# Efficient Set Intersection with Simulation-Based Security 

Michael J. Freedman<br>Department of Computer Science, Princeton University, Princeton, NJ, USA<br>mfreed@cs.princeton.edu<br>Carmit Hazay*<br>Faculty of Engineering, Bar-Ilan University, Ramat Gan, Israel<br>carmit.hazay@biu.ac.il

Kobbi Nissim ${ }^{\dagger}$<br>Department of Computer Science, Ben-Gurion University, Beer-Sheva, Israel<br>kobbi@cs.bgu.ac.il<br>Benny Pinkas ${ }^{\ddagger}$<br>Department of Computer Science, Bar-Ilan University, Ramat Gan, Israel<br>benny@pinkas.net

Communicated by Hugo Krawczyk
Received 6 December 2010
Online publication 24 October 2014


#### Abstract

We consider the problem of computing the intersection of private datasets of two parties, where the datasets contain lists of elements taken from a large domain. This problem has many applications for online collaboration. In this work, we present protocols based on the use of homomorphic encryption and different hashing schemes for both the semi-honest and malicious environments. The protocol for the semi-honest environment is secure in the standard model, while the protocol for the malicious environment is secure in the random oracle model. Our protocols obtain linear communication and computation overhead. We further implement different variants of our semi-honest protocol. Our experiments show that the asymptotic overhead of the protocol is affected by different constants. (In particular, the degree of the polynomials evaluated by the protocol matters less than the number of polynomials that are evaluated.) As a result, the protocol variant with the best asymptotic overhead is not necessarily preferable for inputs of reasonable size.


[^0]
## 1. Introduction

In the setting of secure two-party computation, two parties with private inputs wish to jointly compute some function of their inputs while preserving certain security properties such as privacy and correctness. The standard way of defining security in this setting is via the so-called ideal/real model paradigm $[9,14,31,52]$. Here, an ideal model is first defined where an uncorrupted trusted party is used to compute the function for the parties. Then, a real protocol is said to be secure if no adversary can do more harm in a real protocol execution than in an ideal one (where by definition no harm can be done). Starting with the work of $[11,15,33,65]$, it is by now well known that (in various settings, and considering semi-honest and malicious adversaries) any polynomial-time computation can be generically compiled into a secure function evaluation protocol with polynomial complexity. However, more often than not, the resulting protocols are inefficient for practical uses, and hence, attention was given to constructing efficient protocols for specific functions.

In this work, we consider the basic two-party set intersection problem, where two parties $P_{1}$ and $P_{2}$ hold input sets $X, Y$, respectively, and wish to compute $X \cap Y$. The secure variant of this computation guarantees that nothing beyond $X \cap Y$ leaks within the protocol execution. This problem has been widely studied due to its extensive usage in computations over databases, e.g., for data mining where the data are vertically partitioned between parties (namely, each party has different attributes referring to the same subjects). One could easily envision other potential applications for secure set intersection such as online recommendation services, online dating services, medical databases, and many other applications.

We continue with a survey of the current state of affairs with respect to secure twoparty computation of set intersection.

### 1.1. Background and Related Work

Private equality tests (PET) The simplest non-trivial form of set intersection is where each of the two datasets consists of a single element. A circuit computing this function has $O(\log N)$ gates (where $N$ is the domain size) and therefore can be securely evaluated with this overhead. Specialized protocols for this function, denoted by private equality test (PET) or the socialist millionaires problem, were also suggested in [13,29,47,53] with essentially the same overhead. We note that the solution in [13] also provides fairness. A related problem is that of private authenticated key exchange (PAKE), in which two parties agree on a secure key if both share the same password, taken from a small domain, see [44] and references within.

Private set intersection The following intuitive "solution" is often suggested for private set intersection. Assume the two parties agree on some cryptographic function $H(\cdot)$ such as a one-way hash function or a pseudorandom function. Alice sends to Bob the results of applying $H(\cdot)$ to each element of her input set. Bob then compares these values to the results of applying the function to the elements in his input set. If Bob identifies that, for an input $x$ of his, the value $H(x)$ appears in the list of values sent by Alice, then $x$ must be in the intersection. This simple solution is unfortunately insecure. The reason
is that given Alice's hashed values Bob can test whether an element $x$ appears in her set by searching for $H(x)$ in Alice's hashed set. In particular, when Alice's set comes from a polynomial domain, Bob can recover her entire input set.

Denoting by $n$ the number of elements in each dataset and by $N$ the domain size from which the elements are picked, a trivial solution compares all combinations of items from the two datasets using $n^{2}$ instantiations of a PET protocol (that incurs $O(\log N)$ overhead). The computation of this comparison can be reduced to $O(n \log N)$, while retaining the $O\left(n^{2} \log N\right)$ communication overhead [53]. A circuit with a smaller asymptotic size can be constructed by sorting the inputs of the two parties using a sorting network and then comparing every two adjacent items. The size of a sorting network is $O(n \log n)$ comparators when using the AKS sorting network [6] (which has a large constant factor) or $O\left(n \log ^{2} n\right)$ comparators when using the Batcher sorting network [8], whose overhead induces a more reasonable constant factor. Since our circuit has to handle $N$-bit long inputs, the size of the circuit would be $O(n \log n \log N)$ when using the AKS sorting network or $O\left(n \log ^{2} n \log N\right)$ when using the Batcher sorting network. The $\log N$ factor can be somewhat reduced using hashing.

These solutions can be implemented securely by applying Yao's construction of garbled circuit [65]. For the semi-honest setting, this results in a protocol with communication complexity of the same order as that of the circuit's size times the security parameter [49]. Coping with malicious adversaries is more challenging and requires additional tools for ensuring correctness such as cut-and-choose [50] or zero-knowledge proofs [21,24,42], which inflate the communication/computation costs. In this work, we avoid both costly techniques.

The first concrete construction solving the set intersection problem securely had a cost of $O(n)$ exponentiations [25,35]. However, these constructions were only analyzed in the random oracle model against semi-honest parties and did not provide a full proof with simulation-based security. In [28], Freedman et al. studied set intersection in the standard model and presented a construction for the semi-honest setting, utilizing oblivious polynomial evaluation and balanced allocation hash functions. Their protocol exhibits linear communication and $O(n \log \log n)$ computation overhead (counting modular exponentiations). In addition to their semi-honest protocol, Freedman et al. presented variants of the above protocol for the case where one of the parties is malicious and the other is semi-honest. The protocol for a malicious $P_{1}$ and a semi-honest $P_{2}$ utilizes a cut-andchoose strategy, and therefore, the communication/computation costs are enhanced by a statistical security parameter.

Kissner and Song [46] used polynomials to represent multi-sets. Letting the roots of $Q_{X}(\cdot)$ and $Q_{Y}(\cdot)$ coincide with elements of the multi-sets $X$ and $Y$. They observed that if $r(\cdot), s(\cdot)$ are polynomials chosen at random, then the roots of $r(\cdot) \cdot Q_{X}(\cdot)+s(\cdot) \cdot Q_{Y}(\cdot)$ coincide with high probability with the multi-set $X \cap Y$. Their result is a set intersection protocol for the semi-honest case, where the parties use an additively homomorphic encryption scheme to perform the polynomial multiplication, introducing quadratic computation costs in the set sizes. For the security of the protocol, it is crucial that no party should be able to decrypt on its own. Hence, the secret key must be shared and joint decryption must be deployed. For malicious parties, Kissner and Song [46] introduced the generic zero-knowledge proofs for proving adherence to the prescribed protocol (e.g., zero-knowledge proofs of knowledge for the multiplication of the encrypted $Q_{x}(\cdot)$
with a randomly selected $r(\cdot))$. The costs of these proofs as well as those of setting the shared key for the Paillier scheme are not specified explicitly and can be relatively high. An improved protocol for the malicious setting using more efficient zero-knowledge proofs was presented in [16].

Another tool that can be used for computing set intersection is secure implementation of oblivious pseudorandom function evaluation (OPRF). Namely, having one party holds the keys to a PRF $f_{k}()$ while enabling the other party, who has an input $x$, to compute $f_{k}(x)$ without learning anything else. The idea of using this tool for set intersection was first described in [27]. In [36], Hazay and Lindell presented two OPRF-based protocols for set intersection, one achieving security in the presence of malicious adversaries with one-sided simulatability, whereas the other is secure in the presence of covert adversaries [7]. Having $P_{1}, P_{2}$ hold sets of sizes $m_{1}, m_{2}$, respectively, both protocols in [36] are constant round and incur communication of $O\left(m_{1} \log N+m_{2}\right)$ group elements and computation of $O\left(m_{1} \log N+m_{2}\right)$ modular exponentiations. We note that the protocols in [36] can be made secure in the malicious setup by introducing a secure key selection step for the oblivious PRF and by adding zero-knowledge proofs of knowledge to show correctness at each step. Namely, for proving that the same PRF key is used by party $P_{1}$ for all PRF evaluations and to enable the extraction of the preimages (as a pseudorandom function is not necessarily invertible). While this would preserve the complexity of these protocols asymptotically (in $m_{1}, m_{2}$ ), introducing such proofs would probably make this protocol impractical since there is no efficient known way to design such proofs.

Jarecki and Liu [40] generalized the technique of [36] and presented a very efficient protocol for computing a pseudorandom function with a committed key (informally, this means that the same key is used in all invocations) and showed that it yields an efficient set intersection protocol. The main restriction of this construction is that the input domain size of the PRF must be polynomial in the security parameter (since the proof of security for the set intersection protocol makes use of the ability to exhaustively search over the input domain). Their protocol is secure in the CRS model with a safe RSA modulus placed in the CRS and relies on the Decisional q-Diffie-Hellman Inversion assumption. In a followup work [41], Jarecki and Liu presented a protocol for set intersection that is secure against malicious adversaries under the interactive One-More Gap DiffieHellman assumption in the random oracle model. Their protocol computes an adaptive variant of set intersection for which a receiver is allowed to make adaptive queries, each time revealing whether an item $y_{i}$ belongs to a set $X$. On the other hand, their protocol takes only one round of interaction and its total computational cost is under $3(|X|+|Y|)$ exponentiations (which is better than our protocol presented here).

Dachman-Soled et al. [22] presented a protocol for set intersection in the presence of malicious adversaries without restricting the domain. Their construction uses polynomial evaluation and secret sharing of the inputs. They avoid generic zero knowledge by utilizing the fact that Shamir's secret sharing implies Reed Solomon code. Their protocol incurs communication of $O\left(n k^{2} \log ^{2} n+k n\right)$ group elements and $O\left(n^{2} k \log n+n k^{2} \log ^{2} n\right)$ exponentiations where $k$ is the security parameter.

Finally, Hazay and Nissim [38] investigated protocols in the malicious setting for constructing efficient secure two-party protocols for set intersection and set union. They designed constant-round protocols that exhibit linear communication and a (practically) linear number of exponentiations with simulation-based security. More
explicitly, they get that for sets $X, Y \subseteq\{0,1\}^{\log N}$ of $m_{1}, m_{2}$ elements, respectively, the costs are of sending $O\left(m_{1}+m_{2} \log N\right)$ group elements and the computation of $O\left(m_{1}+m_{2}\left(\log \log m_{1}+\log N\right)\right)$ modular exponentiations. In the heart of these constructions is a technique based on a combination of a perfectly hiding commitment and an oblivious pseudorandom function evaluation protocol with the aim to replace the random oracle used in [28]. Their work does not consider Cuckoo hashing, but our ideas can be applied to their construction as well.

Other variants of the problem were also investigated. Ateniese et al. [3] discussed sizehiding set intersection, where one of the parties can hide the size of its set. Camenisch and Zavrucha [18] investigated the problem of set intersection inputs that are certified by a third party.

Disjointness and set intersection Much attention has been given to bounding the communication complexity of the disjointness function, defined as

$$
\operatorname{Diss}(A, B)=1 \text { if } A \cap B=\phi
$$

It is known $[45,60]$ that if $A, B$ can be arbitrary subsets of $[n]$, then the randomized communication complexity of the disjointness function is $\Theta(n)$. An immediate implication is that computing the intersection of two sets $|X|=|Y|=n$ (over a large enough domain) requires $\Omega(n)$ communication. This follows by a reduction from $\operatorname{DisJ}(\cdot, \cdot)$ to set intersection over domain of size $N \geq 3 n .{ }^{1}$ Let $Z=\left\{z_{1}, \ldots, z_{3 n}\right\}$ be a subset of the domain and let $\phi$ be a one-to-one mapping from $[n]$ to $\left\{z_{1}, \ldots, z_{n}\right\}$. Given a subset $A \subseteq[n]$ let $\hat{X}=\{\phi(a): a \in A\}$ and set $X=\hat{X} \cup\left\{z_{n+1}, \ldots, z_{2 n-|\hat{X}|}\right\}$. Similarly, given $B \subseteq[n]$ let $\hat{Y}=\{\phi(b): b \in B\}$ and set $Y=\hat{Y} \cup\left\{z_{2 n+1}, \ldots, z_{3 n-|\hat{Y}|}\right\}$. Note that $|X \cap Y|=|\hat{X} \cap \hat{Y}|=|A \cap B|$ and hence $\operatorname{DisJ}(A, B)=1$ iff $X \cap Y=\emptyset$. We thus obtain a lower bound of $\Omega(n)$ on the communication needed for computing the intersection even without taking privacy and security into consideration.

Set intersection and oblivious transfer The bit oblivious transfer functionality is defined as

$$
\left(\left(b_{0}, b_{1}\right), \sigma\right) \mapsto\left(\lambda, b_{\sigma}\right),
$$

where $\left(b_{0}, b_{1}\right) \in\{0,1\} \times\{0,1\}$ is the sender's input, $\sigma$ is the receiver's input bit and $\lambda$ denotes the empty string (meaning that the sender has no output). A simple reduction from oblivious transfer (OT) to set intersection shows that implementing set intersection implies implementing oblivious transfer. Given its input $\left(b_{0}, b_{1}\right)$, the sender generates a set of two strings $X=\left\{0\left|b_{0}, 1\right| b_{1}\right\}$. The receiver generates the set $Y=\{\sigma|0, \sigma| 1\}$. The parties run the set intersection protocol on $X, Y$ where at the end the receiver learns $X \cap Y=\left\{\sigma \mid b_{\sigma}\right\}$. By the results of Impagliazzo and Rudich [39], it follows that there is no black-box reduction of oblivious transfer to one-way functions, and therefore, the same holds with regard to reductions of set intersection to one-way functions. Indeed, our set

[^1]intersection protocols use a stronger primitive-an additively homomorphic encryption scheme.

### 1.2. Our Contributions

This paper is an extended and improved version of [28]. We present secure protocols for set intersection in the presence of semi-honest and malicious adversaries with linear costs (with respect to the sets sizes). Our results include the following.

### 1.2.1. Protocols for Computing Set Intersection

These protocols employ a homomorphic encryption scheme and, in particular, the Paillier or ElGamal encryption schemes: (i) a protocol with security against semi-honest adversaries (cf. Sect. 3.1) and (ii) a protocol in the random oracle model with security against malicious adversaries (cf. Sect. 5). Our protocols have simulation-based security (unlike the protocols in [28]), assuming the hardness of DDH/DCR problems. Moreover, they introduce linear (or nearly linear) computation and communication overheads with small constant factors, where the analysis depends on the type of the hash scheme we use. The most efficient result is achieved using a new protocol based on Cuckoo hashing. We further analyze efficiency based on simple hashing and balanced allocation schemes. Our analysis is presented in details in Sect. 3.2.

The semi-honest setting The high-level description of our semi-honest protocol follows by having party $P_{1}$ generating a polynomial $Q(\cdot)$ of degree $m_{1}$, with roots set to the $m_{1}$ elements of $X$, and sending the encrypted coefficients to $P_{2}$ (using a homomorphic encryption). Then, for each element $y \in Y, P_{2}$ replies with the encryption of $r \cdot Q(y)+y$ for a random $r$. This immediately implies that for $y \in X \cap Y$ the result plaintext would be $y$. Otherwise, the plaintext equals a random value that does not leak any information about $y$. Note first that the communication complexity of this protocol is linear in $m_{1}+m_{2}$, yet the work performed by $P_{2}$ is high, as each of the $m_{2}$ oblivious polynomial evaluations includes performing $O\left(m_{1}\right)$ exponentiations totaling in $O\left(m_{1} \cdot m_{2}\right)$ exponentiations.

To save on computational work, we use hashing to map the items into different bins. In that case, the items mapped by $P_{1}$ to a certain bin must only be compared to those mapped by $P_{2}$ to the same bin. Thus, the number of comparisons can be reduced to be in the order of the number of $P_{2}$ 's inputs times the maximum number of items mapped to a bin. In this work, we describe and compare modifications of the basic protocol based on the following different hash schemes: simple hashing, balanced allocations, and Cuckoo hashing. We also provide in Sect. 4 results of experiments with each of these schemes.

The malicious setting Introducing security in the malicious setting raises new concerns. First, the basic scheme introduced above is not secure any longer in the malicious setting since, for instance, a malicious $P_{2}$ can compute the encryption of $r \cdot Q(y)+y^{\prime}$ for distinct $y$ and $y^{\prime}$ (which implies the ciphertext is decrypted to $y^{\prime}$ if and only if $y \in X$ ). In addition, using hashing introduces new attacks since we must ensure that $P_{1}$ computes the small polynomials correctly and that its input is well defined.

Our protocol for the malicious setting avoids the standard solutions that involve zeroknowledge proofs or the cut-and-choose technique for demonstrating correct behavior. Instead, it enables party $P_{1}$ to redo the entire computation supposedly carried out by $P_{2}$ on each element and verify that its outcome is consistent with the messages received from $P_{2}$. In the proof, we show that the probability that $P_{2}$ convinces $P_{1}$ of a correct behavior even though it is not the case is negligible. This technique is implemented in the random oracle model. Importantly, we do not rely on the programmability property of the random oracle which weakens the security random oracle notion that we require. In particular, our simulator only needs to observe the adversary's random oracle queries.

We note that using hash functions to reduce communication cost in the malicious setting introduces new problems as the parties must prove that they used a correct mapping for each element without leaking anything about it. Our solution also deals with this challenge, ensuring that each element is mapped to the correct bin. A more subtle problem that we deal with that was overlooked in prior work, is that with some homomorphic PKEs $P_{1}$ may construct $Q(\cdot)$ such that the evaluation of $r \cdot Q(y)$ (and hence also of $r \cdot Q(y)+y)$ is far from being random in the plaintext space even though $Q(y) \neq 0$ and $r$ is chosen at random. This attack can be carried out with respect to PKEs for which the plaintext space is not a cyclic group of prime order, implying that $r \cdot Q(y)$ may be a random element within a smaller subgroup.

Variants of set intersection Finally, we present in Sect. 3.4 a protocol for computing the cardinality of the intersection. This protocol uses our protocol for computing the set intersection as a main building block. The computation/communication overheads do not change.

