

Oblivious Polynomial Evaluation and Secure Set-Intersection from Algebraic PRFs

Carmit Hazay*

Abstract

In this paper we study the two fundamental functionalities *oblivious polynomial evaluation in the exponent* and *set-intersection*, and introduce a new technique for designing efficient secure protocols for these problems (and others). Our starting point is the [BGV11] technique (CRYPTO 2011) for verifiable delegation of polynomial evaluations, using *algebraic PRFs*. We use this tool, that is useful to achieve *verifiability* in the *outsourced setting*, in order to achieve *privacy* in the *standard two-party* setting. Our results imply new simple and efficient oblivious polynomial evaluation (OPE) protocols. We further show that our OPE protocols are readily used for secure set-intersection, implying much simpler protocols in the plain model. As a side result, we demonstrate the usefulness of algebraic PRFs for various search functionalities, such as keyword search and oblivious transfer with adaptive queries. Our protocols are secure under full simulation-based definitions in the presence of malicious adversaries.

Keywords: Efficient Secure Computation, Oblivious Polynomial Evaluation, Secure Set-Intersection, Committed Oblivious PRF

*Faculty of Engineering, Bar-Ilan University, Israel. Email: carmit.hazay@biu.ac.il. Research partially supported by a grant from the Israel Ministry of Science and Technology (grant No. 3-10883).
© IACR 2015. This article is the final version submitted by the author to TCC 2015.

1 Introduction

Efficient secure two-party computation. Secure two-party computation enables two parties to mutually run a protocol that computes some function f on their private inputs, while preserving a number of security properties. Two of the most important properties are privacy and correctness. The former implies data confidentiality, namely, nothing leaks by the protocol execution but the computed output. The latter requirement implies that the protocol enforces the integrity of the computations made by the parties, namely, honest parties learn the correct output. Feasibility results are well established [Yao86, GMW87, MR91, Bea91], proving that any efficient functionality can be securely computed under full simulation-based definitions (following the ideal/real paradigm). Security is typically proven with respect to two adversarial models: the semi-honest model (where the adversary follows the instructions of the protocol but tries to learn more than it should from the protocol transcript), and the malicious model (where the adversary follows an arbitrary polynomial-time strategy), and feasibility holds in the presence of both types of attacks.

Following these works, many constructions focused on improving the *efficiency* of the computational and communication costs. Conceptually, this line of works can be split into two sub-lines: (1) Improved generic protocols that compute any boolean/arithmetic circuit; see [ST04, IPS08, NO09, LOP11, LP11, BDOZ11, DPSZ12, NNOB12] for just a few examples. (2) Protocols for concrete functionalities. In the latter approach attention is given to constructing efficient protocols for specific functions while exploiting their internal structure. This approach has been proven useful for many different functions in both the semi-honest and malicious settings. Notable examples are calculating the k th ranked element [AMP04], pattern matching and related search problems [HT10, Ver11], set-intersection [JL09, HN12] and oblivious pseudorandom function (PRF) evaluation [FIPR05].

In this paper we study the two fundamental functionalities *oblivious polynomial evaluation in the exponent* and *set-intersection* and introduce a new technique for designing efficient secure protocols for these problems in the presence of semi-honest and malicious attacks with simulation-based security proofs. We further demonstrate that our technique is useful for various search functionalities.

Algebraic PRFs. Informally, an algebraic pseudorandom function (PRF) is a PRF with a range that forms an Abelian group such that group operations are efficiently computable. In addition, certain algebraic operations on these outputs can be computed significantly more efficiently if one possesses the key of the pseudorandom function that was used to generate them. This property is denoted by *closed form efficiency* and allows to compute a batch of l PRF values much more efficiently than by computing the l values separately and then combining them. Algebraic PRFs were exploited in [BGV11] to achieve faster verifiable polynomial evaluations (in the exponent). Specifically, in their setting, a client outsources a d -degree polynomial to an untrusted server together with some authenticating information, while the client stores a short secret key. Next, when the client provides an input for this polynomial the server computes the result and an authentication message that allows the client to verify this computation in sub-linear time in d .

More concretely, let $Q(\cdot) = (q_0, \dots, q_d)$ be the polynomial stored on the server in the clear. Then the client additionally stores a vector of group elements $\{g^{aq_i+r_i}\}_{i=0}^d$ where $a \leftarrow \mathbb{Z}_p$ and p is a prime, and r_i is the i th coefficient of a polynomial $R(\cdot)$ of the same degree as $Q(\cdot)$. Then for every client's input t the server returns $y = Q(t)$ and $u = g^{aQ(t)+R(t)}$ and the client accepts u if and only if $u = g^{ay+R(t)}$. Interestingly, in case $g^{r_i} = \text{PRF}_K(i)$, where PRF is an algebraic PRF, the closed form efficiency property enables the client to compute the value $g^{R(t)}$ in sub-linear time in d . Stated differently, verifiability is achieved by viewing $g^{aq_i+r_i}$ as a (one-time) message authentication code (MAC) for g^{q_i} where batch verification of multiple MACs can be computed more efficiently than verifying each MAC separately.

In this work we demonstrate the usefulness of algebraic PRFs for various two-party problems by designing secure protocols based on this primitive. In particular, we modify the way [BGV11] use algebraic

PRFs so that instead of achieving *verifiability* in the *outsourced setting*, we achieve *privacy* in the *standard two-party* setting. It is worth noting that although the main focus of [BGV11] is correctness, they do discuss how to achieve one-sided privacy by encrypting the coefficients of the polynomial (since the polynomial must be specified explicitly). Nevertheless, it is not clear how to maintain the privacy of the input to the polynomial in their protocol. In this work, we use algebraic PRFs to mask the polynomial in a different way that does not allow the verifiability of the polynomial evaluation but allows the extractability of the polynomial more easily, and demonstrate an alternative way to achieve correctness. We focus our attention on the *plain model* where no trusted setup is required.

Oblivious polynomial evaluation. The oblivious polynomial evaluation (OPE) functionality is an important functionality in the field of secure two-party computation. It considers a setting where party P_0 holds a polynomial $Q(\cdot)$ and party P_1 holds an element t , and the goal is that P_1 obtains $Q(t)$ and nothing else while P_0 learns nothing. OPE has proven to be a useful building block and can be used to solve numerous cryptographic problems; e.g., secure equality of strings, set-intersection, approximation of a Taylor series, RSA key generation, oblivious keyword search, set membership, data entanglement and more [Gil99, LP02, FNP04, FIPR05, NP06, ADDV12].

Despite its broad applicability the study of OPE was demonstrated using only few concrete secure protocols, initiated in [NP99] and further continued in [CL05, ZB05, HL09]. In particular, the only protocol with a complete simulation-based proof in the presence of malicious attacks is the protocol in [HL09]. This protocol evaluates a d -degree polynomial over a composite order group \mathbb{Z}_N with $O(sd)$ modular exponentiations, where N is an RSA composite and s is a statistical security parameter.

The general (and currently the most practical) approach of [DPSZ12, DKL⁺13] for arithmetic circuits follows the preprocessing model: in an offline phase some shared randomness is generated independently of the function and the inputs; in an online phase the actual secure computation is performed. One of the main advantages of these protocols is that the basic operations are almost as cheap as those used in the passively secure protocols. To get good performance, these protocols use the somewhat-homomorphic SIMD approach that handles many values in parallel in a single ciphertext, and thus more applicable for large degree polynomials. Similarly, protocols for Boolean circuits apply the cut-and-choose technique which requires to repeat the computation s times in order to prevent cheating except with probability 2^{-s} [Lin13].

In some applications such as password-based authenticate key exchange protocols or when sampling an element from a d -wise independence space, the polynomial degree is typically small and even a constant. In these cases, our protocols have clear benefits since they are much simpler, efficient and easily implementable.

Secure set-intersection. In the set-intersection problem parties P_0, P_1 , holding input sets X, Y of sizes m_X and m_Y , respectively, wish to compute $X \cap Y$. This problem has been intensively studied by researchers in the last few years mainly due to its potential applications for dating services, datamining, recommendation systems, law enforcement and more; see [FNP04, KS05, DSMRY09, JL09, JL10, HL10a, HN12] for a few examples. For instance, consider two security agencies that wish to compare their lists of suspects without revealing their contents, or an airline company that would like to check its list of passengers against the list of people that are not allowed to go abroad.

