali, S., Wigderson, A.: How to play any mental game or a completeness theorem for protocols with honest majority. In: STOC (1987)
[Gon81] Gonnet, G.H.: Expected length of the longest probe sequence in hash code searching. J. ACM 28(2), 289-304 (1981)
[HCE11] Huang, Y., Chapman, P., Evans, D.: Privacy-preserving applications on smartphones. In: Hot Topics in Security (HotSec) (2011)
[HEK12] Huang, Y., Evans, D., Katz, J.: Private set intersection: are garbled circuits better than custom protocols? In: NDSS (2012)
[HEKM11] Huang, Y., Evans, D., Katz, J., Malka, L.: Faster secure two-party computation using garbled circuits. In: USENIX Security (2011)
[HN10] Hazay, C., Nissim, K.: Efficient set operations in the presence of malicious adversaries. In: Nguyen, P.Q., Pointcheval, D. (eds.) PKC 2010. LNCS, vol. 6056, pp. 312-331. Springer, Heidelberg (2010). https://doi.org/10. 1007/978-3-642-13013-7_19
[HOS17] Hallgren, P., Orlandi, C., Sabelfeld, A.: PrivatePool: privacy-preserving ridesharing. In: Computer Security Foundations Symposium (CSF) (2017)
$[$ IKN+17] Ion, M., et al.: Private intersection-sum protocol with applications to attributing aggregate ad conversions. Cryptology ePrint Archive, Report 2017/738 (2017)
[IKNP03] Ishai, Y., Kilian, J., Nissim, K., Petrank, E.: Extending oblivious transfers efficiently. In: Boneh, D. (ed.) CRYPTO 2003. LNCS, vol. 2729, pp. 145-161. Springer, Heidelberg (2003). https://doi.org/10.1007/978-3-540-45146-4_9
[KKRT16] Kolesnikov, V., Kumaresan, R., Rosulek, M., Trieu, N.: Efficient batched oblivious PRF with applications to private set intersection. In: CCS (2016)
[KLS+17] Kiss, Á., Liu, J., Schneider, T., Asokan, N., Pinkas, B.: Private set intersection for unequal set sizes with mobile applications. PoPETs 2017(4), 177-197 (2017)
[KM18] Kushilevitz, E., Mour, T.: Sub-logarithmic distributed oblivious RAM with small block size. CoRR, abs/1802.05145 (2018)
[KMP+17] Kolesnikov, V., Matania, N., Pinkas, B., Rosulek, M., Trieu, N.: Practical multi-party private set intersection from symmetric-key techniques. In: CCS (2017)
[KMW09] Kirsch, A., Mitzenmacher, M., Wieder, U.: More robust hashing: cuckoo hashing with a stash. SIAM J. Comput. 39(4), 1543-1561 (2009)
[Kre17] Kreuter, B.: Secure multiparty computation at Google. In: RWC (2017)
[Kre18] Kreuter, B.: Techniques for Scalable Secure Computation Systems. Ph.D. thesis, Northeastern University (2018)
[KS08] Kolesnikov, V., Schneider, T.: Improved garbled circuit: free XOR gates and applications. In: Aceto, L., Damgård, I., Goldberg, L.A., Halldórsson, M.M., Ingólfsdóttir, A., Walukiewicz, I. (eds.) ICALP 2008. LNCS, vol. 5126, pp. 486-498. Springer, Heidelberg (2008). https://doi.org/10.1007/ 978-3-540-70583-3_40
[LWN+15] Liu, C., Wang, X. S., Nayak, K., Huang, Y., Shi, E.: ObliVM: a programming framework for secure computation. In: S\&P (2015)
[Mea86] Meadows, C.: A more efficient cryptographic matchmaking protocol for use in the absence of a continuously available third party. In: S\&P (1986)
[MR95] Motwani, R., Raghavan, P.: Randomized Algorithms (1995)
[PR01] Pagh, R., Rodler, F.F.: Cuckoo hashing. In: auf der Heide, F.M. (ed.) ESA 2001. LNCS, vol. 2161, pp. 121-133. Springer, Heidelberg (2001). https://doi.org/10.1007/3-540-44676-1_10
[PSSZ15] Pinkas, B., Schneider, T., Segev, G., Zohner, M.: Phasing: private set intersection using permutation-based hashing. In: USENIX Security (2015)
[PSWW18] Pinkas, B., Schneider, T., Weinert, C., Wieder, U.: Efficient circuit-based PSI via cuckoo hashing. In: Nielsen, J.B., Rijmen, V. (eds.) EUROCRYPT 2018. LNCS, vol. 10822, pp. 125-157. Springer, Cham (2018). https://doi. org/10.1007/978-3-319-78372-7_5
[PSZ14] Pinkas, B., Schneider, T., Zohner, M.: Faster private set intersection based on OT extension. In: USENIX Security (2014)
[PSZ18] Pinkas, B., Schneider, T., Zohner, M.: Scalable private set intersection based on OT extension. TOPS 21(2), 7 (2018)
[RA18] Resende, A.C.D., Aranha, D.F.: Faster unbalanced private set intersection. In: FC (2018)
[RR17a] Rindal, P., Rosulek, M.: Improved private set intersection against malicious adversaries. In: Coron, J.-S., Nielsen, J.B. (eds.) EUROCRYPT 2017. LNCS, vol. 10210, pp. 235-259. Springer, Cham (2017). https:// doi.org/10.1007/978-3-319-56620-7_9
[RR17b] Rindal, P., Rosulek, M.: Malicious-secure private set intersection via dual execution. In: CCS (2017)
[Sha80] Shamir, A.: On the power of commutativity in cryptography. In: de Bakker, J., van Leeuwen, J. (eds.) ICALP 1980. LNCS, vol. 85, pp. 582595. Springer, Heidelberg (1980). https://doi.org/10.1007/3-540-100032_100
[SZ13] Schneider, T., Zohner, M.: GMW vs. Yao? Efficient secure two-party computation with low depth circuits. In: Sadeghi, A.-R. (ed.) FC 2013. LNCS, vol. 7859, pp. 275-292. Springer, Heidelberg (2013). https://doi.org/10. 1007/978-3-642-39884-1_23
[Yao86] Yao, A.C.: How to generate and exchange secrets. In: FOCS (1986)
[Yun15] Yung, M.: From mental poker to core business: why and how to deploy secure computation protocols? In: CCS (2015)
[ZC17] Zhao, Y., Chow, S.S.M.: Are you the one to share? Secret transfer with access structure. PoPETs 2017(1), 149-169 (2017)
[ZC18] Zhao, Y., Chow, S.S.M.: Can you find the one for me? Privacy-preserving matchmaking via threshold PSI. In: WPES (2018)
[ZRE15] Zahur, S., Rosulek, M., Evans, D.: Two halves make a whole: reducing data transfer in garbled circuits using half gates. In: Oswald, E., Fischlin, M. (eds.) EUROCRYPT 2015. LNCS, vol. 9057, pp. 220-250. Springer, Heidelberg (2015). https://doi.org/10.1007/978-3-662-46803-6_8