### 1.2.2. Experimental Results

We implement and test the different variants of the semi-honest protocol and analyze their overhead (cf. Sect. 4). These experiments are new to this work and were not conducted in [28]. Throughout our experiments, we gain some insights regarding the practicality of the cryptographic primitives we use in our constructions and the relation between the asymptotic and actual overhead of the protocols. Somewhat surprisingly, we get that the variant with the best asymptotic overhead is not necessarily preferable for inputs of reasonable size. In particular, the degree of the polynomials evaluated by the protocol matters less than the number of polynomials that are evaluated. More specifically, recalling that the random hashing construction evaluates a single polynomial of degree $O\left(\log m_{1}\right)$, whereas the balanced allocations construction evaluates two polynomials of degree $O\left(\log \log m_{1}\right)$ and the Cuckoo hashing construction evaluates three polynomials: two linear polynomials and a polynomial of degree 2. Asymptotically, the performance of the constructions based on balanced allocations and Cuckoo hashing is preferable, but since these two constructions use more polynomials than the first construction, their overhead is higher than that of random hashing for the input sizes that we tested.

### 1.2.3. A Roadmap

In Sect. 2, we present definitions and tools that are useful for our constructions. In Sect. 3, we present our first construction for the semi-honest setting and generalizations for two
related problems. In Sect. 4, we present our performance evaluation, and in Sect. 5, we describe our protocol for the malicious setting and its proof.

## 2. Definitions and Tools

Basic notations The security parameter is denoted by $k$, and, although not explicitly specified, input lengths are always assumed to be bounded by some polynomial in $k$. A probabilistic machine is said to run in polynomial time (PPT) if it runs in time that is polynomial in the security parameter $k$. A function $\mu(k)$ is called negligible in $k$ (negligible for short) if for every polynomial $p(\cdot)$ there exists a value $k_{0}=k_{0}(p)$ such that $\mu(k)<\frac{1}{p(k)}$ for all $k>k_{0}$; i.e., $\mu(k)=k^{-\omega(1)}$. Let $X=\{X(k, a)\}_{k \in \mathbb{N}, a \in\{0,1\}^{*}}$ and $Y=\{Y(k, a)\}_{k \in \mathbb{N}, a \in\{0,1\}^{*}}$ be distribution ensembles (over strings of length polynomial in $k$ ). We say that $X$ and $Y$ are computationally indistinguishable, denoted $X \xlongequal{\underline{\mathrm{c}}} Y$, if for every polynomial non-uniform distinguisher $D$ there exists a negligible $\mu(\cdot)$ such that

$$
|\operatorname{Pr}[D(X(k, a))=1]-\operatorname{Pr}[D(Y(k, a))=1]|<\mu(k)
$$

for every $k \in \mathbb{N}$ and $a \in\{0,1\}^{*}$.

### 2.1. Secure Two-Party Computation

We briefly present the standard definition for secure multiparty computation and refer to [34, Chapter 7] for more details and motivating discussions.

Two-party computation A two-party protocol problem is cast by specifying a random process that maps pairs of inputs to pairs of outputs (one for each party). We refer to such a process as a functionality and denote it $f:\{0,1\}^{*} \times\{0,1\}^{*} \rightarrow\{0,1\}^{*} \times\{0,1\}^{*}$, where $f=\left(f_{1}, f_{2}\right)$. That is, for every pair of inputs $(x, y)$, the output vector is a random variable $\left(f_{1}(x, y), f_{2}(x, y)\right)$ ranging over pairs of strings where $P_{1}$ receives $f_{1}(x, y)$ and $P_{2}$ receives $f_{2}(x, y)$. We use the notation $(x, y) \mapsto\left(f_{1}(x, y), f_{2}(x, y)\right)$ to describe a functionality. For example, the oblivious transfer functionality is written $\left(\left(x_{0}, x_{1}\right), \sigma\right) \mapsto\left(\lambda, x_{\sigma}\right)$, where $\left(x_{0}, x_{1}\right)$ is the first party's input, $\sigma$ is the second party's input and $\lambda$ denotes the empty string (meaning that the first party has no output). A special case for a two-party functionality is that of zero-knowledge proof of knowledge for a relation $\mathcal{R}_{\mathrm{ZK}}$. This relation can be defined by the inputs $(x,(x, w))$ that are mapped into $(1, \lambda)$ if $\mathcal{R}_{\mathrm{ZK}}(x, w)=1$, or into $(\perp, \lambda)$ otherwise.

Security of protocols We prove the security of our protocols in the settings of semihonest and malicious computationally bounded adversaries. Loosely speaking, the adversary in the semi-honest setting is assumed to act according to its prescribed actions in the protocol, whereas in the malicious setting it may arbitrarily deviate from the specified protocol. Security is analyzed by comparing what an adversary can do in a real protocol execution to what it can do in an ideal scenario. In the ideal scenario, the computation involves an incorruptible trusted third party to whom the parties send their inputs. The trusted party computes the functionality on the inputs and returns to each party its
respective output. Informally, the protocol is secure if any adversary interacting in the real protocol (i.e., where no trusted third party exists) can do no more harm than what it could do in the ideal scenario. There are technical issues that arise, such as that it may be impossible to achieve fairness or guaranteed output delivery. For example, it is possible for an adversarial party to prevent an honest party from receiving outputs.

### 2.1.1. The Semi-Honest Setting

In this model, the adversary controls one of the parties and follows the protocol specification. However, it may try to learn more information than allowed by looking at the transcript of messages that it received and its internal state. The following definition is according to [34].

Let $f=\left(f_{1}, f_{2}\right)$ be a two-party functionality and let $\pi$ be a two-party protocol for computing $f$. The view of the first party in an execution of $\pi$ on inputs $(x, y)$ is

$$
\mathbf{V i e w}_{\pi, 1}(x, y)=\left(x, r_{1}, m_{1}, \ldots, m_{t}\right),
$$

where $r_{1}$ is the content of the first party's internal random tape and $m_{i}$ represents the $i$ th message that it received. The output of the first party in an execution of $\pi$ on $(x, y)$ is denoted Output ${ }_{\pi, 1}(x, y)$ and can be computed from View ${ }_{\pi, 1}(x, y)$. Similarly, $\operatorname{View}_{\pi, 2}(x, y)\left(y, r_{2}, m_{1}, \ldots, m_{t}\right)$ where $r_{2}$ is second party's randomness and $m_{i}$ is the $i$ th message it received. The output of the second party can be computed from her view and is denoted Output ${ }_{\pi, 2}(x, y)$.

Definition 2.1. Let $f$ and $\pi$ be as above. Protocol $\pi$ is said to securely compute $f$ in the presence of semi-honest adversaries if there exist probabilistic polynomial-time algorithms $\mathcal{S}_{1}$ and $\mathcal{S}_{2}$ such that

$$
\begin{aligned}
& \left(\mathcal{S}_{1}\left(x, f_{1}(k, x, y)\right), f_{2}(k, x, y)\right)_{k \in \mathbb{N}, x, y \in\{0,1\}^{*}} \\
& \quad \stackrel{\text { c }}{=}\left\{\left(\mathbf{V i e w}_{\pi, 1}(k, x, y), \text { Output }_{\pi, 2}(k, x, y)\right)\right\}_{k \in \mathbb{N}, x, y \in\{0,1\}^{*}} \\
& \left(f_{1}(k, x, y), \mathcal{S}_{2}\left(y, f_{2}(k, x, y)\right)\right)_{k \in \mathbb{N}, x, y \in\{0,1\}^{*}} \\
& \stackrel{=}{=}\left\{\left(\mathbf{O u t p u t}_{\pi, 1}(k, x, y),\left(\mathbf{V i e w}_{\pi, 2}(k, x, y)\right)\right)\right\}_{k \in \mathbb{N}, x, y \in\{0,1\}^{*}}
\end{aligned}
$$

where $k$ is the security parameter.

### 2.1.2. The Malicious Setting

Execution in the ideal model In an ideal execution, the parties submit inputs to a trusted party that computes the output. An honest party receives its input for the computation and just directs it to the trusted party, whereas a corrupted party can replace its input with any other value of the same length. Since we do not consider fairness, the trusted party first sends the outputs of the corrupted parties to the adversary and the adversary then decides whether the honest parties would receive their outputs from the trusted party or an abort symbol $\perp$. Let $f$ be a two-party functionality where $f=\left(f_{1}, f_{2}\right)$, let $\mathcal{A}$ be a non-uniform probabilistic polynomial-time machine, and let $I \subset[2]$ be the
set of corrupted parties (either $P_{1}$ is corrupted or $P_{2}$ is corrupted or neither). Then, the ideal execution of $f$ on inputs $(x, y)$, auxiliary input $z$ to $\mathcal{A}$ and security parameter $k$, denoted Ideal ${ }_{f, \mathcal{A}(z), I}(k, x, y)$, is defined as the output pair of the honest party and the adversary $\mathcal{A}$ from the above ideal execution.

Execution in the real model In the real model, there is no trusted third party and the parties interact directly. The adversary $\mathcal{A}$ sends all messages in place of the corrupted party and may follow an arbitrary polynomial-time strategy. The honest parties follow the instructions of the specified protocol $\pi$.

Let $f$ be as above and let $\pi$ be a two-party protocol for computing $f$. Furthermore, let $\mathcal{A}$ be a non-uniform probabilistic polynomial-time machine and let $I$ be the set of corrupted parties. Then, the real execution of $\pi$ on inputs $(x, y)$, auxiliary input $z$ to $\mathcal{A}$ and security parameter $k$, denoted $\operatorname{Real}_{\pi, \mathcal{A}(z), I}(k, x, y)$, is defined as the output vector of the honest parties and the adversary $\mathcal{A}$ from the real execution of $\pi$.

Security as emulation of a real execution in the ideal model Having defined the ideal and real models, we can now define security of protocols. Loosely speaking, the definition asserts that a secure party protocol (in the real model) emulates the ideal model (in which a trusted party exists). This is formulated by saying that adversaries in the ideal model are able to simulate executions of the real model protocol.

Definition 2.2. Let $f$ and $\pi$ be as above. Protocol $\pi$ is said to securely compute $f$ with abort in the presence of malicious adversaries if for every non-uniform probabilistic polynomial-time adversary $\mathcal{A}$ for the real model, there exists a non-uniform probabilistic polynomial-time adversary $\mathcal{S}$ for the ideal model, such that for every $I \subset$ [2],

$$
\left\{\operatorname{Ideal}_{f, \mathcal{S}(z), I}(k, x, y)\right\}_{k \in \mathbb{N}, x, y, z \in\{0,1\}^{*}} \stackrel{\mathrm{c}}{\equiv}\left\{\operatorname{Rea}_{\pi, \mathcal{A}(z), I}(k, x, y)\right\}_{k \in \mathbb{N}, x, y, z \in\{0,1\}^{*}}
$$

where $k$ is the security parameter.
The $f$-hybrid model In our constructions, we will use secure two-party protocols as sub-protocols. A standard way of abstracting out the details of the sub-protocols is to work in a "hybrid model" where the two parties interact with each other (as in the real model) and also use trusted help (as in the ideal model). Specifically, an execution of a protocol $\pi$ that uses a sub-protocol for securely computing some functionality $f$ is modeled as if the parties run $\pi$ and issue "ideal calls" to a trusted party for computing $f$ instead of invoking the protocol for $f$. In these calls to $f$, the parties send inputs to the trusted party, which, upon receiving the inputs from the parties, computes $f$ and sends each party its corresponding output. After receiving these outputs, the protocol $\pi$ continues.

We stress that we use the $f$-hybrid model in a sequential composition, i.e., the parties do not send messages in $\pi$ between the time that they send input to the trusted party and the time that they receive back output. The trusted party may be used a number of times throughout the execution of $\pi$. Each time is independent in the sense that the trusted party does not maintain any state between these calls. We call the regular messages of
$\pi$ that are sent among the parties standard messages and the messages that are sent between parties and the trusted party ideal messages.

Let $f$ be a functionality and let $\pi$ be a two-party protocol that uses ideal calls to a trusted party computing $f$. Let $\mathcal{A}$ be a non-uniform probabilistic polynomialtime machine and let $I$ be the set of corrupted parties. Then, the $f$-hybrid execution of $\pi$ on inputs ( $x, y$ ), auxiliary input $z$ to $\mathcal{A}$ and security parameter $k$, denoted $\operatorname{Hybrid}_{\pi, \mathcal{A}(z), I}^{f}(k, x, y)$, is defined as the output vector of the honest parties and the adversary $\mathcal{A}$ from the hybrid execution of $\pi$ with a trusted party computing $f$.

Let $f$ and $\pi$ be as above, and let $\rho$ be a protocol. Consider the real protocol $\pi^{\rho}$ that is defined as follows. All standard messages of $\pi$ are unchanged. When a party $P_{i}$ is instructed to send an ideal message $\alpha_{i}$ to the trusted party, it begins a real execution of $\rho$ with input $\alpha_{i}$ instead. When this execution of $\rho$ concludes with output $\beta_{i}$, party $P_{i}$ continues with $\pi$ as if $\beta_{i}$ was the output received by the trusted party (i.e., as if it were running in the $f$-hybrid model). Then, the composition theorem of [14] states that if $\rho$ securely computes $f$, then the output distribution of a protocol $\pi$ in a hybrid execution with $f$ is computationally indistinguishable from the output distribution of the real protocol $\pi^{\rho}$. Thus, it suffices to analyze the security of $\pi$ when using ideal calls to $f$; security of the real protocol $\pi^{\rho}$ is derived via this composition theorem.

### 2.2. Hardness Assumptions

Our constructions rely on the following hardness assumptions.
Definition 2.3. ( $D D H$ ) We say that the decisional Diffie-Hellman (DDH) problem is hard relative to $\mathbb{G}=\left\{\mathbb{G}_{k}\right\}$ if for all polynomial-sized circuits $\mathcal{A}=\left\{\mathcal{A}_{k}\right\}$ there exists a negligible function negl such that

$$
\left|\operatorname{Pr}\left[\mathcal{A}\left(\mathbb{G}, q, g, g^{x}, g^{y}, g^{z}\right)=1\right]-\operatorname{Pr}\left[\mathcal{A}\left(\mathbb{G}, q, g, g^{x}, g^{y}, g^{x y}\right)=1\right]\right| \leq \operatorname{negl}(k),
$$

where $q$ is the order of $\mathbb{G}$ and the probabilities are taken over the choices of $g$ and $x, y, z \in \mathbb{Z}_{q}$.

We require the DDH assumption to hold for prime-order groups.

Definition 2.4. (DCR) We say that the decisional composite residuosity $(D C R)$ problem is hard if for all polynomial-sized circuits $\mathcal{A}=\left\{\mathcal{A}_{k}\right\}$ there exists a negligible function negl such that

$$
\begin{aligned}
& \mid \operatorname{Pr}\left[\mathcal{A}(N, z)=1 \mid z=y^{N} \bmod N^{2}\right] \\
& \quad-\operatorname{Pr}\left[\mathcal{A}\left(\mathcal{A}(N, z)=1 \mid z=(1+N)^{r} \cdot y^{N} \bmod N^{2}\right] \mid \leq \operatorname{negl}(k),\right.
\end{aligned}
$$

where $N$ is a random $k$-bit RSA composite, $r$ is chosen at random in $\mathbb{Z}_{N}$ and the probabilities are taken over the choices of $N, y$ and $r$.

### 2.3. Public Key Encryption Schemes

We begin with the definitions of public key encryption and semantic security. We then specify the definition of homomorphic encryption and two encryption schemes that meet this definition.

Definition 2.5. ( $P K E$ ) We say that $\Pi=(G, E, D)$ is a public key encryption scheme if $G, E, D$ are polynomial-time algorithms specified as follows:

- $G$, given a security parameter $n$ (in unary), outputs keys ( $p k, s k$ ), where $p k$ is a public key and $s k$ is a secret key. We denote this by $(p k, s k) \leftarrow G\left(1^{k}\right)$.
- E, given the public key $p k$ and a plaintext message $m$, outputs a ciphertext $c$ encrypting $m$. We denote this by $c \leftarrow E_{p k}(m)$; and when emphasizing the randomness $r$ used for encryption, we denote this by $c \leftarrow E_{p k}(m ; r)$.
- $D$, given the public key $p k$, secret key $s k$ and a ciphertext $c$, outputs a plaintext message $m$ s.t. there exists randomness $r$ for which $c=E_{p k}(m ; r)$ (or $\perp$ if no such message exists). We denote this by $m \leftarrow D_{p k, s k}(c)$.

For a public key encryption scheme $\Pi=(G, E, D)$ and a non-uniform adversary $\mathcal{A}=\left(\mathcal{A}_{1}, \mathcal{A}_{2}\right)$, we consider the following Semantic security game:

$$
\begin{aligned}
& (p k, s k) \leftarrow G\left(1^{k}\right) . \\
& \left(m_{0}, m_{1}, \text { history }\right) \leftarrow \mathcal{A}_{1}(p k), \text { s.t. }\left|m_{0}\right|=\left|m_{1}\right| . \\
& c \leftarrow E_{p k}\left(m_{b}\right), \text { where } b \leftarrow R\{0,1\} . \\
& b^{\prime} \leftarrow \mathcal{A}_{2}(c, \text { history }) . \\
& \mathcal{A} \text { wins if } b^{\prime}=b .
\end{aligned}
$$

Denote by $\operatorname{Adv}_{\Pi, \mathcal{A}}(n)$ the probability that $\mathcal{A}$ wins the semantic security game.
Definition 2.6. (Semantic security) A public key encryption scheme $\Pi=(G, E, D)$ is semantically secure, if for every polynomial non-uniform adversary $\mathcal{A}=\left(\mathcal{A}_{1}, \mathcal{A}_{2}\right)$ there exists a negligible function negl such that $\operatorname{Adv}_{\Pi, \mathcal{A}}(k) \leq \frac{1}{2}+\operatorname{negl}(k)$.

### 2.3.1. Building Blocks: Additively Homomorphic PKE

Intuitively, a public key encryption scheme is additively homomorphic if given two ciphertexts $c_{1}=E_{p k}\left(m_{1} ; r_{1}\right)$ and $c_{2}=E_{p k}\left(m_{2} ; r_{2}\right)$ it is possible to efficiently compute $E_{p k}\left(m_{1}+m_{2} ; r\right)$ with independent $r$ and without the knowledge of the secret decryption key. Clearly, this assumes that the plaintext message space is a group; we actually assume that both the plaintext and ciphertext spaces are groups (with respective group operations + and $\cdot)$. We abuse notation and use $E_{p k}(m)$ to denote the random variable induced by $E_{p k}(m ; r)$ where $r$ is chosen uniformly at random. We have the following formal definition.

Definition 2.7. (Homomorphic PKE) A public key encryption scheme ( $G, E, D$ ) is homomorphic if for all $k$ and all $(p k, s k)$ output by $G\left(1^{k}\right)$, it is possible to define groups $\mathcal{M}, \mathbb{C}$ such that:

- The plaintext space is $\mathcal{M}$, and all ciphertexts output by $E_{p k}(\cdot)$ are elements of $\mathbb{C} .^{2}$
- For every $m_{1}, m_{2} \in \mathcal{M}$, it holds that

$$
\begin{equation*}
\left\{p k, c_{1}=E_{p k}\left(m_{1}\right), c_{1} \cdot E_{p k}\left(m_{2}\right)\right\} \equiv\left\{p k, E_{p k}\left(m_{1}\right), E_{p k}\left(m_{1}+m_{2}\right)\right\} \tag{1}
\end{equation*}
$$

where the group operations are carried out in $\mathbb{C}$ and $\mathcal{M}$, respectively, and the randomness for the distinct ciphertexts is independent.