Two common approaches are known to solve this problem securely in the plain model: **(1)** oblivious polynomial evaluation and **(2)** committed oblivious PRF evaluation. In the former approach party P_0 computes a polynomial $Q(\cdot)$ such that $Q(x) = 0$ for all $x \in X$. This polynomial is then encrypted using homomorphic encryption and sent to P_1 , that computes the encryption of $r_y \cdot Q(y) + y$ for all $y \in Y$, and using fresh randomness r_y . This approach (or a variant of it) was taken in [FNP04, KS05, DSMRY09, HN12].

The second approach uses a secure implementation of oblivious pseudorandom function evaluation. Namely, P_0 chooses a PRF key K and computes the set $\text{PRF}_X = \{\text{PRF}_K(x)\}_{x \in X}$. The parties then

execute an oblivious PRF protocol where P_0 inputs K , whereas P_1 inputs the set Y and learns the set $\text{PRF}_Y = \{\text{PRF}_K(y)\}_{y \in Y}$. Finally, P_0 sends the set PRF_X to P_1 that computes $\text{PRF}_X \cap \text{PRF}_Y$ and extracts the actual intersection. This idea was introduced in [FIPR05] and further used in [HL10a, JL09, JL10]. Other solutions in the random oracle model such as [CT10, CKT10, ACT10] take a different approach by applying the random oracle on (one of) the sets members, or apply oblivious transfer extension [DCW13].

In a recent result [PSZ14], the authors overview exiting solutions for set-intersection in the semi-honest setting and compare their efficiency. One of their conclusions is that OPE-based approaches are inferior to oblivious-transfer extension based approaches. It is an interesting question to test whether this conclusion also for the case for the malicious setting as well.

To the best of our knowledge, the most efficient protocol in the malicious plain model that does not require a trusted setup or rely on non-standard assumptions is the protocol of [HN12] that incurs computation of $O(m_X + m_Y \log(m_X + m_Y))$ modular exponentiations. A more efficient protocol with $O(m_X + m_Y)$ communication and computational costs was introduced by [JL09] in the common reference string (CRS) model (where the CRS includes a safe RSA composite that determines the group order and implies high overhead when mutually produced). Another drawback of this protocol is that its security proof runs an exhaustive search on the input domain of the PRF in order to extract P_0 's input. This implies that the proof only works for small domain PRFs and that the complexity of the simulator grows linearly with the size of the PRF's input domain.

Committed oblivious PRF evaluation. The oblivious PRF evaluation functionality \mathcal{F}_{PRF} that *obviously* evaluates a PRF is defined by $(K, x) \mapsto (-, \text{PRF}_K(x))$. This functionality is very important in the context of secure computation since it essentially implements a random oracle. That is, the party with the PRF key, say P_0 , mimics the random oracle role via interaction. Therefore, if the protocol that realizes \mathcal{F}_{PRF} is simulation-based secure then both desirable properties of a random oracle, *programmability* and *observability*, can be achieved by this protocol. First, since the simulator can force any output for a corrupted P_1 , it essentially programs the function's output. In addition, it can also observe (via extraction) the input to the functionality. Nevertheless, the usefulness of oblivious PRF evaluation is reflected via an additional property of committed key that implies that the *same key* is used for multiple PRF evaluations.

Committed oblivious PRF (CPRF) evaluation has been used to compute secure set-intersection [JL09, HL10a], oblivious transfer with adaptive queries [FIPR05], keyword search [FIPR05], pattern matching [HL10a, FHV13] and more. It is therefore highly important to design efficient protocols for this functionality. Current implementations of the [NR97] algebraic PRF, discussed in this paper, employ an oblivious transfer protocol for each input bit [FIPR05, HL10a] and are only secure for a single PRF evaluation. Consequently, the protocol of [HL10a] does not achieve full security against malicious adversaries. In addition, the protocol from [JL09] (that implements a variant of the [DY05] PRF) requires a trusted setup of a safe RSA composite and suffers from the drawbacks specified above.

1.1 Our Results

In this paper we use algebraic PRFs to design alternative simple and efficient protocols for polynomial evaluation, set-intersection, committed oblivious PRF evaluation and search problems. Below, we demonstrate the broad usefulness of our technique for the different functionalities.

Oblivious polynomial evaluation (Section 3). We present secure protocols in the plain model for OPE *in the exponent* with simulation-based security against semi-honest and malicious attacks. We stress that evaluating a polynomial in the exponent has strong applicability in the context of set membership where the goal is to privately verify membership in some secret set, as well as achieving d -wise independence. We use

algebraic PRFs to build simple two-phases OPE protocols as follows. In the first phase party P_0 , holding the polynomial $g^{Q(\cdot)}$, publishes its masked polynomial $g^{Q(\cdot)+R(\cdot)}$ where the set $g^{R(\cdot)}$ is determined by an algebraic PRF. Next, P_1 locally computes $g^{Q(t)+R(t)}$ and the parties run an unmasking secure computation for obliviously evaluating $g^{R(t)}$ for P_1 .

The efficiency of the latter phase is dominated by the overhead of the closed form efficiency property of the specific PRF. In this work, we consider two PRF implementations used by [BGV11]: (1) a PRF with security under the strong-DDH assumption. (2) The Naor-Reingold PRF [NR97] with security under the DDH assumption. More concretely, the efficiency of our protocols is only $d+1$ modular exponentiations for the first phase of sending the masked polynomial, and $d+1+O(1)$ (resp. $O(\log d)$) modular exponentiations for the second phase of obliviously evaluating the pseudorandom polynomial under the strong-DDH (resp. DDH) assumption. For simplicity, we only consider univariate polynomials. Our technique can be applied for multivariate polynomials as well (with total degree d or of degree d in each variable); see [BGV11] for further details. To the best of our knowledge, our protocols are the first to obliviously evaluate both univariate and multivariate polynomials that efficiently.

Secure set-intersection (Section 4). In this work we demonstrate that algebraic PRFs are useful for both approaches of OPE and committed oblivious PRF that enable to design set-intersection protocols. We first show that our protocols for OPE readily induce secure protocols for set-intersection. That is, first P_0 encodes the set X by a polynomial $g^{Q(\cdot)}$ as specified above, and masks it. Next, for each $y \in Y$ party P_1 verifies whether the masked polynomial evaluation of y equals the evaluation $g^{R(y)}$, and concludes whether the element is in the intersection. We stress that this naive approach requires a multiplicative overhead (in the sets sizes) since for each element in its input Y , P_1 needs to evaluate a polynomial of degree m_X . To reduce the computational overhead, Freedman et al. [FNP04] introduced a balanced allocation scheme [ABKU99] into their protocol that splits the elements into $\mathcal{B} = \frac{m_X}{\log \log m_X}$ bins, with maximum number of $\mathcal{M} = O(m_X/\mathcal{B} + \log \log \mathcal{B}) = O(\log \log m_X)$ elements in each bin. In that case, the elements mapped by P_0 to a certain bin must only be compared to those mapped by P_1 to the same bin. Therefore, P_1 should only evaluate an \mathcal{M} -degree polynomial for each $y \in Y$, rather than a polynomial of degree m_X . Nevertheless, their solution with hash functions is only applicable in the semi-honest setting. Following that, Hazay and Nissim [HN12] introduced a maliciously secure protocol which implies the computation of $O(m_X + m_Y \log(m_X + m_Y))$ modular exponentiations. Their construction is fairly complicated and combines both approaches of OPE and oblivious PRF evaluation.

We introduce the hashing technique into our constructions and provide a generic description that can be instantiated with different hash functions. Our protocols are far less complicated and maintain a modular description. Specifically, we devise an alternative zero-knowledge proof for verifying the correctness of the hashed polynomials while exploiting the algebraic properties of the PRF. Under the strong-DDH assumption our protocol matches the communication overhead of the protocol from [JL09] (that also relies on a dynamic hardness assumption) and implies the computation of $O(m_X + m_Y \log \log m_X)$ exponentiations, with the benefits that it operates over prime order groups, it does not require a trusted setup and the proof complexity does not depend on the PRF's input domain size. Under the DDH assumption our protocol, using hash functions, implies the computation of $O(m_X + m_Y \log m_X)$ exponentiations which improves the overhead of the [HN12] protocol. Next we show that algebraic PRFs are useful for applications that rely on committed oblivious PRF evaluation. Our results for set-intersection are summarized in Table 1.