[^0]:    ${ }^{1}$ https://www.bloomberg.com/news/articles/2018-08-30/google-and-mastercard-cut-a-secret-ad-deal-to-track-retail-sales.
    ${ }^{2}$ Note that outputting the intersection is a non-symmetric function. Therefore in that case the order of the elements must be shuffled. However, it is unclear why a circuit-based protocol should be used for computing the intersection, since there are specialized protocols for this which are much more efficient, e.g. [KKRT16, PSZ18].

[^1]:    ${ }^{3}$ We require that each element in $T$ is uniformly random but the elements may be correlated.

[^2]:    ${ }^{4}$ In the actual implementation we use a more general variant of Cuckoo hashing with a parameter $K \in\{2,3\}$ where each item is stored in $K$ locations in the table. The size of the hint will be $K \cdot n$.

[^3]:    ${ }^{5}$ In that case either not all items are stored in the stash - resulting in the protocol ignoring part of the input and potentially computing the wrong output, or $P_{1}$ needs to inform $P_{2}$ that it uses a stash larger than $s$ - resulting in a privacy breach.

[^4]:    ${ }^{6}$ Our implementation is available at https://github.com/encryptogroup/OPPRF-PSI.

[^5]:    ${ }^{7}$ This OPRF protocol has communication that is higher by $10 \%$ to $15 \%$ than the communication of the OPRF protocol of [KKRT16]. But since OPRF requires less than $3 \%$ of the total communication, this additional cost is negligible in our protocol.