Note that any such scheme supports a multiplication of a plaintext by a scalar, that can be achieved by computing multiple additions. In particular, this homomorphic property allows oblivious evaluation of any polynomial $Q(x)=\sum_{i} Q_{i} x^{i}$, given its encrypted coefficients $\left\{E_{p k}\left(Q_{i}\right)\right\}_{i}$.

The Paillier PKE An example of an encryption scheme that meets Definition 2.7 is the encryption scheme of Paillier [55]. In this scheme, the public key $p k$ is an RSA composite $N$ and the corresponding secret key $s k$ is $\phi(N)$. Then, given a message $m \in \mathbb{Z}_{N}$, the encryption procedure selects $r \leftarrow_{R} \mathbb{Z}_{N}^{*}$ (in practice $r \leftarrow_{R} \mathbb{Z}_{N}$ ) and computes $E_{p k}(m ; r)=(1+N)^{m} \cdot r^{N} \bmod N^{2}$. Conversely, given a ciphertext $c \in \mathbb{Z}_{N^{2}}^{*}$, the decryption procedure computes

$$
D_{s k}(c)=\frac{\left[\begin{array}{ll}
c^{\phi(N)} & \bmod N^{2}
\end{array}\right]-1}{N} \cdot \phi(N)^{-1} \bmod N .
$$

This scheme is semantically secure, assuming hardness of the decisional composite residuosity problem. Note that the Paillier scheme is homomorphic with respect to addition modulo $N$, as $D_{s k}\left[E_{p k}\left(m_{1}\right) \cdot E_{p k}\left(m_{2}\right)\right]=m_{1}+m_{2} \bmod N$. Furthermore, $D_{s k}\left[\left(E_{p k}(m)\right)^{c} \bmod N^{2}\right]=c m \bmod N$. Hence, given encryptions $\left\{c_{i}\right\}$ of messages $\left\{m_{i}\right\}$ and $p k$, one can compute an encryption of any linear combination $\sum_{i} \alpha_{i} \cdot m_{i}$ $\bmod N$.

The ElGamal PKE Another encryption scheme suitable for our needs is the ElGamal encryption [26]. Namely, let $\mathbb{G}$ be a group generated by $g$ of prime order $q$, in which the decisional Diffie-Hellman (DDH) problem is hard. A public key is then a pair $p k=$ $\langle g, h\rangle$ and the corresponding secret key is $s=\log _{g}(h)$, i.e., $g^{s}=h$. Then, an encryption of a message $m \in \mathbb{Z}_{q}$ is defined by $E_{p k}(m ; r)=\left\langle g^{r}, h^{r} \cdot g^{m}\right\rangle$ where $r$ is picked uniformly at random from $\mathbb{Z}_{q}$. The decryption algorithm follows by outputting $c_{2} / c_{1}^{s}$, decrypting a ciphertext $\left\langle c_{1}, c_{2}\right\rangle$. This scheme does not directly support the necessary homomorphic operations since the decryption yields $g^{m}$ rather than $m$. Fortunately, the lack of "full" decryption is not an issue since $P_{1}$ only needs to distinguish between an encryption of a uniformly picked $m$ and the case where $m \in X$ (i.e., when $m$ is part of its input).

Our experiments, detailed in Sect. 4, show that the ElGamal PKE is much faster than the Paillier PKE with comparable security, and therefore, one might prefer using the ElGamal-based version of our protocols.

[^2]
## 3. Secure Set Intersection in the Semi-Honest Setting

In this section, we present in detail our construction for a protocol realizing $\mathcal{F}_{\cap}$ in the presence of semi-honest adversaries. The main tool used in our construction is oblivious polynomial evaluation, implemented based on an additively homomorphic PKE. We consider the functionality of set intersection where each party's input consists of a set and the size of the other party's input set (the reason for including the size of the other party's set in each party's input is to model the fact that the protocol leaks the sizes of the input sets). More formally:

Definition 3.1. Let $X$ and $Y$ be subsets of a predetermined domain, ${ }^{3}$ the functionality $\mathcal{F}_{\cap}$ is:

$$
\left(\left(X, m_{2}\right),\left(Y, m_{1}\right)\right) \mapsto\left\{\begin{array}{cl}
(X \cap Y, \lambda) & \text { if }|X|=m_{1} \text { and }|Y|=m_{2} \\
(\lambda, \lambda) & \text { otherwise }
\end{array}\right.
$$

We continue with a high-level description of the basic semi-honest protocol as given in [28]. We then discuss techniques to improve the computational overhead (Sect. 3.2) and provide a formal description of our protocol together with a detailed simulationbased security proof (Sect. 3.3). In Sect. 3.4, we consider two variants of the set intersection problem and discuss solutions for these problems.

### 3.1. A High-Level Description

Recall that $P_{1}$ has $m_{1}$ elements and $P_{2}$ has $m_{2}$ elements. The basic protocol works as follows:

1. Party $P_{1}$ chooses a pair of encryption/decryption keys $(p k, s k) \leftarrow G\left(1^{k}\right)$ for a homomorphic PKE ( $G, E, D$ ) and sends $p k$ to $P_{2}$.
2. $P_{1}$ computes the coefficients of a polynomial $Q(\cdot)$ of degree $m_{1}$, whose roots are set to be the $m_{1}$ elements of $X . P_{1}$ sends the encrypted coefficients of $Q(\cdot)$ to $P_{2}$.
3. For each element $y \in Y$ (in a random order), party $P_{2}$ chooses $r$ at random (from the appropriate plaintext space $\mathcal{M}$ ) and uses the homomorphic properties of the encryption scheme to compute an encryption of $r \cdot Q(y)+y . P_{2}$ sends the encrypted values to $P_{1}$.
4. Upon receiving these encrypted values, $P_{1}$ extracts $X \cap Y$ by decrypting each value and then checking if the result is in $X$.

Security argument (informal) This argument is split into two parts. First note that if $y \in X \cap Y$, then by the construction of the polynomial $Q(\cdot)$ we get that $r \cdot Q(y)+y=$ $r \cdot 0+y=y$. On the other hand, if $y \notin X \cap Y$, we require that the product $r \cdot Q(y)$ corresponds to a random value within the plaintext group so that it reveals no information

[^3]about $y$ and (with high probability) is not in $X .{ }^{4}$ This requirement is obtained almost immediately when using the ElGamal PKE since the plaintext space (excluding zero), $\mathbb{Z}_{q}^{*}$, is a multiplicative group. Therefore, for any element $m \in \mathbb{Z}_{q}^{*}$, the probability that $r \cdot Q(y)=m$ is $1 /(q-1)$ since $r=m / Q(y)$, given that $Q(y) \neq 0$. This, however, is not immediately true when using the Paillier PKE with plaintext space $\mathbb{Z}_{N}$ since $N$ is not a prime. Fortunately, the fraction of elements outside the multiplicative subgroup $\mathbb{Z}_{N}^{*}$ is negligible in $k$ (furthermore, finding an element in $\mathbb{Z}_{N} \backslash \mathbb{Z}_{N}^{*}$ when the secret key is unknown amounts to factoring $N$ ). We thus get that for any element $y \notin X \cap Y$ the probability that $Q(y) \notin \mathbb{Z}_{N}^{*}$ is negligible in $k$.

More generally, we formalize the requirements from the homomorphic PKE as follows:

Definition 3.2. We say that homomorphic PKE is good for plaintext space if the following holds:

1. The plaintext space $\mathcal{M}$ contains the input domain space $\{0,1\}^{p(k)}$.
2. There exists a multiplicative group $\mathbb{G}$ for which all but a negligible fraction of $\mathcal{M}$ is in $\mathbb{G}$. Specifically, for any $X, Y \subseteq\{0,1\}^{p(k)}$, the probability that $Q(y) \notin \mathbb{G}$ for all $y \notin X \cap Y$ is negligible, where the probability is taken over the choice of the parameters for $\mathbb{G}$.

Efficiency The communication complexity of this simple scheme is linear in $m_{1}$ and $m_{2}$, as $m_{1}+1$ ciphertexts are sent from $P_{1}$ to $P_{2}$ (these are the encrypted coefficients of $Q(\cdot)$ ), and $m_{2}$ encrypted values are sent from $P_{2}$ to $P_{1}$ (i.e., $Q(y)$ for every $y \in Y$ ). However, the work performed by $P_{2}$ is relatively high as each of the $m_{2}$ oblivious polynomial evaluations includes performing $O\left(m_{1}\right)$ exponentiations, totaling in $O\left(m_{1}\right.$. $m_{2}$ ) exponentiations.

A first improvement in the computational overhead can be obtained by applying Horner's rule to the evaluation of $Q(\cdot)$. Any polynomial $P(y)=\alpha_{0}+\alpha_{1} y+\alpha_{2} y^{2}+$ $\cdots+\alpha_{m_{1}} y^{m_{1}}$ can be evaluated "from the inside out" as follows,

$$
\alpha_{0}+y\left(\alpha_{1}+y\left(\alpha_{2}+y\left(\alpha_{3}+\cdots y\left(\alpha_{m_{1}-1}+y \alpha_{m_{1}}\right) \cdots\right)\right)\right) .
$$

This method of computation employs $m_{1}$ multiplications by $y$, compared with the straightforward polynomial computation method that multiplies the coefficients with the values $y, y^{2}, \ldots, y^{m_{1}}$. The improvement is significant when the length of $y$ is much shorter than the length of a full exponent. Recall that with the textbook algorithm for modular exponentiation the computational overhead is linear in the length of the exponent. The gain may be significant even when fine-tuned exponentiation algorithms such as Montgomery's method or Karatsuba's technique are used.

For example, the length of the exponent is typically at least 1,024 bits in the case of Paillier and could be as short as 160 bits in the case of ElGamal. There are many applications in which the input $y$ can be rather short, e.g., if $y$ is a social security number

[^4]then it can be encoded using only 30 bits. When the domain from which the inputs $y$ are taken is large, standard hashing techniques can be used as long as collision probability is negligible. Based on the birthday paradox, it is sufficient to hash the values to strings of length about $2 \max \left(\log m_{1}, \log m_{2}\right)$ bits. For the textbook exponentiation algorithm, the improvement is by a factor of $k / 2 \max \left(\log m_{1}, \log m_{2}\right)$, where $k$ is the length of the exponent. For $m_{1}, m_{2} \approx 100,000$, this factor is about 30 for Paillier with $k=1,024$ or about 5 for ElGamal with $k=160$.

### 3.2. Using Hashing to Reduce the Computational Overhead

The main computational overhead of the basic protocol is the work of $P_{2}$, which essentially has to do $m_{1} \cdot m_{2}$ comparisons, in order to compare each of its inputs to each of the inputs of $P_{1}$. This overhead can be reduced using hashing, if both parties use the same hash scheme to map their respective items into different bins. In that case, the items mapped by $P_{1}$ to a certain bin must only be compared to those mapped by $P_{2}$ to the same bin. Thus, the number of comparisons can be reduced to be in the order of the number of $P_{2}$ 's inputs times the maximum number of items mapped to a bin. (Of course, care must be taken to ensure that the result of the hashing does not reveal information about the inputs.) We describe and compare modifications of the basic protocol based on the following different hash schemes: simple hashing, balanced allocations, and Cuckoo hashing. We also provide in Sect. 4 results of experiments with each of these schemes.

Using simple hashing Suppose that the items are hashed into one of $B$ bins as follows. Let $h$ be a randomly chosen hash function mapping elements into bins numbered $1, \ldots, B$. It is well known that if the hash function $h$ maps $m_{1}$ items to random bins, then, if $m_{1} \geq B \log B$, each bin contains with high probability at most $M=\frac{m_{1}}{B}+\sqrt{\frac{m_{1} \log B}{B}}$ (see, e.g., $[61,63]$ ). Setting $B=m_{1} / \log m_{1}$ and applying the Chernoff bound shows that $M=O\left(\log m_{1}\right)$ except with probability $\left(m_{1}\right)^{-s}$, where $s$ is a constant that depends on the exact value of $M .{ }^{5}$

In the protocol, the hash function $h$ will be known to both parties. We would therefore like to hide from each party how many of the inputs of the other party are mapped by $h$ to each bin. Thus, the modified protocol works by having $P_{1}$ define a polynomial of degree $M$ for each bin by fixing its mapped elements to be the set of roots. In addition, $P_{1}$ adds the root 0 sufficiently many times, so that the total degree of the polynomial is $M$. That is, if $h$ maps $\ell$ items to the bin, $P_{1}$ first defines a polynomial whose roots are these $\ell$ items and then multiplies it by $x^{M-\ell}$. (We assume that 0 is not a valid input.) The process results in $B$ polynomials, all of degree $M$, with exactly $m_{1}$ nonzero roots. Note that although only a few of the bins have $M$ items, we have to set all polynomials to this degree and pay the associated overhead, in order to hide the exact number of items that are mapped to every bin.

[^5]Finally, $P_{1}$ sends to $P_{2}$ the encrypted coefficients of the polynomials and the mapping $h$ from elements to bins. (For our purposes, it is sufficient that the mapping is selected pseudorandomly, either jointly or by either party.) For every $y \in Y, P_{2}$ finds the bin into which $y$ is mapped and evaluates the polynomial of this bin. It proceeds as before and responds to $P_{1}$ with the encryptions of $r P(y)+y$ for all $y$. Note that the communication complexity is not affected, as $P_{1}$ sends $B M=O\left(m_{1}\right)$ items and $p_{2}$ replies with $m_{2}$ values. The main gain is that $P_{2}$ now has to perform only $m_{2} M=m_{2} O\left(\log m_{1}\right)$ exponentiations (rather than $m_{1} m_{2}$ exponentiations if no hashing is used).

Using balanced allocations To save on the computational work even more, we introduce a balanced allocation scheme for randomly mapping elements into bins. Loosely speaking, we use the balanced allocation scheme of [1] where elements are inserted into $B$ bins as follows. Let $h_{0}, h_{1}$ be two randomly chosen hash functions mapping elements into bins numbered $1, \ldots, B$. An element $x$ is inserted into the less occupied bin from $\left\{h_{0}(x), h_{1}(x)\right\}$, where ties are broken arbitrarily. If $m_{1}$ elements are inserted, then, except with negligible probability over the choice of the hash functions $h_{0}, h_{1}$, the maximum number of elements allocated to any single bin is at most $M=O\left(m_{1} / B+\log \log B\right)$. (The exact probability is described in a note below.) Setting $B=\frac{m_{1}}{\log \log m_{1}}$, we get that $M=O\left(\log \log m_{1}\right) \cdot{ }^{6}$ Note that in this case the degree of the polynomials and the associated overhead are much smaller than the $\log m_{1}$ degree that was set when simple hashing was used.

Then, upon receiving the encrypted polynomials, party $P_{2}$ obliviously evaluates, for each of its inputs $y$, the encryption of $r_{0} \cdot Q_{h_{0}(y)}(y)+y$ and $r_{1} \cdot Q_{h_{1}(y)}(y)+y$ for each of the two bins $h_{0}(y), h_{1}(y)$ in which $y$ can be allocated. (Note that here $P_{2}$ must evaluate two polynomials for each of its inputs since it has no way of knowing what is the right bin.) Finally, $P_{1}$ decrypts both evaluations and performs the same check as in the highlevel description above, twice. Setting $B=\frac{m_{1}}{\log \log m_{1}}$ and $M=O\left(\log \log m_{1}\right)$, we get that the communication complexity is not affected (neglecting constant factors), as $P_{1}$ now sends $B M=O\left(m_{1}\right)$ encrypted values and $P_{2}$ replies with $2 m_{2}$ encrypted values. There is, however, a reduction in the work performed by $P_{2}$, as each of the oblivious polynomial evaluations amounts to performing $O(M)$ exponentiations, and so $P_{2}$ now performs only $2 m_{2} M=2 m_{2} O$ ( $\left.\log \log m_{1}\right)$ exponentiations overall.

Finally, one may worry about the case that $P_{1}$ is unlucky in its choice of hash functions such that more than $M$ items are mapped to some bin. The bound of [1] only guarantees that this happens with probability $o(1)$. However, Broder and Mitzenmacher [12] have shown that asymptotically, when we map $n$ items into $n$ bins, the number of bins with $i$ or more items falls approximately like $2^{-2.6^{i}}$. Formally, this means that if $M=\omega(\log \log n)$, then except with negligible probability no bin will be of size greater than $M$. Practically, this means that a bound of $M=5$ suffices with probability $1-10^{-35}$. The authors also provide experimental results that confirm the asymptotic bound for the case of $n=32,000$.

[^6]Using Cuckoo hash Cuckoo hash [58] is a multiple-choice hashing scheme with evictions. In its simplest form, $n$ items are mapped to $2(1+\varepsilon) n$ bins (where $\varepsilon$ is a small constant that affects the error probability), using two hash functions $h_{0}, h_{1}$. When item $x$ is inserted, then if either bin $h_{0}(x)$ or bin $h_{1}(x)$ is free, $x$ is mapped to the free bin. Otherwise, we put $x$ in bin $h_{0}(x)$, evict the item $x^{\prime}$ that is already in that location, and try to move $x^{\prime}$ to the location defined by the other hash function (i.e., if $h_{0}(x)=h_{0}\left(x^{\prime}\right)$, then we try to move $x^{\prime}$ to $h_{1}\left(x^{\prime}\right)$ ). If that location is occupied as well, we evict the item that is located there and try to relocate it, and so on. A version of this algorithm described in [58] continues with this process, until there is a series of $O(\log n)$ evictions, in which case it chooses new hash functions $h_{0}, h_{1}$ and tries to re-insert all items. It was shown in [58] that the expected time, i.e., the expected number of insertion trials, needed by the scheme to insert all items is $O(n)$.

The use of Cuckoo hash for our purposes is appealing, since at most a single item is mapped to every bin. Therefore, the number of comparisons can be reduced to be linear in the uninput size. A major obstacle, though, is that there is a pretty high probability, of $1 / n$, that a specific input set of $n$ elements cannot be inserted in the table using a specific pair of hash functions $h_{0}, h_{1}$, and then, a new pair of hash functions must be chosen. Therefore $P_{2}$, which knows the hash functions and knows that they were chosen at random from the set of pairs of functions which do not cause an abort for $P_{1}$ 's input, can identify with probability of about $1 / n$ whether a certain potential input set of $P_{1}$ cannot be the actual input of $P_{1}$ in a specific run of the protocol. This problem occurs even if $P_{1}$ rearranges its input element in a random order.

One approach to reducing the effect of this problem is to use the "Cuckoo hashing with a stash" solution of [43]: In addition to the hash table, this solution keeps a small amount of additional memory, namely a stash of $s$ items. If the insertion algorithm encounters an infinite cycle of evictions, then an element in that cycle is moved to the stash. When we look for an item $x$ in the table, we search for it in locations $h_{0}(x)$ and $h_{1}(x)$ as well as in the stash. It was shown in [43] that, for any constant $s$, using a stash of size $s$ fails with probability $O\left(n^{-s}\right)$ (taken over the choice of the hash functions). Therefore, if we change the algorithm so that it aborts if the original choice of hash functions results in more than $s$ items being moved to the stash, then the algorithm aborts with probability of at most $O\left(n^{-s}\right)$. Consequently, $P_{2}$ can identify with probability $O\left(n^{-s}\right)$ whether a specific potential input of $P_{1}$ does not agree with the hash functions $h_{0}, h_{1}$. This is a low (albeit not negligible) probability which we further characterize experimentally in Sect. 4 for various settings of hash table size $(1+\epsilon), n$ and stash size $s$.