Committed oblivious PRF evaluation (Section 5). Observing that the batch computation for l PRF values $\text{PRF}'_K(x) = \prod_{i=0}^l [\text{PRF}_K(i)]^{x^i}$ is a PRF as well (by fixing l properly), we derive new PRF constructions in prime order groups and more interestingly, simple committed oblivious PRF evaluation protocols. Our strong-DDH based PRF requires constant overhead, and our DDH-based protocol is the *first* committed

Reference	Modeling	Hardness Assumption	Overhead (Number of Exp.)
[JL09]	CRS of a safe prime	Decisional d-DHI	$O(m_X + m_Y)$
[HN12]	plain model	DDH	$O(m_X + m_Y \log(m_X + m_Y))$
[DCW13]	random oracle	random oracle	$O(n)$, where n is sec. parameter
This Work – OPE	plain model	d-strong DDH	$O(m_X + m_Y \log \log m_X)$
This Work – OPE	plain model	DDH	$O(m_X + m_Y \log m_X)$
This Work – CPRF	plain model	d-strong DDH	$O(m_X + m_Y)$
This Work – CPRF	plain model	DDH	$O((m_X + m_Y) \log(m_X + m_Y))$

Table 1: Comparisons with secure set-intersection constructions. We highlight the constructions with the best performance under each assumption.

oblivious PRF implementation for the [NR97] function. Our protocols using committed oblivious PRF imply set-intersection protocols with $O(m_X + m_Y)$ costs under the strong-DDH assumption and $((m_X + m_Y) \log(m_X + m_Y))$ communication and computation costs under the DDH assumption, where the former analysis matches the overhead from [JL09]. In particular, plugging-in our protocols for committed oblivious PRF evaluation in the protocols cited above implies malicious security fairly immediately. Finally, we note that committed oblivious PRF evaluation is also useful for search functionalities that support database search and data retrievals, such as in keyword search and oblivious transfer with adaptive queries.

2 Preliminaries

2.1 Basic Notations

We denote the security parameter by n . We say that a function $\mu : \mathbb{N} \rightarrow \mathbb{N}$ is *negligible* if for every positive polynomial $p(\cdot)$ and all sufficiently large n it holds that $\mu(n) < \frac{1}{p(n)}$. We use the abbreviation PPT to denote probabilistic polynomial-time. We further denote by $a \leftarrow A$ the random sampling of a from a distribution A , by $[d]$ the set of elements $(1, \dots, d)$ and by $[0, d]$ the set of elements $(0, \dots, d)$.

We now specify the definition of computationally indistinguishable.

Definition 2.1 Let $X = \{X(a, n)\}_{a \in \{0,1\}^*, n \in \mathbb{N}}$ and $Y = \{Y(a, n)\}_{a \in \{0,1\}^*, n \in \mathbb{N}}$ be two distribution ensembles. We say that X and Y are computationally indistinguishable, denoted $X \stackrel{c}{\approx} Y$, if for every PPT machine D , every $a \in \{0, 1\}^*$, every positive polynomial $p(\cdot)$ and all sufficiently large n :

$$|\Pr [D(X(a, n), 1^n) = 1] - \Pr [D(Y(a, n), 1^n) = 1]| < \frac{1}{p(n)}.$$

We define a d -degree polynomial $Q(\cdot)$ by its set of coefficients (q_0, \dots, q_d) , or simply write $Q(x) = q_0 + q_1x + \dots + q_dx^d$. Typically, these coefficients will be picked from \mathbb{Z}_p for a prime p . We further write $g^{Q(\cdot)}$ to denote the coefficients of $Q(\cdot)$ in the exponent of a generator g of a multiplicative group \mathbb{G} of prime order p .

2.2 Hardness Assumptions

All our constructions rely on the following hardness assumptions.

DECISIONAL DIFFIE-HELLMAN. The classic decisional Diffie-Hellman assumption is stated as follows.

Definition 2.2 (DDH) We say that the decisional Diffie-Hellman (DDH) problem is hard relative to \mathcal{G} , if for any PPT distinguisher D there exists a negligible function negl such that

$$\left| \Pr [D(\mathbb{G}, p, g, g^x, g^y, g^z) = 1] - \Pr [D(\mathbb{G}, p, g, g^x, g^y, g^{xy}) = 1] \right| \leq \text{negl}(n),$$

where $(\mathbb{G}, p, g) \leftarrow \mathcal{G}(1^n)$ and the probabilities are taken over the choices of $x, y, z \leftarrow_R \mathbb{Z}_p$.

STRONG DIFFIE-HELLMAN. The strong Diffie-Hellman family of assumptions is parameterized by a parameter d , and states that any PPT distinguisher cannot distinguish elements $g, g^x, g^{x^2}, \dots, g^{x^d}$ from a sequence of random elements in a prime order group \mathbb{G} . More formally,

Definition 2.3 (d-SDDH) We say that the d -strong Diffie-Hellman problem is hard relative to \mathcal{G} if for any PPT distinguisher D there exists a negligible function negl such that

$$\left| \Pr [D(\mathbb{G}, p, g, g^x, g^{x^2}, \dots, g^{x^d}) = 1] - \Pr [D(\mathbb{G}, p, g, g^{x_1}, g^{x_2}, \dots, g^{x_d}) = 1] \right| \leq \text{negl}(n),$$

where $(\mathbb{G}, p, g) \leftarrow \mathcal{G}(1^n)$ and the probabilities are taken over the choices of $x, x_1, \dots, x_d \leftarrow_R \mathbb{Z}_p$.

We refer the reader to [BGV11] for further discussions and related work regarding the strong Diffie-Hellman assumption. Fixing the parameter d appropriately enables [BGV11] to evaluate a special class of polynomials of degree d in constant time. Specifically, there is a tight correlation between the value of d and the efficiency of their construction; the stronger the assumption is, the larger the class of polynomials that can be evaluated in constant time. In what follows, we also fix the parameter d in this assumption according to the degree of evaluated polynomial.

2.3 Public Key Encryption (PKE)

We specify the definitions of public key encryption and IND-CPA.

Definition 2.4 (PKE) We say that $\Pi = (\text{Gen}, \text{Enc}, \text{Dec})$ is a public key encryption scheme if $\text{Gen}, \text{Enc}, \text{Dec}$ are polynomial-time algorithms specified as follows:

- Gen , given a security parameter n (in unary), outputs keys (PK, SK) , where PK is a public key and SK is a secret key. We denote this by $(\text{PK}, \text{SK}) \leftarrow \text{Gen}(1^n)$.
- Enc , given the public key PK and a plaintext message m , outputs a ciphertext c encrypting m . We denote this by $c \leftarrow \text{Enc}_{\text{PK}}(m)$; and when emphasizing the randomness r used for encryption, we denote this by $c \leftarrow \text{Enc}_{\text{PK}}(m; r)$.
- Dec , given the public key PK , secret key SK and a ciphertext c , outputs a plaintext message m s.t. there exists randomness r for which $c = \text{Enc}_{\text{PK}}(m; r)$ (or \perp if no such message exists). We denote this by $m \leftarrow \text{Dec}_{\text{PK}, \text{SK}}(c)$.

For a public key encryption scheme $\Pi = (\text{Gen}, \text{Enc}, \text{Dec})$ and a non-uniform adversary $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$, we consider the following IND-CPA game:

$$\begin{aligned} (\text{PK}, \text{SK}) &\leftarrow \text{Gen}(1^n). \\ (m_0, m_1, \text{history}) &\leftarrow \mathcal{A}_1(\text{PK}), \text{ s.t. } |m_0| = |m_1|. \\ c &\leftarrow \text{Enc}_{\text{PK}}(m_b), \text{ where } b \leftarrow_R \{0, 1\}. \\ b' &\leftarrow \mathcal{A}_2(c, \text{history}). \\ \mathcal{A} \text{ wins if } &b' = b. \end{aligned}$$

Denote by $\text{Adv}_{\Pi, \mathcal{A}}(n)$ the probability that \mathcal{A} wins the IND-CPA game.