When the protocol is implemented using Cuckoo hash, then when $P_{2}$ processes an input $x$ it must evaluate the polynomials of the two bins to which $x$ might have been mapped (bins $h_{0}(x)$ and $h_{1}(x)$ ), as well as the polynomial representing the stash. The advantage is that the polynomials of the two bins are linear and the polynomial representing the stash is of small degree. This should be compared to the balanced allocation solution in which the polynomials representing the bins are of degree $\log \log m_{1}$.

### 3.3. The Detailed Protocol

We are now ready to formally present our protocol. We describe the protocol based on balanced allocations, since for that hashing scheme there is a negligibly small bound on
the failure probability. We remind the reader that our construction is secure with respect to both Paillier and ElGamal and any homomorphic PKE that meets Definition 3.2. We describe our construction using a general notation for the homomorphic encryption scheme. In Sect. 4, we describe the results of experiments with each of the hashing schemes.

Protocol 1. ( $\pi_{\cap}^{\mathrm{SH}}-$ secure set intersection in the semi-honest model):

- Inputs: The input of $P_{1}$ is $m_{2}$ and a set $X \subseteq\{0,1\}^{p(k)}$ containing $m_{1}$ items; the input of $P_{2}$ is $m_{1}$ and a set $Y \subseteq\{0,1\}^{p(k)}$ containing $m_{2}$ items (hence, both parties know $m_{1}$ and $\left.m_{2}\right)$. Recall from definition 3.2 that $\{0,1\}^{p(k)}$ is contained in the plaintext space of a homomorphic PKE with a good plaintext space.
- Auxiliary inputs: A security parameter $1^{k}$.
- Convention: Both parties check every received ciphertext for validity and abort if an invalid ciphertext is received.
- The protocol:

1. Key setup: $P_{1}$ chooses $(p k, s k) \leftarrow G\left(1^{k}\right)$ and sends $p k$ to $P_{2}$.
2. Setting the balanced allocation scheme: $P_{1}$ computes the parameters $B, M$ for the scheme and chooses two randomly chosen hash functions $h_{0}, h_{1}$ : $\{0,1\}^{p(k)} \rightarrow[B] .{ }^{7}$ It sends $B, M, h_{0}, h_{1}$ to $P_{2}$.
3. Creating polynomials for the set $X$ : For every $x \in X, P_{1}$ maps $x$ into the less occupied bin from $\left\{h_{0}(x), h_{1}(x)\right\}$ (ties broken arbitrarily). Let $\mathcal{B}_{i}$ denote the set of elements mapped into bin $i$ and let $Q_{i}(x) \stackrel{\text { def }}{=} \sum_{j=0}^{M} Q_{i, j} \cdot x^{j}$ denote a polynomial with the set of roots $\mathcal{B}_{i}$ (if $\mathcal{B}_{i}=\emptyset$, then $P_{1}$ sets $Q_{i}(x)=1$ ). If $\left|\mathcal{B}_{i}\right|<M$, then $P_{1}$ multiplies $Q_{i}(\cdot)$ by $x^{M-\left|B_{i}\right|}$.
$P_{1}$ encrypts the coefficients of the polynomials, setting $q_{i, j}=E_{p k}\left(Q_{i, j} ; r_{i, j}\right)$. It sends the encrypted coefficients to $P_{2}$.
4. Substituting in the polynomials: Let $y_{1}, \ldots, y_{m_{2}}$ be a random ordering of the elements of set $Y . P_{2}$ does the following for all $\alpha \in\left\{1, \ldots, m_{2}\right\}$ :

- It sets $\hat{h}_{0}=h_{0}\left(y_{\alpha}\right), \hat{h}_{1}=h_{1}\left(y_{\alpha}\right)$ (these values denote the bin number). Then, it chooses random plaintexts $r_{0}, r_{1}$ in the domain of $E_{p k}()$ and uses the homomorphic properties of the encryption scheme to evaluate $e_{\alpha}^{0}=E_{p k}\left(r_{0} \cdot Q_{\hat{h}_{0}}\left(y_{\alpha}\right)+y_{\alpha}\right)$ and $e_{\alpha}^{1}=E_{p k}\left(r_{1} \cdot Q_{\hat{h}_{1}}\left(y_{\alpha}\right)+y_{\alpha}\right)$. Finally, it sends $e_{\alpha}^{0}, e_{\alpha}^{1}$ to $P_{1}$.

5. Computing the intersection: For each received $e_{\alpha}^{0}, e_{\alpha}^{1}$, party $P_{1}$ checks iffor some $x \in X, x=D_{s k}\left(e_{\alpha}^{0}\right)$ or $x=D_{s k}\left(e_{\alpha}^{1}\right)$. In this case, $P_{1}$ records $x$ as part of its output.

Correctness We first note that when both parties follow the protocol, $P_{1}$ outputs $X \cap Y$ with probability negligibly close to 1 : (i) For elements $x \in X \cap Y$, there exists $y_{\alpha} \in Y$ that zeros at least one of the polynomials $Q_{h_{0}\left(y_{\alpha}\right)}(\cdot), Q_{h_{1}\left(y_{\alpha}\right)}(\cdot)$. Hence, getting the

[^7]messages sent for $y_{\alpha}, P_{1}$ records $y_{\alpha}$. (ii) For elements $x \in X \backslash Y, P_{1}$ records $x$ only if either $x=D_{s k}\left(e_{\alpha}^{0}\right)$ or $x=D_{s k}\left(e_{\alpha}^{1}\right)$ which occurs with negligible probability due to the randomness of $r_{0}$ and $r_{1}$ and the requirement that the homomorphic PKE is good for plaintext space.

Theorem 3.3. Assume that $(G, E, D)$ is a semantically secure homomorphic PKE that is good for plaintext space (cf. Definition 3.2). Then, Protocol 1 securely computes $\mathcal{F}_{\cap}$ in the presence of semi-honest adversaries (cf. Definition 2.1).

## Proof. We consider two corruption cases.

Simulating the view of party $\quad P_{1}$ The simulation is based on the fact that messages received by $P_{1}$ are encryptions either of elements of the intersection set or are random elements of the message space of the encryption scheme. We construct a simulator $\mathcal{S}$ as follows:

1. $\mathcal{S}$ is given $X, m_{2}, X \cap Y$ and $1^{k}$ and sets $m_{1}=|X|$.
2. $\mathcal{S}$ receives from $P_{1}$ its public key $p k$.
3. $\mathcal{S}$ receives from $P_{1}$ the parameters $B, M$ and the seeds for the two chosen hash functions $h_{0}, h_{1}:\{0,1\}^{p(k)} \rightarrow[B]$ used in the balanced allocation scheme.
4. $\mathcal{S}$ receives from $P_{1}$ the encrypted polynomials $\left\{q_{i, j}\right\}_{i \in\{1, \ldots, B\}, j \in\{0, \ldots, M\}}$.
5. $\mathcal{S}$ completes $X \cap Y$ to size $m_{2}$ by adding random elements from $\{0,1\}^{p(k)}$. Let $\tilde{Y}$ denote this recorded set.
6. $\mathcal{S}$ now plays the role of the honest $P_{2}$ for the rest of the execution, using $\tilde{Y}$ as the input for $P_{2}$.
7. $\mathcal{S}$ outputs the view of $P_{1}$.

We claim that the joint distributions of $P_{1}$ 's view and $P_{2}$ 's output are statistically close in the real and simulated executions. The potential difference between these executions is in Step 5 of the simulation, where $\mathcal{S}$ plays the role of the honest $P_{2}$ with $\tilde{Y}$ instead of $Y$. The difference is that the random elements in $\tilde{Y}^{\prime}=\tilde{Y} \backslash(Y \cap X)$ are used in the simulated execution where members of $Y^{\prime}=Y \backslash X=Y \backslash(Y \cap X)$ are used in the real execution.

Let $y \in Y^{\prime}$ (respectively, $\tilde{y} \in \tilde{Y}^{\prime}$ ) be elements considered in the real (respectively, simulated) execution. Then, $P_{1}$ receives the encryptions $e^{0}=E_{p k}\left(r_{0} \cdot Q_{h_{0}(y)}(y)+y\right)$ and $e^{1}=E_{p k}\left(r_{1} \cdot Q_{h_{1}(y)}(y)+y\right)$ in the real execution and the encryptions $\tilde{e}^{0}=$ $E_{p k}\left(\tilde{r}_{0} \cdot Q_{h_{0}(\tilde{y})}(\tilde{y})+\tilde{y}\right)$ and $\tilde{e}^{1}=E_{p k}\left(\tilde{r}_{1} \cdot Q_{h_{1}(\tilde{y})}(\tilde{y})+\tilde{y}\right)$ in the simulated execution. Note first that due to the requirement that the input domain size is super-polynomial in $k$, the probability that elements from $\tilde{Y}$ are in $X$ as well is negligible in $k$. Moreover, $Q_{h_{0}(y)}(y), Q_{h_{1}(y)}(y), Q_{h_{0}(\tilde{y})}(\tilde{y}), Q_{h_{1}(\tilde{y})}(\tilde{y}) \in \mathbb{G}$ with overwhelming probability (for $\mathbb{G} \subseteq \mathcal{M}$ multiplicative group as specified in Definition 3.2). We thus get that all the encrypted plaintexts sent to $P_{1}$ (by both $P_{2}$ and $\mathcal{S}$ ) are uniformly distributed in $\mathbb{G}$ and therefore $P_{1}$ cannot distinguish between the two cases since they are identical with overwhelming probability.

Simulating the view of party $\quad P_{2}$ The simulation is based on the fact that as $P_{2}$ only receives encrypted values it cannot detect whether these are encryption of values sent in
the real protocol or encryptions of arbitrary messages sent by the simulator. Construct a simulator $\mathcal{S}$ as follows:

1. $\mathcal{S}$ is given $Y, m_{1}, X \cap Y$ and $1^{k}$ and invokes $P_{2}$ on these inputs. $\mathcal{S}$ sets $m_{2}=|Y|$.
2. $\mathcal{S}$ chooses $(p k, s k) \leftarrow G\left(1^{k}\right)$ and sends $p k$ to $P_{2}$.
3. $\mathcal{S}$ computes the parameters $B, M$ for the balanced allocation scheme and chooses random seeds for the hash functions $h_{0}, h_{1}$. These are then sent to $P_{2}$.
4. $\mathcal{S}$ sends to $P_{2}, B M$ encryptions of the value 0 , under the key $p k$ (i.e., encryptions of the coefficients of $B$ polynomials which are identically equal to 0 ). Each encryption is done with fresh randomness.
5. $\mathcal{S}$ completes the execution and outputs whatever $P_{2}$ does.

In the following, we define a sequence of hybrid games and denote by the random variable $\mathbf{H}_{\ell}(k, X, Y)$ (for a fixed $k$ ) the joint output of $P_{2}$ and $P_{1}$ in hybrid game $\mathbf{H}_{\ell}$.
Game $\mathbf{H}_{0}$ : The simulated execution.
Game $\mathbf{H}_{1}$ : The simulator $\mathcal{S}_{1}$ acts identically to $\mathcal{S}$ except that in Step 2 of the simulation $S_{1}$ does not get to know the secret key $s k$. The random variables $\mathbf{H}_{0}(k, X, Y)$ and $\mathbf{H}_{1}(k, X, Y)$ are identical as $\mathcal{S}$ ignores $s k$.
Game $\mathbf{H}_{2}$ : In this game, there is no trusted party and no honest $P_{1}$. Instead, the simulator $\mathcal{S}_{2}$ is given as input $P_{1}$ 's real input $X . \mathcal{S}_{2}$ works exactly like $\mathcal{S}_{1}$ except that instead of sending encryptions of zero polynomials it computes the polynomials as in the real execution using the set $X$.

We prove that the random variables $\mathbf{H}_{1}(k, X, Y)$ and $\mathbf{H}_{2}(k, X, Y)$ are computationally indistinguishable via a reduction to the security of $(G, E, D) .{ }^{8}$ Assume, for contradiction, the existence of a distinguisher circuit $D$ for the distribution induced by games $\mathbf{H}_{1}(k, X, Y)$ and $\mathbf{H}_{2}(k, X, Y)$ and choose hash functions $h_{0}, h_{1}$ that maximize $D$ 's distinguishing advantage. Construct a distinguisher circuit $D_{\mathrm{E}}$ that distinguishes between the encryptions of two sets of messages: (i) coefficients of zero polynomials as constructed in game $\mathbf{H}_{1}$ and (ii) coefficients of polynomials corresponding to $X$ and hash functions $h_{0}, h_{1}$ as constructed in game $\mathbf{H}_{2} . D_{\mathrm{E}}$ receives a public key $p k$ and vector of encryptions $\bar{c}$ under $p k$ and works exactly like $\mathcal{S}_{2}$ except that it uses the public key given to it as input in the semantic security game, instead of generating it by itself, and forwards to $P_{2}$ the vector of encryptions $\bar{c}$ instead of the encrypted polynomials. Note that the output distribution generated by $D_{\mathrm{E}}$ is either identical to $\mathbf{H}_{1}(k, X, Y)$ or to $\mathbf{H}_{2}(k, X, Y)$ (conditioned on our choice of $h_{0}, h_{1}$ ), and hence, the existence of $D$ contradicts the indistinguishability of ciphertexts of the encryption scheme ( $G, E, D$ ).
Game $\mathbf{H}_{3}$ : The simulator $\mathcal{S}_{3}$ acts identically to $\mathcal{S}_{2}$ except that $\mathcal{S}_{3}$ is given $s k$. The random variables $\mathbf{H}_{2}(k, X, Y)$ and $\mathbf{H}_{3}(k, X, Y)$ are identical as $\mathcal{S}_{2}$ and $\mathcal{S}_{3}$ do not use $s k$.

Note that game $\mathbf{H}_{3}$ is identical to the real execution with respect to the adversary's view.

[^8]
### 3.4. Computing Set Intersection Cardinality

In a protocol for computing the cardinality of the set intersection, $P_{1}$ learns $|X \cap Y|$ but not the actual elements of this set. In order to compute this functionality, $P_{2}$ needs only slightly change his behavior from that in Protocol 1. That is, instead of computing the encryptions of $r_{0} \cdot Q_{\hat{h}_{0}}\left(y_{\alpha}\right)+y_{\alpha}$ and $r_{1} \cdot Q_{\hat{h}_{1}}\left(y_{\alpha}\right)+y_{\alpha}, P_{2}$ now only encodes $r_{0} \cdot Q_{\hat{h}_{0}}\left(y_{\alpha}\right)$ and $r_{1} \cdot Q_{\hat{h}_{1}}\left(y_{\alpha}\right)$. Then, in Step 5 of the protocol, $P_{1}$ counts the number of ciphertexts received from $P_{2}$ that are decrypted to zero and locally outputs this number. The proof of security for this protocol trivially follows from that of Protocol 1.

## 4. Performance Evaluation

In order to give concrete performance numbers to our protocol, we provide in this section performance benchmarks for an implementation of the private matching protocol described in Sect. 3.1. As described below, the benchmarks show the advantage of the hashing-based approaches over the basic approach which uses no hashing. On the other hand, the performance benefit of using the solutions based on balanced allocations and on Cuckoo hashing, over that of a single hash function, does not become noticeable until the input sets are pretty large, due to additional factors that affect the overhead.

Our implementation uses the cryptographic libraries that are part of the SFS toolkit [51] (although we wrote the implementation of the Paillier cryptosystem), which in turn uses the GNU Multi-Precision library 4.3.0 [32] for large arithmetic operations. Performance benchmarks were performed on a Sun X2200 M2 server running Fedora core 12 in 64-bit mode, using a single core of an AMD Opteron 2376 processor ( 2.3 GHz ).

Table 1 shows the performance of the encryption schemes of Paillier [55] in a fast decryption mode, as described in Section 6 of [55], and ElGamal, using the variant described in Sect. 2.3.1. (Recall that in that variant of ElGamal we encrypt values in the exponent, to support the homomorphic addition property. This makes decryption harder, but that issue is irrelevant to our protocol since we only need to check whether the decrypted plaintext is equal to 0 .) We used key and modulus sizes that offered comparable security: a 1,024-bit value for $N$ in the Paillier cryptosystem (and thus a 2,048-bit modulus for $N^{2}$ ) and a 160-bit key and 1,024-bit prime modulus for ElGamal. The numbers correspond to the mean value (in microseconds) of 100 runs across five different keys. The column HAdd corresponds to the cost of performing a homomorphic addition of plaintexts, which is akin to modular multiplication. Due to its superior performance, we use ElGamal to instantiate our set intersection protocol in this section's subsequent benchmarks. We also note that the ElGamal encryption can be implemented even more efficiently based on elliptic curve cryptography.

Table 1. Average speed ( $\mu \mathrm{s}$ ) of public key operations with a 1,024-bit El Gamal modulus and a 2,048-bit Paillier modulus.

|  | Encrypt | Decrypt | HAdd |
| :--- | :---: | :--- | :---: |
| Paillier | 3,742 | 273 | 4 |
| El Gamal | 266 | 155 | 2 |



Fig. 1. In Cuckoo hashing, fraction of randomized trials that exceed a given stash size, for varying hash table sizes $2 \cdot(1+\epsilon) \cdot n$. Solid vertical line at 0.02 corresponds to the $\epsilon$ used in set intersection experiments.

For the Cuckoo hash construction, we need to instantiate the protocol with appropriate parameters for the stash size $s$ and table size $2 \cdot(1+\epsilon) \cdot n$. As stated earlier, it was shown in [43] that using a stash of size of size $s$ fails with probability $O\left(n^{-s}\right)$; here, we also characterize this probability experimentally. Figure 1 shows the fraction of randomized trials (out of a total of 100,000 trials) for which the stash was not large enough to prevent cycles, for varying choices of $\epsilon$ and $s$. We evaluate input sizes $n=10^{i}$ for $i=1 \ldots 6$, of which three sizes are shown in the figure. In our set intersection experiments, we instantiate the Cuckoo hashing version with $\epsilon=0.02$ (shown by the solid vertical line in the figure) and a stash size of 2 . With this configuration, $n=1,000$ has the highest fraction of failed constructions at $\sim 0.55 \%$, while inputs that were very small or large had a much lower rate of failure ( $n=10$ has $\sim 0.005 \%, n=10^{6}$ has none). We conclude that these chosen parameters are adequate to balance safety and performance across various values of $n$, although an implementation might choose to vary both parameters as a function of $n$, in order to better tune this tradeoff to a particular setting. As we will see next, however, this still will not result in any real benefit compared with the simpler hashing strategies.

Recall that the set intersection protocol consists of three main stages: (1) $P_{1}$ 's setup to construct an encrypted polynomial. The overhead of this step is $O\left(m_{1}\right)$ exponentiations in the basic protocol, and $O(B M)=O\left(m_{1}\right)$ exponentiations in each of the hashingbased constructions, where $B$ is the number of bins and $m_{1}$ is the size of $P_{1}$ 's input. (2) $P_{2}$ 's evaluation of this polynomial on its inputs. This step takes $O\left(m_{1} m_{2}\right)$ exponentiations in the basic method, $O\left(m_{2} \log m_{1}\right)$ exponentiations in the random hash method, $O\left(m_{2} \log \log m_{1}\right)$ exponentiations in the balanced allocations construction, and only $O\left(m_{2}\right)$ exponentiations in the Cuckoo hash construction. (3) $P_{1}$ 's subsequent recovery of the intersection. This step requires $O\left(m_{2}\right)$ decryptions in each of the methods.