Definition 2.5 (IND-CPA) A public key encryption scheme $\Pi = (\text{Gen}, \text{Enc}, \text{Dec})$ has indistinguishable encryptions under chosen plaintext attacks (IND-CPA), if for every non-uniform adversary $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$ there exists a negligible function negl such that $\text{Adv}_{\Pi, \mathcal{A}}(n) \leq \frac{1}{2} + \text{negl}(n)$.

2.3.1 Additively Homomorphic PKE

A public key encryption scheme is additively homomorphic if given two ciphertexts $c_1 = \text{Enc}_{\text{PK}}(m_1; r_1)$ and $c_2 = \text{Enc}_{\text{PK}}(m_2; r_2)$ it is possible to efficiently compute $\text{Enc}_{\text{PK}}(m_1 + m_2; r)$ with independent r , and without the knowledge of the secret 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 $+$ or \cdot). We abuse notation and use $\text{Enc}_{\text{PK}}(m)$ to denote the random variable induced by $\text{Enc}_{\text{PK}}(m; r)$ where r is chosen uniformly at random. We have the following formal definition,

Definition 2.6 (Homomorphic PKE) We say that a public key encryption scheme $(\text{Gen}, \text{Enc}, \text{Dec})$ is homomorphic if for all k and all (PK, SK) output by $\text{Gen}(1^n)$, it is possible to define groups \mathcal{M}, \mathcal{C} such that:

- The plaintext space is \mathcal{M} , and all ciphertexts output by $\text{Enc}_{\text{PK}}(\cdot)$ are elements of \mathcal{C} .¹
- For every $m_1, m_2 \in \mathcal{M}$ it holds that

$$\{\text{PK}, c_1 = \text{Enc}_{\text{PK}}(m_1), c_1 \cdot \text{Enc}_{\text{PK}}(m_2)\} \equiv \{\text{PK}, \text{Enc}_{\text{PK}}(m_1), \text{Enc}_{\text{PK}}(m_1 + m_2)\}$$

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

Note that any such a scheme supports a multiplication of a plaintext by a scalar.

2.4 The El Gamal PKE

A useful implementation of homomorphic PKE is the El Gamal [Gam85] scheme that has two variations of additive and multiplicative definitions (where the former is only useful for small domains plaintexts). In this paper we exploit the additive variation. Let \mathbb{G} be a group of prime order p in which DDH is hard. Then the public key is a tuple $\text{PK} = \langle \mathbb{G}, p, g, h \rangle$ and the corresponding secret key is $\text{SK} = s$, s.t. $g^s = h$. Encryption is performed by choosing $r \leftarrow \mathbb{Z}_p$ and computing $\text{Enc}_{\text{PK}}(m; r) = \langle g^r, h^r \cdot g^m \rangle$. Decryption of a ciphertext $C = \langle \alpha, \beta \rangle$ is performed by computing $g^m = \beta \cdot \alpha^{-s}$ and then finding m by running an exhaustive search.

2.4.1 Distributed El Gamal

In a distributed scheme, the parties hold shares of the secret key so that the combined key remains a secret. In order to decrypt, each party uses its share to generate an intermediate computation which are eventually combined into the decrypted plaintext. We consider two functionalities: One for securely generating a secret key while keeping it a secret from both parties, whereas the second functionality jointly decrypts a given ciphertext. We denote the key generation functionality by \mathcal{F}_{KEY} , which is defined as follows,

$$(1^n, 1^n) \mapsto \left((\text{PK}, \text{SK}_1), (\text{PK}, \text{SK}_2) \right)$$

¹The plaintext and ciphertext spaces may depend on PK; we leave this implicit.

where $(\text{PK}, \text{SK}) \leftarrow \text{Gen}(1^n)$, and SK_1 and SK_2 are random shares of SK . The decryption functionality \mathcal{F}_{DEC} is defined by,

$$(c, \text{PK}) \mapsto \left((m : c = \text{Enc}_{\text{PK}}(m)), - \right).$$

In El Gamal the parties first agree on a group \mathbb{G} of order p and a generator g . Then, each party P_i picks $s_i \leftarrow \mathbb{Z}_p$ and sends $h_i = g^{s_i}$ to the other. Finally, the parties compute $h = h_1 h_2$ and set $\text{PK} = \langle \mathbb{G}, p, g, h \rangle$. Clearly, the secret key $s = s_1 + s_2$ associated with this public key is correctly shared amongst the parties. In order to ensure correct behavior, the parties must prove knowledge of their s_i by running on (g, h_i) the zero-knowledge proof π_{DL} , specified in Section 2.5. To ensure simulation based security, P_1 must commit to its share first and decommit this commitment after P_2 sends its share. Moreover, decryption of a ciphertext $c = \langle c_1, c_2 \rangle$ follows by computing $c_2 \cdot (c_1^{x_1} \cdot c_1^{x_2})^{-1}$, where each party sends c_i to the power of its share. We denote these protocols by π_{KEY} and π_{DEC} , respectively, and assume that they can be carried out with simulation-based security in the presence of malicious attacks.

2.5 Zero-Knowledge Proofs

To prevent malicious behavior, the parties must demonstrate that they are well-behaved. To achieve this, our protocols utilize zero-knowledge (ZK) proofs of knowledge. Our proofs are Σ -protocols with a constant overhead. A generic efficient technique that enables to transform any Σ -protocol into a zero-knowledge proof of knowledge can be found in [HL10b]. This transformation requires additional 6 exponentiations.

1. π_{DL} , for demonstrating the knowledge of a solution x to a discrete logarithm problem [Sch89].

$$\mathcal{R}_{\text{DL}} = \{ ((\mathbb{G}, g, h), x) \mid h = g^x \}.$$

2. π_{DDH} , for demonstrating that an El Gamal ciphertext is an encryption of zero [CP92].

$$\mathcal{R}_{\text{DDH}} = \{ ((\mathbb{G}, g, h, g_1, h_1), x) \mid g_1 = g^x \wedge h_1 = h^x \}.$$

3. π_{MULT} , for proving that a ciphertext c_2 encrypts a product of two plaintexts values. Namely,

$$\mathcal{R}_{\text{MULT}} = \left\{ ((\mathbb{G}, \text{PK}, c_0, c_1, c_2), (a_0, a_1, r_0, r_1, r_2)) \mid \begin{array}{l} c_i = \text{Enc}_{\text{PK}}(a_i; r_i) \text{ for } i \in \{0, 1\} \wedge \\ c_2 = \text{Enc}_{\text{PK}}(a_0 \cdot a_1; r_2) \end{array} \right\}$$

where multiplication is performed in the corresponding plaintext group. A zero-knowledge proof for the El Gamal PKE, that is based on the Damgård and Jurik technique [DJN10], can be found in [HN12].

4. π_{Eq} , for demonstrating equality of two exponentiations. Namely,

$$\mathcal{R}_{\text{Eq}} = \{ ((\text{PK}, c_1, c'_1, c_2, c'_2), (m, r_1, r_2)) \mid c'_1 = c_1^m \cdot \text{Enc}_{\text{PK}}(0; r_1) \wedge c'_2 = c_2^m \cdot \text{Enc}_{\text{PK}}(0; r_2) \}$$

where exponentiation, as well as multiplication with an encryption of zero, are computed component-wise. A variant of this zero-knowledge proof was presented and discussed in [HMRT12] for Paillier encryption scheme and can be easily extended for this relation as well. We specify the details of this proof for the El Gamal encryption scheme (in a honest verifier format).

Protocol 1 (Zero-knowledge proof of knowledge for \mathcal{R}_{Eq}):

- **Joint statement:** Public key PK and ciphertexts (c_1, c'_1, c_2, c'_2) over group \mathbb{G} of prime order p .

- **Auxiliary input for the prover:** Elements (m, r_1, r_2) such that the conditions within \mathcal{R}_{Eq} are met.
- **The protocol:**
 - (a) P chooses at random $t, s_1, s_2 \leftarrow \mathbb{Z}_p$ and sets $\tilde{c}_1 = c_1^t \cdot \text{Enc}_{\text{PK}}(0; s_1)$ and $\tilde{c}_2 = c_2^t \cdot \text{Enc}_{\text{PK}}(0; s_2)$. That is, let $c_1 = \langle c_1^1, c_1^2 \rangle$ and $c_2 = \langle c_2^1, c_2^2 \rangle$, then $\tilde{c}_1 = \langle (c_1^1)^t \cdot g^{s_1}, (c_1^2)^t \cdot h^{s_1} \rangle$ and $\tilde{c}_2 = \langle (c_2^1)^t \cdot g^{s_2}, (c_2^2)^t \cdot h^{s_2} \rangle$. P sends V the values $\tilde{c}_1 = \langle \tilde{c}_1^1, \tilde{c}_1^2 \rangle, \tilde{c}_2 = \langle \tilde{c}_2^1, \tilde{c}_2^2 \rangle$.
 - (b) V chooses $c \leftarrow \mathbb{Z}_p$ at random and sends it to P .
 - (c) P responds with $cm + t, cr_1 + s_1$ and $cr_2 + s_2$.
 - (d) Upon receiving z_1, z_2 and z_3, V accepts if
 - $\langle (c_1^1)^{z_1} \cdot g^{z_2}, (c_1^2)^{z_1} \cdot h^{z_2} \rangle = \langle (c_1^1)^c \cdot \tilde{c}_1^1, (c_1^2)^c \cdot \tilde{c}_1^2 \rangle$ and
 - $\langle (c_2^1)^{z_1} \cdot g^{z_3}, (c_2^2)^{z_1} \cdot h^{z_3} \rangle = \langle (c_2^1)^c \cdot \tilde{c}_2^1, (c_2^2)^c \cdot \tilde{c}_2^2 \rangle$,
 where $c'_1 = \langle c_1^1, c_1^2 \rangle$ and $c'_2 = \langle c_2^1, c_2^2 \rangle$.

Proposition 2.1 *Protocol 1 is a zero-knowledge honest verifier proof for \mathcal{R}_{Eq} .*

Proof: Completeness is easy to verify since the conditions that the verifier checks are always met. Next, we prove zero-knowledge by designing a simulator \mathcal{S}_{Eq} as follows. \mathcal{S}_{Eq} fixes the last message $z_1, z_2, z_3 \leftarrow \mathbb{Z}_p$ at random and sets

$$\begin{aligned} \tilde{c}_1^1 &= (c_1^1)^{z_1} \cdot g^{z_2} / (c_1^1)^c, & \tilde{c}_1^2 &= (c_1^2)^{z_1} \cdot h^{z_2} / (c_1^2)^c \\ \tilde{c}_2^1 &= (c_2^1)^{z_1} \cdot g^{z_3} / (c_2^1)^c, & \tilde{c}_2^2 &= (c_2^2)^{z_1} \cdot h^{z_3} / (c_2^2)^c. \end{aligned}$$

Note that the simulated \tilde{c}_1^1 equals $(c_1^1)^{z_1 - mc} \cdot g^{z_2 - r_1 c}$ since $c_1^1 = (c_1^1)^m \cdot g^{r_1}$. A similar argument can be made for \tilde{c}_1^2 and \tilde{c}_2 . Finally, soundness holds due to the fact that the proof is a proof of knowledge as well, where two views with the same pair of messages $(\tilde{c}_1, \tilde{c}_2)$ and two challenges c, c' enable to extract m, r_1, r_2 . ■

3 Protocols for Oblivious Polynomial Evaluation

In this section we introduce our new constructions for oblivious polynomial evaluation (OPE) in the exponent, implementing functionality $\mathcal{F}_{\text{OPE}} : (g^{Q(\cdot)}, t) \mapsto (-, g^{Q(t)})$ for $Q(\cdot) = (q_0, \dots, q_d)$. In particular, we assume common knowledge of the public parameters: a multiplicative group \mathbb{G} of order p and a generator g for \mathbb{G} , and that the polynomial coefficients are in \mathbb{Z}_p . In our solution, party P_0 generates these parameters and publishes its masked polynomial $g^{Q(\cdot) + R(\cdot)}$, where the set of values $g^{R(\cdot)}$ is determined by an algebraic PRF that has a closed form efficient computation for univariate polynomials (see Section 3.1). Next, P_1 computes $g^{Q(t) + R(t)}$ and the parties run an unmasking secure computation for obliviously evaluating $g^{R(t)}$ for P_1 . Importantly, the closed form efficiency property of the PRF allows the parties to mutually compute $g^{R(t)}$ in *sub-linear* time in d . Before presenting our OPE constructions we formally define algebraic pseudorandom functions.

3.1 Algebraic Pseudorandom Functions [BGV11]

Algebraic PRFs are PRFs with two additional algebraic properties. First, they map their inputs into some Abelian group, where certain algebraic operations on these outputs can be computed significantly faster if one possesses the PRF key. These properties were exploited in [BGV11] to achieve faster polynomial evaluations (in the exponent), where the coefficients of these polynomials lie in the PRF range. Several constructions, implying different overheads, were introduced in [BGV11]; we focus our attention on their constructions for univariate polynomials. Our protocols can be applied for multivariate polynomials as well (with total degree d or of degree d in each variable). We begin with the formal definition of algebraic PRFs.

Definition 3.1 (Algebraic PRFs) We say that $\mathcal{PRF} = (\text{KeyGen}, \text{PRF}, \text{CFEval})$, is an algebraic PRF if $\text{KeyGen}, \text{PRF}$ are polynomial-time algorithms specified as follows:

- KeyGen , given a security parameter 1^n , and a parameter $m \in \mathbb{N}$ that determines the domain size of the PRF, outputs a pair $(K, \text{param}) \leftarrow \mathcal{K}_n$, where \mathcal{K}_n is the key space for a security parameter n . K is the secret key of the PRF, and param encodes the public parameters.
- PRF , given a key K , public parameters param , and an input $x \in \{0, 1\}^m$, outputs a value $y \in Y$, where Y is some set determined by param .
- In addition, the following properties hold:

Pseudorandomness. We say that \mathcal{PRF} is pseudorandom if for every PPT adversary \mathcal{A} , and every polynomial $m = m(n)$, there exists a negligible function negl such that

$$|\Pr[\mathcal{A}^{\text{PRF}_K(\cdot)}(1^n, \text{param}) = 1] - \Pr[\mathcal{A}^{f_n(\cdot)}(1^n, \text{param}) = 1]| \leq \text{negl}(n),$$

where $(K, \text{param}) \leftarrow \text{KeyGen}(1^n, m)$ and $f_n : \{0, 1\}^m \mapsto Y$ is a random function.

Algebraic. We say that \mathcal{PRF} is algebraic if the range Y of $\text{PRF}_K(\cdot)$ for every $n \in \mathbb{N}$ and $(K, \text{param}) \leftarrow \mathcal{K}_n$ forms an Abelian multiplicative group. We require that the group operation on Y be efficiently computable given param .

Closed form efficiency. Let N be the order of the range sets of PRF for security parameter n . Let $z = (z_1, \dots, z_l) \in (\{0, 1\}^m)^l$, $k \in \mathbb{N}$, and efficiently computable $h : \mathbb{Z}_N^k \mapsto \mathbb{Z}_N^l$ with $h(x) = \langle h_1(x), \dots, h_l(x) \rangle$. We say that (h, z) is closed form efficient for PRF if there exists an algorithm $\text{CFEval}_{h,z}$ such that for every $x \in \mathbb{Z}_N^k$,

$$\text{CFEval}_{h,z}(x, K) = \prod_{i=1}^l [\text{PRF}_K(z_i)]^{h_i(x)}$$

and the running time of CFEval is polynomial in n, m, k but sublinear in l .

The last property is very important for our purposes since it allows to run certain computations very fast when the secret key is known. We next describe two implementations for algebraic PRFs introduced in [BGV11].