Measured performance trends First, let us analyze the high-level performance trends of our various set intersection constructions. Figure 2 shows the performance of the set intersection protocol for these three stages in log-log scale, where both client and server have the same input set size ranging from 10 to 1 million elements. The figure demonstrates the poor performance implications of using a single polynomial constructed across all of $P_{1}$ 's input (i.e., "No Hashing"), as compared to the hashing-based schemes. The latter's performance is almost linear in $n$.


Fig. 2. Computational set intersection performance for $P_{1}$ 's setup and precomputation (left), $P_{2}$ 's evaluation (center), and $P_{1}$ 's corresponding recovery of the intersection (right).

Explaining asymptotic performance We now explain the measured performance trends of the three hashing-based approaches, as it differs from their asymptotic computational overhead. To simplify the notation, we consider the case of $m_{1}=m_{2}=n$. Recall that asymptotically, the computation overhead of the constructions based on random hashing, balanced allocations, and Cuckoo hashing is $O(n \log n), O(n \log \log n)$, and $O(n)$, respectively, where the bulk of the computation is done in Step 2. Our experiments, however, show the opposite: The run time of the random hashing-based approach, in Stages 2 and 3, is better than that of the balanced allocations scheme, which is in turn better than the run time of the Cuckoo hashing approach. In order to explain these results, we must examine the constants that affect the overhead.

Table 2 details the run times for the case of $n=10,000$. We ran experiments for values of $n$ up to $n=100,000$, where we parameterized the experiments based on the results from initial experiments to find values which do not result in the hash functions overflowing the capacity of the bins. The results for $n=10,000$ are representative of the results in the range we examined. The random hash scheme used $B=n / \log n=753$ bins, each of size 26. The balanced allocation scheme used $B=n / \log \log n=2,680$ bins, each of size 6 . The Cuckoo hashing scheme used $2 \cdot 1.02 \cdot n=20,400$ bins, each with a single element, and a stash of size 2 .

In Stage 1, $P_{1}$ constructs a polynomial for each bin in the hash table (corresponding to polynomials of degree 26,6 , or 1 , respectively for each scheme). The task of $P_{1}$ is to encrypt each of the coefficients of the polynomials, and, therefore the overhead per polynomial is linear in the degree. An analysis based on the number of polynomials used

Table 2. Speed (ms) of each of the hashing-based solutions, for $m_{1}=m_{2}=n=10,000$, computed as the average of 10 runs.

| Construction | Stage 1 <br> $P_{1}$ 's setup | Stage 2 <br> $P_{2}$ 's evaluation | Stage 3 <br> decryption by $P_{1}$ |
| :--- | :--- | :--- | :--- |
| Random hash | 10,936 | 52,334 | 2,153 |
| Balanced allocations | 7,833 | 68,336 | 4,692 |
| Cuckoo hash | 8,503 | 86,193 | 6,375 |

The random hash construction required a single evaluation of a polynomial of degree 26 . The balanced allocations construction evaluated two polynomials of degree 6. The Cuckoo hash construction evaluated two linear polynomials and a degree two polynomial (representing a stash of size 2)
in every experiment and their degrees shows that the overhead per polynomial of degree $d$ is about $0.55 d \mathrm{~ms}$.

Let us now focus on the overhead of Stage 2, which consists of the evaluation of the polynomials by $P_{2}$ and consumes the bulk of the overhead. Recall that $P_{2}$ must compute expressions of the form $E_{p k}(r \cdot P(y)+y)$. The overhead of computing these expressions includes a component that is linear in the degree of $P$ and some constant overhead, which does not depend on the degree and is caused by multiplying $P(y)$ by $r$ and adding $y$. The random hash based construction requires computing $n$ polynomials of a relatively high degree ( 26 when $n=10,000$ ). The balanced allocations experiment requires computing $2 n$ polynomials of a lower degree (degree 6 when $n=10,000$ ), whereas the Cuckoo hash experiment has $2 n$ evaluations of linear polynomials and $n$ evaluations of polynomials of degree 2 (corresponding to the stash).

Linear regression based on the experiment results shows that the evaluation time of a polynomial of degree $d$ is about $2.8+0.1 d \mathrm{~ms}$. Namely, that it has a large additive constant, which is comparable to increasing the degree by 28 . This large additive constant is mostly due to the combined effect of Horner's rule and of the short length of value $y$ used as input to the polynomials (which is only 5 bytes long). The Horner's rulebased implementation uses homomorphic multiplications by $y$, that are implemented as exponentiations with a short, 5-byte long exponent, whereas the multiplication by $r$ (that is done once per polynomial) is implemented as a full exponentiation. Therefore, the effect of that single exponentiation is quite dominant.

As a result, for $n=10,000$, the runtime of Step 2 in our implementation for the random hash experiment, which evaluates a single polynomial of degree 26 , is about $(2.8+2.6) \cdot n=5.4 n$. The runtime of the balanced allocations experiment, which evaluates two polynomials of degree 6 , is about $2 \cdot(2.8+.6) \cdot n=6.8 n$. The runtime of the Cuckoo hashing experiment, which evaluates for each input two linear polynomials and one polynomial of degree 2 , is about $(2 \cdot(2.8+.1)+(2.8+.2)) \cdot n=8.8 n$. The balanced allocations construction is therefore slower than the random hashing construction, and the run time of the Cuckoo hashing construction is slower than that of the balanced allocations construction.

When $n$ increases, the degrees of the polynomials of the randomized hashing and balanced allocations constructions increase and dominate the constant additive factor. As a result, the runtime should approach its asymptotic behavior, where the Cuckoo hashing construction is favorable, followed by the balanced allocations construction. To verify that this phenomenon indeed takes place, we charted the ratio of the runtimes of the randomized hashing construction and the balanced allocations construction for different values of $n$ ranging from $n=100$ to $n=100,000$. This ratio is decreasing and is almost 1 when $n=100,000$, and we assume that it will become smaller than 1 for larger values of $n$. See Fig. 3 .

As for Stage 3, note that in the three constructions $P_{2}$ sends back $n, 2 n$, and $3 n$ encryptions, respectively. $P_{1}$ must decrypt all these encryptions, and its actual run time in this stage is indeed linear in the number of decryptions that it must perform.

We note that the running time might be further optimized by increasing the sizes of the tables and by reducing the size of the bins. This will increase the overhead of Stages 1 and 3, as well the communication overhead, but should decrease the overhead of Stage 2.


Fig. 3. The ratio of the runtimes of Step 2 in the randomized hashing and balanced allocations constructions, for $n=100, \ldots, 100,000$.

To summarize our observations about the actual runtime, we found that the running time of evaluating each polynomial is linear in the degree, but it also has a constant factor which increases the degree by about 28 . As a result, each additional polynomial that is evaluated has a considerable effect. The random hashing construction evaluates a single polynomial of degree $O(\log n)$, whereas the balanced allocations construction evaluates two polynomials of degree $O(\log \log n)$ and the Cuckoo hashing construction evaluates three polynomials: two linear polynomials and a polynomial of degree 2 . Asymptotically, the performance of the constructions based on balanced allocations and Cuckoo hashing is preferable, but since these two constructions use more polynomials than the first construction, their overhead is higher than that of random hashing for the input sizes that we tested.

## 5. Security in the Presence of Malicious Adversaries

We now modify the secure set intersection protocol to accommodate malicious adversaries. We present in detail a protocol that is based on a balanced allocation scheme. A similar transformation can be applied to the protocol based on Cuckoo hashing. The overhead of the resulting protocol is of the same order as that of the corresponding semi-honest protocol. Considering our construction for semi-honest parties, a number of problems need to be taken care of:

1. $P_{1}$ can construct polynomials that would enable it to learn the value of elements that are not in the intersection. For instance, by setting the polynomial $Q_{i}$ to be identically zero, $P_{1}$ learns all elements $\left\{y \in Y: h_{0}(y)=i\right.$ or $\left.h_{1}(y)=i\right\}$. We solve this problem by presenting a zero-knowledge protocol for verifying that $\sum_{i \in\{1, \ldots, B\}} \operatorname{deg}\left(Q_{i}\right)=m_{1}$, and $Q_{i}(\cdot) \not \equiv 0$ for all $i \in\{1, \ldots, B\}$ (Sect. 5.1).
2. A more subtle problem that was overlooked in prior work is that with some homomorphic encryption PKEs, $P_{1}$ may construct $Q(\cdot)$ such that the evaluation of
$r \cdot Q(y)$ (and hence also of $r \cdot Q(y)+y$ ) is far from being random in $\mathcal{M}$, even though $Q(y) \neq 0$ and $r$ is chosen at random. This attack can be carried out with respect to encryption schemes for which the plaintext space is not a cyclic group of prime order. Therefore, in this case, $r \cdot Q(y)$ may be a random element within a smaller subgroup.
For a concrete example, consider the Paillier encryption scheme and note that since $P_{1}$ knows both the public and secret keys ( $p k=N, s k=\phi(N)$ ), it can construct a polynomial $Q(\cdot)$ such that $Q(y) \notin \mathbb{Z}_{N}^{*}$ for some specific value(s) of $y$ of its interest. ${ }^{9}$ This implies that $r \cdot Q(y) \notin \mathbb{Z}_{N}^{*}$ and so is very far from being random in $\mathbb{Z}_{N}$, hence failing to hide $y$. This problem does not exist when using ElGamal.
We are aware of two ways to address this problem: The first solution would be to have $P_{1}$ generating the polynomials after it learns $p k$ but before it learns $s k$. That is, the parties first run a secure protocol for a mutual generation of the public key in which the secret key is shared between them. Then, $P_{1}$ generates its polynomials and finally, $P_{2}$ reveals its share so that $P_{1}$ will be able to learn the secret key. For the Paillier encryption scheme, we get that coming up with a polynomial $Q(\cdot)$ and $y$ such that $Q(y) \notin \mathbb{Z}_{N}^{*}$ amounts to factoring the product $N$. The public key generation can be computed using the efficient protocol of Hazay et al. [37] that is proven with simulation-based security in the malicious setting. To the best of our knowledge, this is the only non-generic protocol that guarantees simulation-based security in the two-party setting. We further note that generating such a public key for the ElGamal scheme is easier following the underlying ideas of Diffie and Hellman [19].
Another solution by [48] would be to have $P_{2}$ encrypting some padding of the payload value added to the polynomial evaluation. This corresponds to a randomized encoding in order to compensate the loss of entropy. Looking ahead, this implies that $P_{2}$ first encodes $s$ and then adds it to the encrypted polynomial evaluation. This yields an encryption to an element in $\mathbb{Z}_{N}^{*}$ that is statistically close to a random element in this group, where the statistical difference depends on the encoding parameters. For simplicity, we employ the former solution in our protocol below.
3. A malicious $P_{1}$ may choose the hash functions adversarially, e.g., so that an overflow is caused conditioned on $P_{2}$ having certain input set. This issue can be overcome by having the two parties choose the hash functions jointly so that as long as one of them is honest the result is a randomly chosen hash function. There exist well-known efficient protocols for joint coin tossing, secure against malicious adversaries. We therefore reduce the task of choosing the hash function to joint coin tossing. It is known in advance that the hash function will be applied to some $n$ values, and we wish to ensure that the hash function that is chosen would be close to uniform on the set of $n$ inputs to which it will be applied. A straightforward method of defining this function is by choosing at random the coefficients of an $(n-1)$ degree polynomial. However, this function has the drawback that

[^9]each evaluation of the polynomial takes $O(n)$ time. A result of [57] describes a function defined by $O(n \log n)$ random bits (actually $O(n \log |V|)$ bits, where $|V|$ is the size of the range of function), where the function can be computed in $O(1)$ time, and has the property that for any set $S$ of $n$ items the function is uniform on $S$ except with probability that can be bounded by an arbitrarily small polynomial in $n$. A very recent construction for the case of a Cuckoo hash with a stash shows that the hash functions used can be defined by only $2 n^{1 / 2} \log n+O(s \log n)$ bits, where $s$ is the size of the stash, which is also $O(1)$ [4]. The run time of the resulting hash function was further improved in $[23,64]$, although it still requires the same number of random bits.
4. Lastly, while party $P_{2}$ is supposed to send $m_{2}$ pairs of encryptions resulting from substituting a value $y$ (known to $P_{2}$ ) in the (encrypted) polynomials $Q_{h_{0}(y)}$ and $Q_{h_{1}(y)}$, it may deviate from its prescribed computation. Thus, its input to the protocol may be ill-defined.
The standard solution to this problem involves the usage of zero-knowledge proofs for demonstrating the correct behavior by $P_{2}$. The cost of such a proof for each $y \in Y$ would be proportional to the size of X since $P_{2}$ cannot disclose the identity of $y$ 's bin. Therefore, the overall cost would be quadratic in the size of $X$. Instead, we introduce a technique that enables $P_{1}$ to redo the entire computation supposedly carried out by $P_{2}$ on any value $y$ supposedly in the intersection and verify that its outcome is consistent with the messages received from $P_{2}$. This is where we incorporate the usage of a random oracle $\mathcal{H}$ in our protocol.
We first describe how to use this technique in the case where $P_{2}$ evaluates a single polynomial first (the "No Hashing" case) and then extend it for the hashing-based protocols.
For each $y \in Y$ in $P_{2}$ 's input set, it chooses a random element $s$ and computes the encryption of $r \cdot Q(y)+s$, where the randomness used for this computation is taken from $\mathcal{H}(s)$. Clearly, if $Q(y)=0$, then $P_{1}$ learns $s$ and can be easily verified whether there exists $x \in X$ such that together with $s$ yields exactly the same encryption. The security argument shows that this is true only for elements in the intersection set. In particular, when $P_{2}$ deviates from the prescribed computation, $P_{1}$ records an incorrect output with only a negligible probability.
When using a balanced allocation scheme (or Cuckoo hashing), $P_{2}$ needs to evaluate more than one polynomial on the same value $y$. Furthermore, when $y$ is in the intersection, exactly one of these polynomials evaluates to zero. To force $P_{2}$ to act according to the protocol, we let $P_{2}$ repeat the above computation twice so that it ends up sending two encryptions for $r_{0} \cdot Q_{h_{0}(y)}(y)+s_{1}$ and $r_{1} \cdot Q_{h_{1}(y)}(y)+s_{0}$. By learning $s_{b} \in\left\{s_{0}, s_{1}\right\}, P_{1}$ is able to extract $s_{1-b}$ and hence verify $P_{2}$ 's computation. Loosely speaking, by applying the random oracle on $s_{b}$, the adversary learns $r_{1-b}$ and thus is able to extract $s_{1-b}$ from the plaintext that is combined with $s_{1-b}$. This method is different than the one applied in [38] that uses a single value of $s$ for both bins, where the randomness is extracted from the outcome of a PRF evaluated on $s$. Specifically, in [38], $P_{2}$ commits to each $y$ and the payload value $s$ first. Later, $P_{2}$ uses $s$ to generate the randomness used to evaluate the corresponding polynomial with $y$.

### 5.1. Properties of Homomorphic PKE

In this section, we consider the properties of the homomorphic PKE that we need for proving the security of our protocol. We demonstrate that these properties hold for both Paillier and ElGamal PKEs.

1. Given a random $p k$ (but not $s k$ ), it is infeasible to come up with a message $m \neq 0$ such that $m \notin \mathbb{G}$, for $\mathbb{G} \subseteq \mathcal{M}$ a multiplicative group for which only a negligible fraction of $\mathcal{M}$ is not in $\mathbb{G}$ (as specified in Definition 3.2). For Paillier $\mathbb{G}$ denotes $\mathbb{Z}_{N}^{*}$, whereas for ElGamal $\mathbb{G}$ denotes $\mathbb{Z}_{q}^{*}$.
2. There exists a polynomial-time algorithm that on input ( $p k, c$ ) outputs 1 iff $c$ is valid, i.e., in the range of $E_{p k}(\cdot ; \cdot)$. For Paillier one only needs to check that $c \in \mathbb{Z}_{N^{2}}^{*}$. For ElGamal, one can use the subgroup of quadratic residues modulo $q^{\prime}=2 q+1$.
3. There exists a protocol $\pi_{\text {KEY }}$ for securely computing the key generation functionality $\mathcal{F}_{\text {KEY }}$ in the presence of malicious adversary, where $\mathcal{F}_{\text {KEY }}$ is defined by

$$
\begin{equation*}
\left(1^{k}, 1^{k}\right) \mapsto\left(\left(p k, s k_{1}\right),\left(p k, s k_{2}\right)\right) \tag{2}
\end{equation*}
$$

where $(p k, s k) \leftarrow G\left(1^{k}\right), s k_{1}$ and $s k_{2}$ are random shares of $s k$ that do not leak any information about the secret key, and (efficient) reconstruction of $s k$ is with respect to the specific PKE via a secure two-party protocol. We also require that given $s k$, a simulator can efficiently compute shares $s k_{1}, s k_{2}$ such that the distribution over these shares when output by the protocol and by a simulator is computationally indistinguishable.
For Paillier, one can use the protocol of [37] for which its efficiency is dominated by the number of trial divisions for testing the candidates composites. This protocol ensures an improved analysis based on the analysis from [10] for which a random number of length 1,024 is a prime with probability $1 / 44$, condition that it passed a trial division with some threshold parameter $B$. Thus, the expected number of attempts is 1,936 . For ElGamal, the key setup protocol by Diffie and Hellman [19] (that can be made secure for the malicious setting) can be used here.
4. There exists an efficient computational zero-knowledge proof of knowledge $\pi_{\text {mult }}$ for proving the multiplication of two plaintexts. The relation $\mathcal{R}_{\text {MULT }}$ is formalized as follows,
$\mathcal{R}_{\text {MULT }}=\left\{\left(\left(p k,\left(e_{a}, e_{b}, e_{c}\right)\right)\right) \left\lvert\, \exists\left(a, r_{a}, b, r_{b}, r_{c}\right) s . t \begin{array}{c}e_{a}=E_{p k}\left(a ; r_{a}\right) \wedge e_{b}=E_{p k}\left(b ; r_{b}\right) \\ \wedge e_{c}=E_{p k}\left(a b ; r_{c}\right)\end{array}\right.\right\}$.