3.1.1 Algebraic PRFs From Strong DDH

Let \mathcal{G} be a computational group scheme. The following construction \mathcal{PRF}_1 is an algebraic PRF with polynomial sized domains.

$\text{KeyGen}(1^n, m)$: Generate a group description $(\mathbb{G}, p, g) \leftarrow \mathcal{G}(1^n)$. Choose $k_0, k_1 \leftarrow \mathbb{Z}_p$. Output $\text{param} = (m, p, g, \mathbb{G}), K = (k_0, k_1)$.

$\text{PRF}_K(x)$: Interpret x as an integer in $\{0, \dots, D = 2^m\}$ where D is polynomial in n . Compute and output $g^{k_0 k_1^x}$.

Closed form efficiency for polynomials of degree d . We now show an efficient closed form for \mathcal{PRF}_1 for polynomials of the form (where evaluation is computed in the exponent)

$$Q(x) = \text{PRF}_K(0) \cdot \text{PRF}_K(1)^x \cdot \dots \cdot \text{PRF}_K(d)^{x^d} = \prod_{i=0}^d \text{PRF}_K(i)^{x^i}$$

where $d \leq D$. Let $h : \mathbb{Z}_p \mapsto \mathbb{Z}_p^{d+1}$, be defined as $h(x) \stackrel{\text{def}}{=} (1, x, \dots, x^d)$ and $(z_0, \dots, z_d) = (0, \dots, d)$. Then, we can define

$$\text{CFEval}_h(x, K) \stackrel{\text{def}}{=} g^{\frac{k_0(k_1^{d+1}x^{d+1}-1)}{k_1x-1}}.$$

Specifically, we write

$$\prod_{i=0}^d [\text{PRF}_K(z_i)]^{h_i(x)} = \prod_{i=0}^d [g^{k_0 k_1^i}]^{x^i} = g^{k_0 \sum_{i=0}^d k_1^i x^i}.$$

Correctness of CFEval follows by the identity $\sum_{i=0}^d k_0 k_1^i x^i = \frac{k_0((k_1 x)^{d+1}-1)}{k_1 x-1}$.

Theorem 3.2 ([BGV11]) *Suppose that the D -Strong DDH assumption holds. Then, \mathcal{PRF}_1 is a pseudorandom function.*

3.1.2 Algebraic PRFs From DDH

Let \mathcal{G} be a computational group scheme. Define \mathcal{PRF}_2 as follows.

KeyGen($1^n, m$): Generate a group description $(p, g, \mathbb{G}) \leftarrow \mathcal{G}(1^n)$. Choose $k_0, k_1, \dots, k_m \leftarrow \mathbb{Z}_p$. Output $\text{param} = (m, p, g, \mathbb{G}), K = (k_0, k_1, \dots, k_m)$.

PRF $_K$ (x): Interpret $x = (x_1, \dots, x_m)$ as an m -bit string. Compute and output $g^{k_0 \prod_{i=1}^m k_i^{x_i}}$.

This function is known by the Naor-Reingold function [NR97].

Closed form for polynomials of degree d . We describe an efficient closed form for \mathcal{PRF}_2 for computing polynomials of the same form as above. That is,

$$Q(x) = \text{PRF}_K(0) \cdot \text{PRF}_K(1)^x \cdot \dots \cdot \text{PRF}_K(d)^{x^d} = \prod_{i=0}^d \text{PRF}_K(i)^{x^i}.$$

Let $h : \mathbb{Z}_p \mapsto \mathbb{Z}_p^{d+1}$, defined as $h(x) = (1, x, \dots, x^d)$ and let $z = (z_1, \dots, z_l) = (0, \dots, d)$ then

$$\text{CFEval}_{h,z}(x, K) \stackrel{\text{def}}{=} g^{k_0(1+k_1x)(1+k_2x^2)\dots(1+k_mx^{2^m})}$$

with $m = \lceil \log d \rceil$ (clearly, d must be a power of 2).

Theorem 3.3 ([NR97]) *Suppose that the DDH assumption holds. Then, \mathcal{PRF}_2 is a pseudorandom function.*

To this end, we only consider $z = (0, \dots, d)$ and omit z from the subscript, writing $\text{CFEval}_h(x, K)$ instead.

3.2 Our OPE Constructions

We describe our protocol for oblivious polynomial evaluation in the $\mathcal{F}_{\text{MaskPoly}}$ -hybrid setting, where the parties have access to a trusted party that computes functionality $\mathcal{F}_{\text{MaskPoly}} : (K, t) \mapsto (-, g^{R(t)})$ relative to some prime order group \mathbb{G} and generator g that are picked by P_0 , for $g^{R(\cdot)} = (g^{r_0}, \dots, g^{r_d})$ and $g^{r_i} = \text{PRF}_K(i)$ for all i . For simplicity, we first describe a semi-honest variant of our protocol and then show how to enhance its security into the malicious setting. Formally, let $\mathcal{PRF} = \langle \text{KeyGen}, \text{PRF}, \text{CFEval} \rangle$ denote an algebraic PRF with a range group \mathbb{G} (cf. Definition 3.1), then our semi-honest protocol follows.

Protocol 2 (Protocol π_{OPE} with semi-honest security.)

- **Input:** Party P_0 is given a d -degree polynomial $g^{Q(\cdot)} = (g^{q_0}, \dots, g^{q_d})$ with coefficients q_i 's from \mathbb{Z}_p with respect to prime order group \mathbb{G} and generator g . Party P_1 is given an element t from \mathbb{Z}_p . Both parties are given a security parameter 1^n , group description \mathbb{G}, p and g .
- **The protocol:**
 1. **Masking the polynomial.** P_0 invokes $(K, \text{param}) \leftarrow \text{KeyGen}(1^n, \lceil \log d \rceil)$ where param includes a group description \mathbb{G} of prime order p and a generator g . It next defines a sequence of d elements $\tilde{R}(\cdot) = (\tilde{r}_0, \dots, \tilde{r}_d)$ over \mathbb{G} where $\tilde{r}_i = \text{PRF}_K(i)$ for all i .
 P_0 sends P_1 param and the masked polynomial $C(\cdot) = (g^{q_0 \tilde{r}_0}, \dots, g^{q_d \tilde{r}_d})$, where multiplication is implemented (componentwise) in \mathbb{G} .
 2. **Unmasking the result.** Upon receiving the masked polynomial $C(\cdot) = (c_0, \dots, c_d)$, party P_1 computes the polynomial evaluation $C(t) = \prod_{i=0}^d (c_i)^{t^i}$. I.e., $C(\cdot)$ is evaluated in the exponent. Next, the parties invoke an ideal execution of $\mathcal{F}_{\text{MaskPoly}}$ where the input of P_0 is K and the input of P_1 is t . Let Z denote the output of P_1 from this ideal call, then P_1 's output is $C(t)/Z$ where division is implemented in \mathbb{G} .

Note first that correctness holds since party P_1 computes in Step 2 the polynomial evaluation

$$C(t) = \prod_{i=0}^d (c_i)^{t^i} = \prod_{i=0}^d (g^{q_i \tilde{r}_i})^{t^i} = \prod_{i=0}^d (g^{q_i} g^{r'_i})^{t^i} = g^{Q(t)+R(t)}$$

and then “fixes” its computation by dividing out $Z = g^{R(t)}$. In addition, privacy holds due to the pseudo-randomness of \mathcal{PRF} that hides the coefficients of $Q(\cdot)$. We prove the following theorem,

Theorem 3.4 *Assume $\mathcal{PRF} = \langle \text{KeyGen}, \text{PRF}, \text{CFEval} \rangle$ is an algebraic PRF, then Protocol 2 securely realizes functionality \mathcal{F}_{OPE} in the presence of semi-honest adversaries in the $\mathcal{F}_{\text{MaskPoly}}$ -hybrid model.*

Proof: We prove security for each corruption case separately.