A constant-round zero-knowledge proof for $\mathrm{Ł}_{\text {mult }}$ with 15 exponentiations can be found in [20]. ${ }^{10}$

[^10]5. There exists an efficient computational zero-knowledge proof $\pi_{\mathrm{NZ}}$ for proving that some ciphertext encrypts a nonzero plaintext. Formally,
$$
\mathrm{Ł}_{\mathrm{NZ}}=\left\{\left(p k, e_{a}\right) \mid e=E_{p k}\left(a ; r_{a}\right) \text { for some } a \neq 0, r_{a}\right\} .
$$

A simple reduction to $\pi_{\text {MULT }}$ is as follows. The prover chooses random $b, r_{b}, r_{c}$, sets $e_{b}=E_{p k}\left(b ; r_{b}\right), e_{c}=E_{p k}\left(a b ; r_{c}\right)$ and proves that $\left(\left(p k,\left(e_{a}, e_{b}, e_{c}\right)\right),\left(a, r_{a}, b, r_{b}\right.\right.$, $\left.\left.r_{c}\right)\right) \in \mathcal{R}_{\text {mult }}$. If the proof is accepted, then the prover sends $a b, r_{c}$, and the verifier accepts if $e_{c}=E_{p k}(a b ; r c)$ and $a b \neq 0$.
6. There exists an efficient computational zero-knowledge proof $\pi_{\text {ZERO }}$ for proving that some ciphertext encrypts the zero plaintext. Formally,

$$
\mathrm{Ł}_{\text {ZERO }}=\left\{\left(p k, e_{a}\right) \mid e=E_{p k}\left(0 ; r_{a}\right) \text { for some } r_{a}\right\} .
$$

A constant-round zero-knowledge proof for $Ł_{\text {ZERO }}$ with 13 exponentiations for ElGamal can be found in [17] and with 8 exponentiations for Paillier [20]. ${ }^{11}$

### 5.2. Zero-Knowledge Proof of Knowledge for $\mathcal{R}_{\text {PoLY }}$

In addition to the above, we will need a zero-knowledge proof of knowledge for proving the correctness of the polynomials sent by $P_{1}$. We recall that in the set intersection protocol presented below, $P_{1}$ generates $B$ polynomials $\left\{Q_{i}(\cdot)\right\}_{i \in\{1, \ldots, B\}}$ representing its input set $X$ and sends their encryptions to $P_{2}$. We present an efficient zero-knowledge proof of knowledge $\pi_{\text {PoLy }}$ that is used for checking these encrypted polynomials. ${ }^{12}$ The basic idea of our proof is to count the number of nonzero coefficients and compare this result with $m-B$ (we subtract $B$ since there are $B$ bins and the number of roots of a polynomial $Q_{i}(\cdot)$ is upper bounded by $\left.\operatorname{deg} Q_{i}(\cdot)-1\right)$. To prove this, the prover sends an encryption $Z_{i, j}$ of 1 for every $0 \leq j \leq \operatorname{deg}\left(Q_{i}(\cdot)\right)$, and 0 otherwise. It then proves that it computed these encryptions correctly by proving that $Z_{i, j}$ is indeed an encryption of a value within $\{0,1\}$, and that $Q_{i, j} \cdot\left(1-Z_{i, j}\right)=0$ for all $j \in\{0, \ldots, M\}$. However, a problem arises if $Q_{i}(\cdot)$ has zero coefficients. In this case the prover can convince the verifier that $\sum_{i} \operatorname{deg}\left(Q_{i}(\cdot)\right)=m$ even if it is actually larger (as we only count the nonzero elements). To solve this problem, we add an additional check that the set $Z_{i, 0}, Z_{i, 1}, \ldots, Z_{i, M}$ is monotonically non-increasing, which guarantees that $\forall i$, and $j \geq \operatorname{deg}\left(Q_{i}(\cdot)\right)$, the event in which $Z_{i, j}=0$ and $Z_{i, j+1}=1$ does not happen. The proof is concluded by having the parties sum up these values using homomorphic addition, and having the prover proving that the result is an encryption of $m-B$. We note that the attack in which $P_{1}$ encrypts the zero polynomials is prevented by assuming that the coefficient of the highest degree equals 1 .

Formally, the relation is defined by,

$$
\mathcal{R}_{\text {POLY }}=\left\{\left(\left\{q_{i, j}\right\}_{i, j}, m, p k\right),\left(\left\{Q_{i, j}, r_{i, j}\right\}_{i, j}\right) \left\lvert\, \begin{array}{c}
\forall i, j q_{i, j}=E_{p k}\left(Q_{i, j} ; r_{i, j}\right) \wedge \\
\sum_{i} \operatorname{deg}\left(Q_{i}(\cdot)\right)=m \wedge \forall i, Q_{i}(\cdot) \neq 0
\end{array}\right.\right\}
$$

[^11]where $Q_{i}(x)=\sum_{j=0}^{M} Q_{i, j} \cdot x^{j}$, and $i \in\{1, \ldots, B\}, j \in\{0, \ldots, M\}$.
A protocol for $\mathcal{R}_{\text {POLY }}$, as well as a complete proof, can be found in [38]. For the sake of completeness, we give a slightly modified description of their protocol.

Protocol 2. (Zero-knowledge proof of knowledge $\pi_{\text {POLY }}$ for $\mathcal{R}_{\text {POLY }}$ ):
Joint statement: A collection of $B$ sets, each set is of $M+1$ encryptions $\left\{q_{i, j}\right\}_{i \in\{1, \ldots, B\}, j \in\{0, \ldots, M\}}$, a public key pk and an integer $m$.

- Auxiliary inputs for the prover: A collection of $B$ sets, each set is of $M+1$ values $\left\{Q_{i, j}, r_{i, j}\right\}_{i \in\{1, \ldots, B\}, j \in\{0, \ldots, M\}}$ such that the conditions in $\mathcal{R}_{\text {POLY }}$ are met.
- Convention: Both parties check every received ciphertext for validity and abort if an invalid ciphertext is received. Unless written differently, $i \in\{1, \ldots, B\}$ and $j \in\{0, \ldots, M\}$.


## - The protocol:

1. For every $i$ let $Q_{i}(x)=\sum_{j=0}^{M} Q_{i, j} \cdot x^{j}=\sum_{j=0}^{M} D_{s k}\left(q_{i, j}\right) \cdot x^{j}$, and recall that by our convention (cf. Footnote 12), $\sum_{i} \operatorname{deg}\left(Q_{i}(\cdot)\right)=m$. Let $Z_{i, j}=1$ for $0 \leq j \leq \operatorname{deg}\left(Q_{i}(\cdot)\right)$, and otherwise $Z_{i, j}=0$ (i.e., $Z_{i, 0}=1, Z_{i, 1}, \ldots, Z_{i, M}$ is monotonically non-increasing, and $\left.\sum_{j \in\{0, \ldots, M\}} Z_{i, j}=\operatorname{deg}\left(Q_{i}(\cdot)\right)+1\right)$.
2. The prover $P$ computes $z_{i, j}=E_{p k}\left(Z_{i, j} ; \tilde{r}_{i, j}\right)$ and sends $\left\{z_{i, j}\right\}_{i, j}$ to the verifier $V$.
$\forall i, P$ performs the following:
(a) For each $z_{i, j}$ encrypting a value $Z_{i, j}$, it proves that $Z_{i, j} \cdot\left(1-Z_{i, j}\right)=0$ (and hence $\left.Z_{i, j} \in\{0,1\}\right) .{ }^{13}$ Let $m_{i, j}=E_{p k}\left(0 ; \hat{r}_{i, j}\right)$. Then, $P$ proves that $\left(\left(p k, z_{i, j}, z_{i, j}^{\prime}, m_{i, j}\right),\left(Z_{i, j}, \tilde{r}_{i, j}, 1-Z_{i, j}, \tilde{r}_{i, j}^{\prime}, \hat{r}_{i, j}\right)\right) \in \mathcal{R}_{\text {MULT }}$ and that $m_{i, j}$ is an encryption of zero by decrypting $m_{i, j}$.
(b) $P$ proves that $Z_{i, 0}, Z_{i, 1}, \ldots, Z_{i, M}$ is monotonically non-increasing, i.e., that $Z_{i, j}=0$ and $Z_{i, j+1}=1$ does not happen for any value of $j \in$ $\left\{\operatorname{deg}\left(Q_{i}(\cdot)\right), \ldots, M-1\right\}$. For that, $P$ and $V$ compute an encryption of $2 Z_{i, j}+Z_{i, j+1}-1$ (note that both parties can compute this encryption) and P proves in zero-knowledge that $\left(p k, E_{p k}\left(2 Z_{i, j}+Z_{i, j+1}-1\right)\right) \in Ł_{\mathrm{Nz}}$.
(c) $P$ completes the proof that the values $Z_{i, j}$ were constructed correctly by proving (similarly to Step $2 a$ above) that $Q_{i, j} \cdot\left(1-Z_{i, j}\right)=0$ for all $j \in\{0, \ldots, M\}$.
3. Finally, to prove that the sum of degrees of the polynomials $\left\{Q_{i}(\cdot)\right\}_{i}$ equals $m$, both parties compute an encryption $t$ of $T=\sum_{i, j} Z_{i, j}-B$, and $P$ proves that $\left(p k, t / E_{p k}(m)\right) \in$ Ł $_{\text {ZERO }}$.
4. V verifies all the zero-knowledge proofs and decryptions. If any of the verifications fails, $V$ outputs 0 ; otherwise, it outputs 1 .

Note that Protocol $\pi_{\text {poly }}$ runs in a constant number of rounds because each of the zero-knowledge proofs can be implemented in constant rounds and can be invoked in

[^12]parallel. The parties compute and exchange $O(B M)=O(m)$ encryptions and execute $O(B M)=O(m)$ zero-knowledge proofs (for $\mathcal{R}_{\mathrm{MULT}}, \mathrm{Ł}_{\mathrm{NZ}}, \mathrm{Ł}_{\text {ZERO }}$ ). Overall, these amount to performing $O(m)$ exponentiations and exchanging $O(m)$ group elements.

Proposition 5.1. Assume that $\pi_{\mathrm{MULT}}, \pi_{\mathrm{ZERO}}$, and $\pi_{\mathrm{NZ}}$ are as described above and that $(G, E, D)$ is homomorphic PKE that is good for plaintext space (cf. Definition 3.2). Then Protocol 2 is a computational zero-knowledge proof of knowledge for $\mathcal{R}_{\text {poly }}$ with perfect completeness.

A complete proof can be found in [38]. Intuitively, the zero-knowledge property follows from the fact that all the sub-protocols are zero-knowledge as well. Knowledge extraction follows from the fact that $\pi_{\text {MULT }}$ is a proof of knowledge proof, so that the plaintexts and randomness can be extracted in Step 2c. Finally, soundness follows from soundness of sub-protocols, ensuring that the degree of the polynomial cannot exceed $M$.

### 5.3. Secure Set Intersection in the Random Oracle Model

We are now ready to present a formal description of our protocol in a setting where the parties are malicious and have oracle access to a random function $\mathcal{H}: \mathbb{G} \rightarrow \mathcal{M}\|\mathcal{M}\| R$, i.e., the outcome is split into three values: (1) The first value is used to mask the encrypted polynomial evaluation (for the message space specified for either Paillier or ElGamal), (2) the second value is used to mask the input to the polynomial, (3) and the third value is used to encrypt the values $s_{0}$, $s_{1}$ (to rerandomize the ciphertext encrypting the outcome of the polynomial evaluation), for $R$ denoting the randomness space from which the random coins for the specified PKE are sampled from. A high-level description is presented in Fig. 4.

Protocol 3. ( $\pi_{\cap}^{\mathrm{RO}}-$ secure set intersection in the random oracle model):

- Inputs: The input of $P_{1}$ is $m_{2}$ and a set $X \subseteq\{0,1\}^{p(k)}$ containing $m_{1}$ items; the input of $P_{2}$ is $m_{1}$ and a set $Y \subseteq\{0,1\}^{p(k)}$ containing $m_{2}$ items (hence, both parties know $m_{1}$ and $m_{2}$ ).
- Auxiliary inputs: A security parameter $1^{k}$.
- Convention: Both parties check every received ciphertext for validity and abort if an invalid ciphertext is received.


## - The protocol:

1. Key setup: $P_{1}$ and $P_{2}$ run an execution of $\pi_{\mathrm{KEY}}$ for generating $p k$ and random shares $s k_{1}, s k_{2}$ of sk. Let $\left(p k, s k_{1}\right)$ and $\left(p k, s k_{2}\right)$ denote the respective outputs of $P_{1}$ and $P_{2}$ from this execution.
2. Setting the balanced allocation scheme: $P_{1}$ computes the parameters $B, M$ for the scheme and the parties run a coin tossing protocol in order to choose the seeds for two randomly chosen hash functions $h_{0}, h_{1}:\{0,1\}^{p(k)} \rightarrow[B]$. It sends $B, M, h_{0}, h_{1}$ to $P_{2} . P_{2}$ checks that the parameters $B, M$ and the seeds $h_{0}, h_{1}$ were computed correctly and aborts otherwise.
3. Creating polynomials for the set $X$ : For every $x \in X, P_{1}$ maps $x$ into the less occupied bin from $\left\{h_{0}(x), h_{1}(x)\right\}$ (ties broken arbitrarily). Let $\mathcal{B}_{i}$ denote


Fig. 4. A high-level diagram of $\pi_{\cap}^{\mathrm{RO}}$ (random oracle model).
the set of elements mapped into bin $i$ and let $Q_{i}(x) \stackrel{\text { def }}{=} \sum_{j=0}^{M} Q_{i, j} \cdot x^{j}$ denote a polynomial with the set of roots $\mathcal{B}_{i}$ (if $\mathcal{B}_{i}=\emptyset$ then $P_{1}$ sets $Q_{i}(x)=1$ ). If $\left|\mathcal{B}_{i}\right|<M$ then $P_{1}$ sets the $M+1-\left|\mathcal{B}_{i}\right|$ highest degree coefficients of $Q_{i}(\cdot)$ to zero.
$P_{1}$ encrypts the coefficients of the polynomials, setting $q_{i, j}=E_{p k}\left(Q_{i, j} ; r_{i, j}\right)$. It sends the encrypted coefficients to $P_{2}$.
4. Checking the polynomials: $P_{1}$ and $P_{2}$ engage in an execution of $\pi_{\text {Poly }}$ for which $P_{1}$ enters $p k$ and the sets $\left\{Q_{i, j}, r_{i, j}\right\}_{i \in\{1, \ldots, B\}, j \in\{0, \ldots, M\}}$ and $P_{2}$ enters

5. Key setup completion: $P_{2}$ sends his share of the private key $s k_{2}$, as well as the randomness used for generating this share, to $P_{1} . P_{1}$ reconstructs sk from the shares sk$k_{1}, s k_{2}$ and checks that the result is a valid private key that corresponds to pk and that $P_{2}$ sent the correct share. If it is not, it aborts.
6. Substituting in the polynomials: Let $y_{1}, \ldots, y_{m_{2}}$ be a random ordering of the elements of set $Y . P_{2}$ does the following for all $\alpha \in\left\{1, \ldots, m_{2}\right\}$ :
(a) It sets $\hat{h}_{0}=h_{0}\left(y_{\alpha}\right), \hat{h}_{1}=h_{1}\left(y_{\alpha}\right)$ (these values denote the bin number). Then, it chooses $s_{0}^{\alpha}, s_{1}^{\alpha} \leftarrow_{R} \mathcal{M}$ and parses $\mathcal{H}\left(s_{j}^{\alpha}\right)$ to obtain random strings $r_{j}, \tilde{r}_{j}, \hat{r}_{j}$ of appropriate lengths for their usage below, for all $j \in\{0,1\}$ (i.e., $r_{j}, \tilde{r}_{j}, \hat{r}_{j}=\mathcal{H}\left(s_{j}^{\alpha}\right)$ ).
(b) $P_{2}$ uses the homomorphic properties of the encryption scheme to evaluate $E_{p k}\left(r_{0} \cdot Q_{\hat{h}_{0}}\left(y_{\alpha}\right)\right)$ and $E_{p k}\left(r_{1} \cdot Q_{\hat{h}_{1}}\left(y_{\alpha}\right)\right)$. It then applies an homomorphic operation on these ciphertexts and $E_{p k}\left(s_{1}^{\alpha} ; \hat{r}_{0}\right)$ and $E_{p k}\left(s_{0}^{\alpha} ; \hat{r}_{1}\right)$, respectively. This results in $e_{0}^{\alpha}=E_{p k}\left(r_{0} \cdot Q_{\hat{h}_{0}}\left(y_{\alpha}\right)+s_{1}^{\alpha}\right)$ and $e_{1}^{\alpha}=$ $E_{p k}\left(r_{1} \cdot Q_{\hat{h}_{1}}\left(y_{\alpha}\right)+s_{0}^{\alpha}\right)$.
Finally, $P_{2}$ sends $e_{0}^{\alpha}$, $e_{1}^{\alpha}$ together with $t_{0}^{\alpha}=\tilde{r}_{0} \oplus y_{\alpha}, t_{1}^{\alpha}=\tilde{r}_{1} \oplus y_{\alpha}$ to $P_{1}$.
7. Computing the intersection: $P_{1}$ checks that $P_{2}$ sent $m_{2}$ tuples and aborts otherwise. For each received $e_{0}^{\alpha}, e_{1}^{\alpha}, t_{0}^{\alpha}, t_{1}^{\alpha}$, party $P_{1}$ computes $s_{0}^{\prime}=D_{s k}\left(e_{1}^{\alpha}\right)$ and $s_{1}^{\prime}=D_{s k}\left(e_{0}^{\alpha}\right)$ and sets $r_{j}^{\prime}, \tilde{r}_{j}^{\prime}, \hat{r}_{j}^{\prime}=\mathcal{H}\left(s_{j}^{\prime}\right)$ for all $j$.
If $t_{j}^{\alpha}=\tilde{r}_{j}^{\prime} \oplus x$ for some $x \in X$ and $j \in\{0,1\}$, then $P_{1}$ records $s_{j}^{x}=s_{j}^{\prime}$ and $s_{1-j}^{x}=s_{1-j}^{\prime}-r_{j}^{\prime} \cdot Q_{h_{j}(x)}(x) . P_{1}$ uses these recorded values and $x$ in order to recompute the set $e_{0}^{\alpha}, e_{1}^{\alpha}, t_{0}^{\alpha}, t_{1}^{\alpha}$. If its outcome agrees with $P_{2}$ 's messages, it records $x$.
Namely, $P_{1}$ recomputes $e_{j}^{\alpha}$ and $t_{j}^{\alpha}$ using $P_{2}$ 's procedure in Step 3, but with $s_{0}^{x}, s_{1}^{x}$ in the role of $s_{0}^{\alpha}, s_{1}^{\alpha}$.

Correctness Notice first that if both parties are honest, then $P_{1}$ outputs $X \cap Y$ with probability negligibly close to $1:$ (i) For elements $x \in X \cap Y$, there exists a $y_{\alpha}$ that zeros at least one of the polynomials $Q_{h_{0}\left(y_{\alpha}\right)}(\cdot), Q_{h_{1}\left(y_{\alpha}\right)}(\cdot)$, say w.l.o.g. the first polynomial. Hence, getting the messages sent for $y_{\alpha}, P_{1}$ recovers $s_{1}^{\prime}=s_{1}^{\alpha}$ immediately. Moreover, since it already knows $y_{\alpha}$ from $t_{1}^{\alpha}$, and $r_{1}$ from $\mathcal{H}\left(s_{1}^{\alpha}\right)$, it can recover $s_{0}^{\alpha}$ as well by simply computing $s_{0}^{\alpha}=s_{0}^{\prime}-r_{1} \cdot Q_{h_{1}\left(y_{\alpha}\right)}\left(y_{\alpha}\right)$ and verify that the recomputed encryptions $e_{0}^{\alpha}, e_{1}^{\alpha}, t_{0}^{\alpha}, t_{1}^{\alpha}$ agree with $P_{2}$ 's messages. Consequently, $P_{1}$ records $x$ in its output. (ii) For elements $x \in X \backslash Y, P_{1}$ records $x$ only if for some $y_{\alpha}$ and $j$ the values $r_{j}^{\prime}, \tilde{r}_{j}^{\prime}, \hat{r}_{j}^{\prime}=\mathcal{H}\left(s_{j}^{\prime}\right)$ are such that $t_{j}^{\alpha}=\tilde{r}_{j}^{\prime} \oplus x$ and furthermore, the conditions on $e_{0}^{\alpha}, e_{1}^{\alpha}, t_{0}^{\alpha}, t_{1}^{\alpha}$ are met. As proven below, this occurs with only a negligible probability.