P_0 is corrupted. Let \mathcal{A} be a PPT adversary corrupting party P_0 , we design a PPT simulator SIM that invokes \mathcal{A} by playing the role of the honest P_1 and generating a view that is indistinguishable from a hybrid view. More concretely,

1. Given input $(1^n, g^{Q(\cdot)}, z)$, SIM invokes \mathcal{A} on this input and receives \mathcal{A} 's first message, (\mathbb{G}, p, g) and a d -degree polynomial $C(\cdot)$ (where d is also the degree of $Q(\cdot)$).
2. SIM emulates the ideal calls of $\mathcal{F}_{\text{MaskPoly}}$, playing the role of the trusted party that receives from \mathcal{A} a PRF key K .
3. SIM outputs whatever \mathcal{A} does.

We prove the following claim,

Claim 3.1 For any tuple of inputs $(t, g^{Q(\cdot)})$ and auxiliary input z ,

$$\{\mathbf{IDEAL}_{\mathcal{F}_{\text{OPE}}, \mathcal{SIM}(z)}(n, (t, g^{Q(\cdot)}))\}_{n \in \mathbb{N}} \equiv \{\mathbf{HYBRID}_{\pi, \mathcal{F}_{\text{MaskPoly}}, \mathcal{A}(z)}(n, (t, g^{Q(\cdot)}))\}_{n \in \mathbb{N}}.$$

Proof: It is trivial to verify that \mathcal{A} 's view in the hybrid execution is identically distributed to its view in the ideal execution, since it does not receive any message in the hybrid execution. Now, since correctness holds due to the above argument, then the joint output distribution of both parties is identical in both executions. \square

P_1 is corrupted. Let \mathcal{A} be a PPT adversary corrupting party P_1 , we design a PPT simulator \mathcal{SIM} that invokes \mathcal{A} by playing the role of the honest P_0 and generating a view that is indistinguishable from a hybrid view. More concretely,

1. Given input $(1^n, t, g^{Q(t)}, z)$, \mathcal{SIM} invokes \mathcal{A} on $(1^n, t, z)$ and continues as follows.
2. \mathcal{SIM} invokes first $(\mathbb{G}, p, g) \leftarrow \text{KeyGen}(1^n, \lceil \log d \rceil)$. It then picks a random d -degree polynomial $S(\cdot) = g^{S(\cdot)}$ in \mathbb{G} (we assume that d is part of \mathcal{A} 's auxiliary input, alternatively, d can be part of P_1 's output within \mathcal{F}_{OPE}), and sends it to \mathcal{A} .
3. \mathcal{SIM} then emulates the ideal call of $\mathcal{F}_{\text{MaskPoly}}$ by playing the role of the trusted party that receives t from \mathcal{A} . \mathcal{SIM} replies with $g^{S(t)}/g^{Q(t)}$.
4. \mathcal{SIM} outputs whatever \mathcal{A} does.

We prove the following claim,

Claim 3.2 For any tuple of inputs $(t, g^{Q(\cdot)})$ and auxiliary input z ,

$$\{\mathbf{IDEAL}_{\mathcal{F}_{\text{OPE}}, \mathcal{SIM}(z)}(n, (t, g^{Q(\cdot)}))\}_{n \in \mathbb{N}} \stackrel{c}{\approx} \{\mathbf{HYBRID}_{\pi, \mathcal{F}_{\text{MaskPoly}}, \mathcal{A}(z)}(n, (t, g^{Q(\cdot)}))\}_{n \in \mathbb{N}}.$$

Proof: We now prove that the adversary's view in both hybrid and ideal executions is computationally indistinguishable by a reduction to the security of \mathcal{PRF} . First, define a modified simulator \mathcal{SIM}' that generates \mathcal{A} 's view as follows. Assume that \mathcal{SIM}' knows the polynomial in the exponent $g^{Q(\cdot)}$, then \mathcal{SIM}' first picks a truly random d -degree polynomial $T(\cdot)$ in \mathbb{Z}_p and sends \mathcal{A} the coefficients of the masked polynomial $g^{Q(\cdot)+T(\cdot)}$. Later, it hands \mathcal{A} the value $g^{T(t)}$ within $\mathcal{F}_{\text{MaskPoly}}$. We claim that \mathcal{A} 's view in the modified simulation with \mathcal{SIM}' is identical to its view in the simulation above with \mathcal{SIM} . This is because the random polynomial $S(\cdot) = g^{S(\cdot)}$ is distributed identically to $g^{T(\cdot)+Q(\cdot)}$ for some polynomial $T(\cdot)$ in \mathbb{Z}_p . Therefore, the message $g^{S(t)}/g^{Q(t)}$ sent by simulator \mathcal{SIM} within $\mathcal{F}_{\text{MaskPoly}}$, is identically distributed to $g^{T(t)+Q(t)}/g^{Q(t)} = g^{T(t)}$.

Next, we prove that the adversary's view in the simulation and in the hybrid execution is computationally indistinguishable, relying on the pseudorandomness of \mathcal{PRF} . More specifically, assume by contradiction the existence of an adversary \mathcal{A} and a distinguisher D_{OPE} that distinguishes \mathcal{A} 's simulated and hybrid views for infinitely many n 's, construct a distinguisher D_F that distinguishes F from a truly random function f that maps $\{0, 1\}^m$ to \mathbb{G} , for infinitely many n 's. Namely, upon given access to an oracle \mathcal{O} that implements either F or f , and auxiliary input $(1^n, g^{Q(\cdot)}, t, z)$, D_F invokes \mathcal{A} on $1^n, t, z$. It then computes the first message of protocol Π_{OPE} as follows. It first invokes its oracle on the set $(0, \dots, d)$; let f_0, \dots, f_d be the oracle responses. D_F then sends \mathcal{A} the coefficients $\{g^{a_i} f_i\}_{i=0}^d$, for $g^{Q(\cdot)} = (g^{a_0}, \dots, g^{a_d})$. Next, D_F

emulates the trusted party for $\mathcal{F}_{\text{MaskPoly}}$, receiving a value t from \mathcal{A} and replying with $\prod_{i=0}^d f_i^{t^i}$. Finally, D_F invokes D_{OPE} on \mathcal{A} 's view and outputs whatever D_{OPE} does.

Note that if oracle \mathcal{O} implements function F , then \mathcal{A} 's view distributes as in the hybrid execution. On the other hand, if \mathcal{O} implements f then \mathcal{A} 's view distributes as in the simulation. Therefore, any gap in distinguishing these two views is immediately translated into a gap distinguishing F from f . This concludes the proof. \square \blacksquare

Efficiency. In the first phase P_0 computes $d + 1$ modular exponentiations as it can first compute the PRF evaluations in \mathbb{Z}_p (using the PRF key K) and then raise the outcomes to the power of g . Next, P_0 multiplies each PRF evaluation $\text{PRF}_K(i)$ with g^{q^i} (where these computations can be combined into a single exponentiation per index i). Efficiency of the second phase is dominated by the degree of $Q(\cdot)$ and the implementation of functionality $\mathcal{F}_{\text{MaskPoly}}$. In Section 3.3 we discuss several ways to realize $\mathcal{F}_{\text{MaskPoly}}$. (1) Assuming the strong-DDH assumption, our protocol requires a constant number of modular exponentiations. (2) Assuming the DDH assumption our protocol requires $O(\log d)$ modular exponentiations. Therefore, the overall cost is $2(d + 1) + O(1)$ (resp. $O(\log d)$) exponentiations.

Extension for a reactive OPE functionality. In a more realistic setting, we may consider a reactive variant of the OPE functionality where P_0 first commits to its polynomial *once*, similarly to the first message in π_{OPE} , and then receives multiple requests for polynomial evaluations. These requests can be adaptively sent by either a single or distinct parties. In this case, the amortized work of P_0 would be $O(1)$ (resp. $O(\log d)$) exponentiations for $O(d)$ evaluation requests. Note that such a reactive functionality captures a broader class of applications, e.g., when a single authority communicates with several different entities that wish to securely check membership in the authority's list. This further implies that correctness must be enforced by ensuring that P_0 uses the same PRF key for all of its evaluations; see further discussion in Section 4.