Theorem 5.1. Assume that $\pi_{\mathrm{KEY}}$ and $\pi_{\mathrm{POLY}}$ are as described above, that $(G, E, D)$ is a semantically secure homomorphic PKE that is good for plaintext space (cf. Definition 3.2), and that $\mathcal{H}$ is a random oracle. Then, Protocol 3 securely computes $\mathcal{F}_{\cap}$ in the presence of malicious adversaries in the random oracle model.

Proof. We separately prove security in the case that $P_{1}$ is corrupted and the case that $P_{2}$ is corrupted. Our proof is in a hybrid model where a trusted party computes the ideal functionalities $\mathcal{F}_{\text {KEY }}$ and the zero-knowledge proof of knowledge $\mathcal{R}_{\text {POLY }}$.
$P_{1}$ is corrupted. Intuitively, all a corrupted $P_{1}$ can do in protocol $\pi_{\cap}^{\mathrm{RO}}$ is to try and construct the polynomials $Q_{i}(\cdot)$ such that it gains information about $P_{2}$ 's input $Y$. There are two provisions in the protocol against such an attempt: (i) The application of $\pi_{\text {Poly }}$ ensures that as long as all the homomorphic operations are within elements of $\mathbb{G}$ (as specified in Definition 3.2), $P_{1}$ should fail; and (ii) Letting $P_{1}$ learn $s k$ only after it constructs the polynomials ensures that this would indeed be the case. The following proof makes this intuition formal. Let $\mathcal{A}$ denote an adversary controlling $P_{1}$. Construct a simulator $\mathcal{S}$ for $\mathcal{A}$ as follows:

Convention: During the entire execution, $\mathcal{S}$ evaluates queries to the random oracle $\mathcal{H}$. Such queries are made by $\mathcal{A}$ or by $\mathcal{S}$ during its simulation of $P_{2}$. To evaluate $\mathcal{H}(s), \mathcal{S}$ first checks if it has already recorded a pair $(s, r)$, in which case $\mathcal{H}(s)$ evaluates to the value $r$. Otherwise, $\mathcal{S}$ chooses a random string $r$ of the appropriate length, records $(s, r)$ and evaluates $\mathcal{H}(s)$ to $r$.

1. $\mathcal{S}$ is given $X, m_{2}$, and $1^{k}$ and invokes $\mathcal{A}$ on these inputs. $\mathcal{S}$ sets $m_{1}=|X|$.
2. $\mathcal{S}$ receives from $\mathcal{A}$ its input $1^{k}$ for the ideal functionality $\mathcal{F}_{\text {KEY }}$ and computes $(p k, s k) \leftarrow G\left(1^{k}\right)$. Next, $\mathcal{S}$ computes random shares $s k_{1}, s k_{2}$ such that the reconstruction with these shares yields $s k$, and sends $\left(p k, s k_{1}\right)$ to $\mathcal{A}$, emulating the response of a trusted party for $\mathcal{F}_{\text {Key }}$.
3. Upon engaging in a coin-tossing protocol with $\mathcal{A}, \mathcal{S}$ receives from $\mathcal{A}$ the parameters $B, M$ and the seeds for the two random chosen hash functions $h_{0}, h_{1}:\{0,1\}^{p(k)} \rightarrow$ [ $B$ ] used in the balanced allocation scheme. If the parameters $B, M$ or $h_{0}, h_{1}$ were not computed correctly, $\mathcal{S}$ sends $\perp$ to the trusted party for $\mathcal{F}_{\cap}$ and aborts.
4. $\mathcal{S}$ receives from $\mathcal{A}$ the encrypted polynomials $\left\{q_{i, j}\right\}_{i \in\{1, \ldots, B\}, j \in\{0, \ldots, M\}}$.
5. $\mathcal{S}$ receives from $\mathcal{A}$ its input $\left\{Q_{i, j}, r_{i, j}\right\}_{i \in\{1, \ldots, B\}, j \in\{0, \ldots, M\}}$ and $p k$ for the ideal computation $\mathcal{R}_{\text {POLY }}$. If the conditions of $\mathcal{R}_{\text {POLY }}$ for outputting $(\lambda, 1)$ are not met, then $\mathcal{S}$ sends $\perp$ to the trusted party for $\mathcal{F}_{\cap}$ and aborts.
6. $\mathcal{S}$ sends $s k_{2}$ to $\mathcal{A}$.
7. $\mathcal{S}$ sets $\tilde{X}=\cup_{i=1}^{B}\left\{x: Q_{i}(x)=0 \wedge\left(h_{0}(x)=i \vee h_{1}(x)=i\right)\right\}$ (note that for Paillier, $\mathcal{S}$ would know the factorization $p$ and $q$ of the public key $N$, so that it can factor $Q_{i}(\cdot)$ over the fields $\mathbb{Z}_{q}$ and $\mathbb{Z}_{p}$, recovering the roots over $\mathbb{Z}_{N}$ using the Chinese Reminder Theorem) and completes $\tilde{X}$ to size $m_{1}$ by adding random elements from $\{0,1\}^{p(k)} . \mathcal{S}$ sends $\tilde{X}$ to the trusted party for $\mathcal{F}_{\cap}$ and receives as answer a set $Z=\tilde{X} \cap Y . \mathcal{S}$ sets $\tilde{Y}$ to $Z$ and completes $\tilde{Y}$ to the size $m_{2}$ by adding random elements from $\{0,1\}^{p(k)}$.
8. $\mathcal{S}$ plays the role of the honest $P_{2}$ for the rest of the execution, using $\tilde{Y}$ as the input for $P_{2}$, with the exception that $\mathcal{S}$ itself evaluates queries to $\mathcal{H}$ as described above.
9. $\mathcal{S}$ outputs whatever $\mathcal{A}$ does.

We claim that $\mathcal{A}$ 's output distributions in the hybrid and the simulated executions are statistically close. The potential difference between these executions is in Step 7 of the simulation, where $\mathcal{S}$ plays the role of the honest $P_{2}$ with $\tilde{Y}$ instead of $Y$. Ignoring the case where $Y \cap \tilde{X} \neq Y \cap X$ (which happens with only a negligible probability), the difference is that members of $\tilde{Y}^{\prime}=\tilde{Y} \backslash(Y \cap X)$ are used in the simulated execution, whereas members of $Y^{\prime}=Y \backslash X=Y \backslash(Y \cap X)$ are used in the hybrid execution.

Let $y \in Y^{\prime}\left(y \in \tilde{Y}^{\prime}\right)$ be the $\alpha$ th element considered in the hybrid (simulated) execution, then $\mathcal{A}$ receives four values: $e_{0}^{\alpha}, e_{1}^{\alpha}, t_{0}^{\alpha}, t_{1}^{\alpha}$. Note first that unless $\mathcal{A}$ recovers either $s_{0}^{\alpha}$ or $s_{1}^{\alpha}$ the plaintexts encrypted under $e_{0}^{\alpha}, e_{1}^{\alpha}$ are statistically close to random messages. Similarly, the strings $t_{0}^{\alpha}$ and $t_{1}^{\alpha}$ are statistically close to uniformly selected strings in $\mathcal{M}$. Hence, it suffices to show that $s_{0}^{\alpha}, s_{1}^{\alpha}$ cannot be recovered, unless with negligible probability.

Recall that $e_{j}^{\alpha}=E_{p k}\left(r_{j} \cdot Q_{h_{j}(y)}(y)+s_{1-j}^{\alpha} ; \hat{r}_{j}\right)$. Then if $Q_{h_{0}(y)}(y), Q_{h_{1}(y)}(y) \in \mathbb{G}$, except with negligible probability, $r_{j}$ is uniformly distributed in $\mathbb{G}$ and we get that $r_{j} \cdot Q_{h_{j}(y)}(y)$ is uniformly distributed in $\mathbb{G}$. Therefore, the probability of guessing $s_{1-j}^{\alpha}$ is negligible in the order of $\mathbb{G}$. Moreover, the probability that $Q_{h_{0}(y)}(y) \notin \mathbb{G}$ or
$Q_{h_{1}(y)}(y) \notin \mathbb{G}$ is negligible both in the hybrid and in the simulated executions by Item 1 in Sect. 5.1. This concludes the proof.
$P_{2}$ is corrupted. Intuitively, we should prevent a corrupted $P_{2}$ from learning about $P_{1}$ 's input and from making $P_{1}$ output a wrong output. The latter concern is dealt with using a technique we mentioned above, i.e., by having $P_{1}$ recover $P_{2}$ 's randomness from $\mathcal{H}(s)$ and verify that $P_{2}$ 's messages that result in an element supposedly in the intersection are in accordance with what a honest $P_{2}$ would have sent.

Getting to the formal proof, let $\mathcal{A}$ denote an adversary controlling $P_{2}$ and construct a simulator $\mathcal{S}$ for $\mathcal{A}$ as follows:

Convention: During the entire execution, $\mathcal{S}$ evaluates queries to the random oracle $\mathcal{H}$. Such queries are made by $\mathcal{A}$ or by $\mathcal{S}$ during its simulation of $P_{1}$. To evaluate $\mathcal{H}(s), \mathcal{S}$ first checks if it has already recorded a pair $(s, r)$, in which case $\mathcal{H}(s)$ evaluates to the value $r$. Otherwise, $\mathcal{S}$ chooses a random string $r$ of the appropriate length, records $(s, r)$ and evaluates $\mathcal{H}(s)$ to $r$.

1. $\mathcal{S}$ is given $Y, m_{1}$, and $1^{k}$ and invokes $\mathcal{A}$ on these inputs. $\mathcal{S}$ sets $m_{2}=|Y|$.
2. $\mathcal{S}$ receives from $\mathcal{A}$ its input $1^{k}$ for the ideal functionality $\mathcal{F}_{\text {KEY }}$ and computes $(p k, s k) \leftarrow G\left(1^{k}\right)$. Next, $\mathcal{S}$ chooses random shares such that the reconstruction with these shares yields $s k$, and sends ( $p k, s k_{2}$ ) to $\mathcal{A}$, emulating the response of a trusted party for $\mathcal{F}_{\text {KEY }}$.
3. Upon engaging in a coin-tossing protocol with $\mathcal{A}, \mathcal{S}$ computes the parameters $B, M$ for the balanced allocation scheme and the seeds for the hash functions $h_{0}, h_{1}$. These are then sent to $\mathcal{A}$.
4. $\mathcal{S}$ sends to $\mathcal{A}, B M$ encryptions of the value 0 , under the key $p k$ (i.e., encryptions of $B$ zero polynomials). Each encryption is done with fresh randomness.
5. $\mathcal{S}$ emulates the ideal computation of $\mathcal{R}_{\text {poly }}$. It receives from $\mathcal{A}$ a set of $B M$ coefficients and $p k$. If $\mathcal{A}$ 's input is the exact set of encryptions that it received from $\mathcal{S}$ in the previous step and $p k$ then $\mathcal{S}$ returns 1 ; otherwise, it returns 0 .
6. $\mathcal{S}$ receives from $\mathcal{A}$ its private-key share. If $\mathcal{A}$ does not send $s k_{2}, \mathcal{S}$ aborts, sending $\perp$ to the trusted party for $\mathcal{F}_{\cap}$.
7. In case, $\mathcal{A}$ sends more than $m_{2}$ tuples in Step 3 of the protocol $\mathcal{S}$ aborts, sending $\perp$ to the trusted party.
8. For every $e_{0}^{\alpha}, e_{1}^{\alpha}, t_{0}^{\alpha}, t_{1}^{\alpha}$ received from $\mathcal{A}$ in the simulation of Step 3 of the protocol, every $j \in\{0,1\}$ and every pairs $\left(s_{0}, r_{0}\right),\left(s_{1}, r_{1}\right)$ recorded by $\mathcal{S}$ as part of its evaluation of the random oracle $\mathcal{H}, \mathcal{S}$ performs the following: (i) it parses $r_{j}$ as $r_{j}^{\prime}\left\|\tilde{r}_{j}\right\| \hat{r}_{j}$ and sets $y_{j}=t_{j}^{\alpha} \oplus \tilde{r}_{j}$. (ii) For all $j \in\{0,1\}$, it checks whether $e_{0}^{\alpha}, e_{1}^{\alpha}, t_{0}^{\alpha}, t_{1}^{\alpha}$ are consistent with what $P_{2}$ should have sent on $y_{\alpha}=y_{j}$ given $s_{1-j}^{\alpha}=s_{1-j}$ and $r_{j}^{\prime}, \tilde{r}_{j}, \hat{r}_{j}$. That is, $\mathcal{S}$ recomputes these encryptions using $s, r$, and $y_{\alpha}$ as the honest $P_{2}$ would in the real execution, and checks whether the result equals $e_{0}^{\alpha}, e_{1}^{\alpha}, t_{0}^{\alpha}, t_{1}^{\alpha}$. If the check succeeds, $\mathcal{S}$ locally records the value $y_{\alpha}$.
Naively implementing the above algorithm, the computational complexity of the simulator is quadratic in the number of queries. It is possible to reduce this cost into a linear overhead by additional encoding a common payload in the ciphertexts. Namely, the input to the random oracle would be the $s_{i}$ 's concatenated with additional string that is in common for both $s_{0}$ and $s_{1}$.
9. $\mathcal{S}$ sets $\tilde{Y}$ to the set of values recorded in the previous step and completes $\tilde{Y}$ to the size $m_{2}$ by adding random elements from $\{0,1\}^{p(k)}$.
10. $\mathcal{S}$ sends $\tilde{Y}$ to the trusted party and outputs whenever $\mathcal{A}$ does.

In the following, we define a sequence of hybrid games and denote by the random variable $\mathbf{H}_{\ell}^{\mathcal{A}(z)}(k, X, Y)($ for a fixed $k)$ the joint output of $\mathcal{A}$ and $P_{1}$ in hybrid game $\mathbf{H}_{\ell}$. Game $\mathbf{H}_{0}$ : The simulated execution.
Game $\mathbf{H}_{1}$ : The simulator $\mathcal{S}_{1}$ acts identically to $\mathcal{S}$ except that it does not get to know $s k$. The random variables $\mathbf{H}_{0}^{\mathcal{A}(z)}(k, X, Y)$ and $\mathbf{H}_{1}^{\mathcal{A}(z)}(k, X, Y)$ are identical as $\mathcal{S}$ and $\mathcal{S}_{1}$ do not use $s k$.
Game $\mathbf{H}_{2}$ : In this game, there is no trusted party and no honest $P_{1}$. Instead, the simulator $\mathcal{S}_{2}$ is given as input $P_{1}$ 's real input $X$. $\mathcal{S}_{2}$ works exactly like $\mathcal{S}_{1}$, except that instead of sending zero polynomials, it computes the polynomials as in the hybrid execution using the set $X$. In addition, $\mathcal{S}_{2}$ does not send $\tilde{Y}$ to the trusted party, but uses $X$ to compute and output $X \cap \tilde{Y}$. The random variables $\mathbf{H}_{1}^{\mathcal{A}(z)}(k, X, Y)$ and $\mathbf{H}_{2}^{\mathcal{A}(z)}(k, X, Y)$ are computationally indistinguishable via a reduction to the security of $(G, E, D)$. Assume, for contradiction, the existence of a distinguisher circuit $D$ for $\mathbf{H}_{1}^{\mathcal{A}(z)}(k, X, Y)$ and $\mathbf{H}_{2}^{\mathcal{A}(z)}(k, X, Y)$ and choose hash functions $h_{0}, h_{1}$ that maximize $D$ 's distinguishing advantage. Construct a distinguisher circuit $D_{\mathrm{E}}$ that distinguishes between the encryptions of two sets of messages: (i) coefficients of zero polynomials as constructed in game $\mathbf{H}_{1}$ and (ii) coefficients of polynomials corresponding to $X$ and hash functions $h_{0}, h_{1}$ as constructed in game $\mathbf{H}_{2}$. $D_{\mathrm{E}}$ receives a public key $p k$ and a vector of encryptions $\bar{c}$ under $p k$ and works exactly like $\mathcal{S}_{2}$ except that it uses the public key given to it as input instead of the public key it generated within $\mathcal{F}_{\text {KEY }}$, and forwards to $\mathcal{A}$ the vector of encryptions $\bar{c}$ instead of the encrypted polynomials. Note that the output distribution generated by $D_{\mathrm{E}}$ is identical either to $\mathbf{H}_{1}^{\mathcal{A}(z)}(k, X, Y)$ or to $\mathbf{H}_{2}^{\mathcal{A}(z)}(k, X, Y)$ (conditioned on our choice of $h_{0}, h_{1}$ ), and hence, the existence of $D$ contradicts the indistinguishability of ciphertexts of the encryption scheme ( $G, E, D$ ).
Game $\mathbf{H}_{3}$ : The simulator $\mathcal{S}_{3}$ acts identically to $\mathcal{S}_{2}$ except that $\mathcal{S}_{3}$ is given $s k$. The random variables $\mathbf{H}_{2}^{\mathcal{A}(z)}(k, X, Y)$ and $\mathbf{H}_{3}^{\mathcal{A}(z)}(k, X, Y)$ are identical as $\mathcal{S}_{2}$ and $\mathcal{S}_{3}$ do not use $s k$.
Game $\mathbf{H}_{4}$ : In this game, $\mathcal{S}_{4}$ acts identically to $\mathcal{S}_{3}$ except that $\mathcal{S}_{4}$ performs in Step 3 of the protocol the same check as in the hybrid execution, i.e., first decrypt and then recompute. To conclude the proof, we show that the random variables $\mathbf{H}_{3}^{\mathcal{A}(z)}(k, X, Y)$ and $\mathbf{H}_{4}^{\mathcal{A}(z)}(k, X, Y)$ are statistically close. Neglecting the event that new elements added to $\tilde{Y}$ are in $X$ as well and ignoring this set observe first that if an element $y_{\alpha}$ satisfies the conditions for being included in the output in Game $\mathbf{H}_{3}$, it also satisfies the conditions for being included in the output in Game $\mathbf{H}_{4}$. Namely, if $\mathcal{S}_{3}$ outputs an element $y_{\alpha}$, then it must be that $y_{\alpha} \in X$, and either $e_{0}^{\alpha}$ or $e_{1}^{\alpha}$ encrypts $s_{j}^{\alpha}$, and hence, $\mathcal{S}_{4}$ would have outputted it.

Consider now the reverse direction. Let bad denote the event where there exists an element $x \in X$ such that $\mathcal{S}_{4}$ decided to output, but should not have been outputted by $\mathcal{S}_{3}$. We show that $\operatorname{Pr}[\mathrm{bad}]$ is negligible. Note that for bad to occur it must be that for some $j$, one of $e_{0}^{\alpha}, e_{1}^{\alpha}$ is decrypted into $s_{j}^{\alpha}$ such that for some $x \in X$ the values $e_{0}^{\alpha}, e_{1}^{\alpha}, t_{0}^{\alpha}, t_{1}^{\alpha}$ are consistent with setting $y=x$ and the randomness obtained from $\mathcal{H}\left(s_{0}^{\alpha}\right)$ and $\mathcal{H}\left(s_{1}^{\alpha}\right)$,
and this occurs only with negligible probability due to the fact that the output by $\mathcal{H}$ is truly random.
Game $\mathbf{H}_{5}$ : The hybrid execution. The random variables $\mathbf{H}_{4}^{\mathcal{A}(z)}(k, X, Y)$ and $\mathbf{H}_{5}^{\mathcal{A}(z)}(k, X, Y)$ are identical as the only difference between the executions is that in $\mathbf{H}_{4}$ the outcome of the random oracle is chosen upon request, whereas in $\mathbf{H}_{5}$ it is chosen before the execution. In both cases, the outcome of the execution does not depend on entries of the random oracle which are not accessed.

Efficiency Protocol $\pi_{\cap}^{\mathrm{RO}}$ runs in a constant number of rounds, because all the zeroknowledge proofs can be implemented in constant rounds, and also the implementations of $\pi_{\mathrm{KEY}}$ and $\pi_{\text {POLY }}$ are constant round. Implementation of $\pi_{\mathrm{KEY}}$ depends on the homomorphic PKE that is used. For ElGamal, the key setup of [19] can be used here with additional zero-knowledge proofs of knowledge for proving a knowledge of discrete logarithm. The overhead induced by this protocol is relatively small since each such zero-knowledge proof of knowledge requires 9 exponentiations. A protocol for Paillier requires an increased overhead since its costs depend on the number of attempts to successfully generate a pair of primes which takes an expected number of 1,936 trials, where for each trial the number of exponentiations for generating and testing the composite is constant [10]. Clearly, ElGamal gives better efficiency with a shorter security parameter (i.e., 160 bits compared with $1,024 / 2,048$ ).

The overhead induced by protocol $\pi_{\text {POLY }}$ is $O(B M)$ which amounts to performing $O\left(m_{1}\right)$ exponentiations and exchanging $O\left(m_{1}\right)$ group elements. In the last two steps of the protocol, $P_{2}$ substitutes $m_{2}$ values in the encrypted polynomials and sends the results of the substitution to $P_{1}$. Neglecting the costs of invoking the random oracle $\mathcal{H}$, this amounts to performing $O\left(m_{2} M\right)=O\left(m_{2} \log \log m_{1}\right)$ modular exponentiations and communicating $O\left(m_{2}\right)$ group elements. We get that the overall communication costs are of sending $O\left(m_{1}+m_{2}\right)$ group elements, and the computation costs are of performing $O\left(m_{1}+m_{2} \log \log m_{1}\right)$ modular exponentiations.

## References

[1] Yossi Azar, Andrei Z. Broder, Anna R. Karlin, and Eli Upfal.Balanced allocations. SIAM Journal on Computing, 29(1):180-200, 1999.
[2] Proc. Twentieth Annual ACM Symposium on Theory of Computing, Chicago, Illinois, 2-4 May 1988.
[3] Giuseppe Ateniese, Emiliano De Cristofaro, and Gene Tsudik. (if) size matters: Size-hiding private set intersection. In Dario Catalano, Nelly Fazio, Rosario Gennaro, and Antonio Nicolosi, editors, Public Key Cryptography, volume 6571 of Lecture Notes in Computer Science, pages 156-173. Springer, 2011.
[4] Martin Aumller, Martin Dietzfelbinger, and Philipp Woelfel. Explicit and efficient hash families suffice for cuckoo hashing with a stash. 2012.
[5] Bill Aiello, Yuval Ishai, and Omer Reingold. Priced oblivious transfer: How to sell digital goods. In Advances in Cryptology-EUROCRYPT 2001, Innsbruck, Austria, May 2001.
[6] Miklós Ajtai, János Kolmós, and Endre Szemerédi. An $O(n \log n)$ sorting network. In STOC, pages 1-9, 1983.
[7] Yonatan Aumann and Yehuda Lindell. Security against covert adversaries: Efficient protocols for realistic adversaries. In Salil P. Vadhan, editor, TCC, volume 4392 of Lecture Notes in Computer Science, pages 137-156. Springer, 2007.
[8] Kenneth E. Batcher. Sorting networks and their applications. In AFIPS Spring Joint Computing Conference, pages 307-314, 32(1968).
[9] Donald Beaver. Foundations of secure interactive computing. CRYPTO, 576:377-391, 1991.
[10] Dan Boneh and Matthew K. Franklin. Efficient generation of shared rsa keys. J. ACM, 48(4):702-722, 2001.
[11] Michael Ben-Or, Shafi Goldwasser, and Avi Wigderson. Completeness theorems for non-cryptographic fault-tolerant distributed computation (extended abstract). In ACM [2], pages 1-10.
[12] Andrei Z. Broder and Michael Mitzenmacher. Using multiple hash functions to improve ip lookups. In IEEE INFOCOM 2001. 20th Ann. Joint Conference of the IEEE Computer and Communications Societies. Proceedings, 2001.
[13] Fabrice Boudot, Berry Schoenmakers, and Jacques Traore. A fair and efficient solution to the socialist millionaires' problem. Discrete Applied Mathematics, 111(1-2):23-036, 2001.
[14] Ran Canetti. Security and composition of multi-party cryptographic protocols. Journal of Cryptology, 13(1):143-202, 2000.
[15] David Chaum, Claude Crépeau, and Ivan Damgård. Multiparty unconditionally secure protocols (extended abstract). In ACM [2], pages 11-19.
[16] Jung Hee Cheon, Stanislaw Jarecki, and Jae Hong Seo. Multi-party privacy-preserving set intersection with quasi-linear complexity. Cryptology ePrint Archive, Report 2010/512, 2010.
[17] David Chaum and Torben P. Pedersen. Wallet databases with observers. In CRYPTO, pages 89-105, 1992.
[18] Jan Camenisch and Gregory M. Zaverucha. Private intersection of certified sets. In Roger Dingledine and Philippe Golle, editors, Financial Cryptography, volume 5628 of Lecture Notes in Computer Science, pages 108-127. Springer, 2009.
[19] Whitfield Diffie and Martin E. Hellman. New directions in cryptography. IEEE Transactions on Information Theory, 22(6):644-654, November 1976.
[20] Ivan Damgård and Mads Jurik. A generalisation, a simplification and some applications of Paillier's probabilistic public-key system. In 4th International Workshop on Practice and Theory in Public Key Cryptosystems (PKC 2001), pages 13-15, Cheju Island, Korea, February 2001.
[21] Ivan Damgård, Valerio Pastro, Nigel P. Smart, and Sarah Zakarias. Multiparty computation from somewhat homomorphic encryption. In CRYPTO, pages 643-662, 2012.
[22] Dana Dachman-Soled, Tal Malkin, Mariana Raykova, and Moti Yung. Efficient robust private set intersection. In Michel Abdalla, David Pointcheval, Pierre-Alain Fouque, and Damien Vergnaud, editors, ACNS, volume 5536 of Lecture Notes in Computer Science, pages 125-142, 2009.
[23] Martin Dietzfelbinger and Philipp Woelfel. Almost random graphs with simple hash functions. In Proceedings of the 35th Annual ACM Symposium on Theory of Computing, pages 629-638, 2003.
[24] Ivan Damgård and Sarah Zakarias. Constant-overhead secure computation of boolean circuits using preprocessing. In TCC, pages 621-641, 2013.
[25] Alexandre Evfimievski, Johannes Gehrke, and Ramakrishnan Srikant. Limiting privacy breaches in privacy preserving data mining. In Proc. 22nd ACM Symposium on Principles of Database Systems (PODS 2003), pages 211-222, San Diego, CA, June 2003.
[26] Taher ElGamal. A public key cryptosystem and a signature scheme based on discrete logarithms. IEEE Transactions on Information Theory, 31(4):469-472, 1985.
[27] Michael J. Freedman, Yuval Ishai, Benny Pinkas, and Omer Reingold. Keyword search and oblivious pseudorandom functions. In Joe Kilian, editor, TCC, volume 3378 of Lecture Notes in Computer Science, pages 303-324. Springer, 2005.
[28] Michael J. Freedman, Kobbi Nissim, and Benny Pinkas. Efficient private matching and set intersection. In Christian Cachin and Jan Camenisch, editors, Advances in Cryptology-EUROCRYPT 2004, volume 3027 of LNCS, pages 1-19. Springer-Verlag, 2-6 May 2004.
[29] Ronald Fagin, Moni Naor, and Peter Winkler. Comparing information without leaking it. Communications of the ACM, 39(5):77-85, 1996.
[30] Oded Goldreich and Ariel Kahan. How to construct constant-round zero-knowledge proof systems for NP. Journal of Cryptology, 9(3):167-189, 1996.
[31] Shafi Goldwasser and Leonid A. Levin. Fair computation of general functions in presence of immoral majority. CRYPTO, 537:77-93, 1990.
[32] GMP. GNU Multiple Precision Arithmetic Library. gmplib.org, 2009.
[33] Oded Goldreich, Silvio Micali, and Avi Wigderson. How to play any mental game or a completeness theorem for protocols with honest majority. In Proc. Nineteenth Annual ACM Symposium on Theory of Computing, pages 218-229, New York City, 25-27 May 1987.
[34] Oded Goldreich. Foundations of cryptography: Basic applications. Cambridge Univ Pr, 2004.
[35] Bernardo A. Huberman, Matt Franklin, and Tad Hogg. Enhancing privacy and trust in electronic communities. In Proc. ACM Conference on Electronic Commerce, pages 78-86, Denver, Colorado, November 1999.
[36] Carmit Hazay and Yehuda Lindell. Efficient protocols for set intersection and pattern matching with security against malicious and covert adversaries. In TCC, pages 155-175, 2008.
[37] Carmit Hazay, Gert Læssøe Mikkelsen, Tal Rabin, and Tomas Toft. Efficient rsa key generation and threshold paillier in the two-party setting. In CT-RSA, pages 313-331, 2012.
[38] Carmit Hazay and Kobbi Nissim. Efficient set operations in the presence of malicious adversaries. In Public Key Cryptography, pages 312-331, 2010.
[39] Russell Impagliazzo and Steven Rudich. Limits on the provable consequences of one-way permutations. In Proc. 21st Annual ACM Symposium on Theory of Computing, pages 44-61, Seattle, Washington, May 1989.
[40] Stanislaw Jarecki and Xiaomin Liu. Efficient oblivious pseudorandom function with applications to adaptive ot and secure computation of set intersection. TCC, 5444:577-594, 2009.
[41] Stanislaw Jarecki and Xiaomin Liu. Fast secure computation of set intersection. SCN, 6280:418-435, 2010.
[42] Stanislaw Jarecki and Vitaly Shmatikov. Efficient two-party secure computation on committed inputs. EUROCRYPT, 4515:97-114, 2007.
[43] Adam Kirsch, Michael Mitzenmacher, and Udi Wieder. More robust hashing: Cuckoo hashing with a stash. In Proceedings of the 16th annual European symposium on Algorithms, pages 611-622. Springer, 2008.
[44] Efficient password-authenticated key exchange using human-memorable passwords. In Advances in Cryptology - EUROCRYPT 2001, Innsbruck, Austria, May 6-10, 2001, pages 475-494, 2001.
[45] Bala Kalyanasundaram and Georg Schnitger. The probabilistic communication complexity of set intersection. SIAM J. Discrete Mathematics, 5(4):545-557, 1992.
[46] Lea Kissner and Dawn Song. Private and threshold set-intersection. In Proceedings of CRYPTO '05, August 2005.
[47] Helger Lipmaa. Verifiable homomorphic oblivious transfer and private equality test. In Advances in Cryptology-ASIACRYPT 2003, pages 416-433, Taipei, Taiwan, November 2003.
[48] Sven Laur and Helger Lipmaa. A new protocol for conditional disclosure of secrets and its applications. In $A C N S$, pages 207-225, 2007.
[49] Yehuda Lindell and Benny Pinkas. A proof of security of yaos protocol for two-party computation. Journal of Cryptology, 22(2):161-188, 2009.
[50] Yehuda Lindell and Benny Pinkas. Secure two-party computation via cut-and-choose oblivious transfer. J. Cryptology, 25(4):680-722, 2012.
[51] David Mazières. A toolkit for user-level file systems. In USENIX Technical Conference, June 2001.
[52] Silvio Micali and Phillip Rogaway. Privacy preserving data mining. Unpublished manuscript, 576: 392-404, 1991.
[53] Moni Naor and Benny Pinkas. Oblivious transfer and polynomial evaluation. In Proc. 31st Annual ACM Symposium on Theory of Computing, pages 245-254, Atlanta, Georgia, May 1999.
[54] Moni Naor and Benny Pinkas. Efficient oblivious transfer protocols. In SIAM Symposium on Discrete Algorithms (SODA), pages 448-457, Washington, D.C., January 2001.
[55] Pascal Paillier. Public-key cryptosystems based on composite degree residuosity classes. In Advances in Cryptology-EUROCRYPT '99, pages 223-238, Prague, Czech Republic, May 1999.
[56] Torben Pryds Pedersen. Non-interactive and information-theoretic secure verifiable secret sharing. In J. Feigenbaum, editor, Advances in Cryptology-CRYPTO '91, volume 576 of LNCS, pages 129-140. Springer-Verlag, 1992, 11-15 August 1991.
[57] Anna Pagh and Rasmus Pagh. Uniform hashing in constant time and optimal space. SIAM J. Comput., 38(1):85-96, 2008.
[58] Rasmus Pagh and Flemming Friche Rodler. Cuckoo hashing. J. Algorithms, 51(2):122-144, 2004.
[59] Mihai Patrascu and Mikkel Thorup. The power of simple tabulation hashing. J. ACM, 59(3):14, 2012.
[60] Alexander A. Razborov. Application of matrix methods to the theory of lower bounds in computational complexity. Combinatorica, 10(1):81-93, 1990.
[61] Martin Raab and Angelika Steger. Balls into Bins - A Simple and Tight Analysis. Randomization and Approximation Techniques in Computer Science, pages 159-170, 1998.
[62] Berthold Vöcking. How asymmetry helps load balancing. Journal of the ACM (JACM), 50:568-589, July 2003.
[63] Udi Wieder. Balanced allocations with heterogenous bins. In Proceedings of the nineteenth annual ACM symposium on Parallel algorithms and architectures, page 193. ACM, 2007.
[64] Philipp Woelfel. Asymmetric balanced allocation with simple hash functions. In Proceedings of the seventeenth annual ACM-SIAM symposium on Discrete algorithm, SODA '06, pages 424-433, 2006.
[65] Andrew C. Yao. Protocols for secure computations. In 23rd Annual Symposium on Foundations of Computer Science, pages 160-164, Chicago, Illinois, 3-5 November 1982. IEEE.


[^0]:    * Research partially supported by a grant from the Israel Ministry of Science and Technology (Grant No. 3-10883).
    ${ }^{\dagger}$ Research partially supported by the Israel Science Foundation (Grant No. 860/06).
    $\ddagger$ Research partially supported by the European Unions 7th Framework Program (FP7/2007-2013) under Grant Agreement No. 609611 (PRACTICE) and by a grant from the Israel Ministry of Science and Technology (Grant No. 3-9094).

[^1]:    ${ }^{1}$ The reduction is to a domain of size $3 n$ in order to ensure that the inputs to the set intersection problem are always of size $\Theta(n)$. This simplifies the description of the result that is proved by the reduction.

[^2]:    ${ }^{2}$ The plaintext and ciphertext spaces may depend on $p k$; we leave this implicit.

[^3]:    ${ }^{3}$ W.l.o.g., we assume $X, Y \subseteq\{0,1\}^{p(k)}$ for some polynomial $p(\cdot)$ such that $2^{p(k)}$ is super-polynomial in k.

[^4]:    ${ }^{4}$ This construction can be considered a generalization of the oblivious transfer protocols of [5,47,54]. In those, a chooser retrieving item $i$ sends to the sender a predicate which is 0 if and only if $i=j$, where $j \in[N]$.

[^5]:    ${ }^{5}$ By setting $B=m_{1} \log \log m_{1} / \log m_{1}$, we can make the error probability negligible in $m_{1}$. However, any actual implementation will have to examine the exact value of $B$ which results in a sufficiently small error probability for the input sizes that are expected. As for theoretical analysis, the subsequent construction, based on balanced allocation hashing, presents a negligible error probability.

[^6]:    ${ }^{6}$ A constant factor improvement is achieved using the Always Go Left scheme in [62] where $h_{0}:\{0,1\}^{*} \rightarrow$ $\left[1, \ldots, \frac{B}{2}\right], h_{1}:\{0,1\}^{*} \rightarrow\left[\frac{B}{2}+1, \ldots, B\right]$. In that scheme, an element $x$ is inserted into the less occupied bin from $\left\{h_{0}(x), h_{1}(x)\right\}$; in case of a tie, $x$ is inserted into $h_{0}(x)$.

[^7]:    ${ }^{7}$ Any implementation of a hashing scheme must replace the idealized random hash function (that is used for the analysis) with an actual construction of a hash function that works well in practice. See, e.g., [59] and related work.

[^8]:    ${ }^{8}$ We consider an extension of the semantic security game (cf. Definition 2.6), where a distinguisher $D_{\mathrm{E}}$ is given a public key $p k$, outputs two vectors of plaintexts $\bar{m}_{0}, \bar{m}_{1}$, and receives back a vector of ciphertexts $\bar{c}$ comprised of the encryptions of $\bar{m}_{b}$ under $E_{p k}$, where $b \leftarrow_{R}\{0,1\} . D_{\mathrm{E}}$ then outputs a bit $b^{\prime}$, and we say that it distinguishes successfully in this game if $b^{\prime}=b$.

[^9]:    ${ }^{9}$ Learning $s k$ allows $P_{1}$ to efficiently decrypt messages it receives from $P_{2}$. Otherwise, $P_{1}$ and $P_{2}$ would have to engage in a protocol for a joint decryption.

[^10]:    ${ }^{10}$ The original construction of [20] is presented in the honest verifier setting. Deriving a statistical zeroknowledge proof can be achieved by instantiating Pedersen's commitment scheme [56] with the technique of Goldreich and Kahan [30]. The analysis of 15 exponentiations already takes into account this adjustment.

[^11]:    ${ }^{11}$ See previous footnote.
    ${ }^{12}$ We will use the convention that the degree of a polynomial $Q_{i}(\cdot)$ can be chosen to be any integer $j^{\prime}$ such that $Q_{i, j}=0$ for all $j \geq j^{\prime}$; hence, equality with $m$ can always be achieved.

[^12]:    ${ }^{13}$ The proof can be easily computed since both parties can compute an encryption $z^{\prime}$ of $1-Z_{i, j}$, where $P$ can recover randomness $\tilde{r}_{i, j}^{\prime}$ that is consistent with this encryption. We also assume that $P$ and $V$ agree on an encryption of 1 , for which both know the randomness.