3.2.1 Security in the Presence of Malicious Adversaries

We next prove the security of Protocol 2 in the presence of malicious attacks. We observe that if the protocol that implements $\mathcal{F}_{\text{MaskPoly}}$ is secure in the presence of malicious corruptions then the entire protocol is secure against malicious attacks as well. Intuitively, security against a corrupted P_1 is immediately implied since a corrupted P_1 does not learn anything beyond $g^{R(t')}$, where t' is P_1 's input to $\mathcal{F}_{\text{MaskPoly}}$. More concretely, in the security proof the simulator publishes a random polynomial $\tilde{S}(\cdot) = g^{S(\cdot)}$ first, and then extracts P_1 's input t' to π_{MaskPoly} . Finally, the simulator forces P_1 's output within π_{MaskPoly} to be $g^{S(t')}/g^{Q(t')}$.

In case P_0 is corrupted we need to demonstrate how to extract the coefficients of $g^{Q(\cdot)}$. This is achieved by the fact that P_0 is committed to the PRF key K within $\mathcal{F}_{\text{MaskPoly}}$. Specifically, extraction is obtained by first calculating the sequence of values $\tilde{R}(\cdot)$ using K , and then dividing them out (component-wise) from the polynomial $C(\cdot)$ sent by P_0 within the first message. The technical part is to prove that P_1 indeed learns the same value in both ideal and real executions. Intuitively, by the correctness of the ideal execution for functionality $\mathcal{F}_{\text{MaskPoly}}$ we are ensured that P_1 learns $g^{R(t)}$ such that $\tilde{R}(x) = g^{R(x)} = \tilde{r}_0 \cdot \tilde{r}_1^x \cdot \dots \cdot \tilde{r}_d^{x^d}$ where $\tilde{r}_i = \text{PRF}_K(i)$ for all i and K is the PRF key as above. This implies that P_1 outputs $C(t)/g^{R(t)}$, which corresponds to a polynomial evaluation with the same coefficients as defined above by the simulator.

To conclude, in order to obtain malicious security the only modification we need to consider with respect to π_{OPE} is to employ a maliciously secure implementation of functionality $\mathcal{F}_{\text{MaskPoly}}$. In the hybrid setting this does not make a difference for the protocol description. In Section 3.3 we discuss secure implementations of functionality $\mathcal{F}_{\text{MaskPoly}}$. The proof for the following theorem formalizes the intuition above.

Theorem 3.5 Assume $\mathcal{PRF} = \langle \text{KeyGen}, \text{PRF}, \text{CFEval} \rangle$ is an algebraic PRF, then Protocol 2 securely realizes functionality \mathcal{F}_{OPE} in the presence of malicious adversaries in the $\mathcal{F}_{\text{MaskPoly}}$ -hybrid model.

Proof: We prove security for each corruption case separately.

P_0 is corrupted. Let \mathcal{A} be a PPT adversary corrupting party P_0 , we design a PPT simulator \mathcal{SIM} that extracts the coefficients of \mathcal{A} 's input polynomial in the exponent of some generator g , and sends them to the trusted party. We prove that the joint output distribution in both ideal and hybrid executions is computationally indistinguishable.

1. Given input $(1^n, g^{Q(\cdot)}, z)$, \mathcal{SIM} invokes \mathcal{A} on this input and receives \mathcal{A} 's first message, (\mathbb{G}, p, g) and a d -degree polynomial $C(\cdot) = (c_0, \dots, c_d)$.
2. \mathcal{SIM} emulates the ideal calls of $\mathcal{F}_{\text{MaskPoly}}$, playing the role of the trusted party that receives from \mathcal{A} a PRF key K .
3. \mathcal{SIM} defines polynomial $g^{Q'(\cdot)}$ with the following set of coefficients. Namely, \mathcal{SIM} fixes $g^{q_i} = c_i/\tilde{r}_i$ for all $0 \leq i \leq d$ where $\tilde{r}_i = \text{PRF}_K(i)$. \mathcal{SIM} sends the trusted party the set $\{g^{q_i}\}_i$.
4. \mathcal{SIM} outputs whatever \mathcal{A} does.

Note that \mathcal{A} receives no messages in the hybrid version of our protocol. Therefore security is implied by correctness of the closed form efficiency property that implies that the coefficients defined by the simulator implies a polynomial evaluation that is consistent with the real result.

P_1 is corrupted. Let \mathcal{A} be a PPT adversary corrupting party P_1 , we design a PPT simulator \mathcal{SIM} that generates the view of \mathcal{A} as in the semi-honest case with the following difference. \mathcal{SIM} invokes $\mathcal{A}(1^n, t, z)$ and sends it (\mathbb{G}, p, g) and a random polynomial $\tilde{S}(\cdot) = g^{S(\cdot)}$. Upon receiving the adversary's input t' to $\mathcal{F}_{\text{MaskPoly}}$, \mathcal{SIM} forwards it to the trusted party for \mathcal{F}_{OPE} . Let Z denotes the output returned by the trusted party, then \mathcal{SIM} completes the simulation, forcing the output of $\mathcal{F}_{\text{MaskPoly}}$ to be $g^{S(t')}/Z$. We note that security follows exactly as in the semi-honest proof since \mathcal{A} sees the same message distribution. ■

3.3 Secure Protocols for π_{MaskPoly}

In this section we describe a concrete protocol that implements functionality $\mathcal{F}_{\text{MaskPoly}} : (K, t) \mapsto (-, g^{R(t)})$, used as a subprotocol within our main protocol π_{OPE} for oblivious polynomial evaluation from Section 3.2. This computation corresponds to the polynomial evaluation $\tilde{R}(x) = \text{PRF}_K(0) \cdot \text{PRF}_K(1)^x \cdot \dots \cdot \text{PRF}_K(d)^{x^d}$ with respect to function PRF. In what follows, we discuss a detailed secure implementation for \mathcal{PRF}_1 that is described in Section 3.1.1 and then briefly discuss how to implement function \mathcal{PRF}_2 , formally described in Section 3.1.2, using similar ideas.

We recall that when implementing functionality $\mathcal{F}_{\text{MaskPoly}}$ relative to \mathcal{PRF}_1 the parties compute the value $g^{\sum_{i=0}^d k_0 k_1^i x^i} = g^{\frac{k_0((k_1 x)^{d+1}-1)}{k_1 x - 1}}$, so that P_0 enters a PRF key $K = (k_0, k_1)$ and learns nothing and P_1 enters $x = t$ and learns this outcome. This is a simple computation that requires a constant number of exponentiations and can be easily implemented securely. Achieving malicious security requires to ensure correctness of computations which we obtain using simple zero-knowledge proofs of knowledge. Loosely speaking, the parties first generate a joint public key for the additive El Gamal PKE such that no party knows the secret key (this protocol is discussed in Section 2.4.1). Next, each party commits to its input and the

parties jointly compute $k_1 t$. A slight complication arises since the parties need to compute the inverse of $k_1 t - 1$. Relying on the fact that $(k_1 t - 1)^{-1} = (k_1 t - 1)^{p-2} \pmod p$ and that

$$\frac{k_0((k_1 t)^{d+1} - 1)}{k_1 t - 1} = \frac{k_0(k_1 t)^{d+1} - k_0}{k_1 t - 1} = \frac{k_0(k_1 t)^{d+1}}{k_1 t - 1} - \frac{k_0}{k_1 t - 1} = \frac{k_0 k_1^{d+1} t^{d+1}}{k_1 t - 1} - \frac{k_0}{k_1 t - 1},$$

we let the parties compute the inverse of $k_1 t - 1$ first and then complete the computation by multiplying the result with $k_0(k_1 t)^{d+1}$ and k_0 . Formally, our protocol uses the following tools:

1. Distributed additive El Gamal; see Section 2.4 for the details of these schemes. We denote this scheme by $\Pi = (\pi_{\text{KeyGen}}, \text{Enc}, \pi_{\text{DEC}})$.
2. Zero-knowledge proofs of knowledge: π_{DL} for proving a discrete logarithm and π_{Eq} for proving consistency of exponents, which are formally stated in Section 2.5. For simplicity, we implicitly assume that the parties continue to the next step only if all the proofs were verified correctly thus far.

Finally, we implicitly assume that a party rerandomizes its homomorphic computations on the ciphertexts. Such that rerandomization is carried out by multiplying the outcome with a random encryption of zero. We now describe our protocol in details.

Protocol 3 (Protocol π_{MaskPoly} with malicious security.)

- **Input:** Party P_0 is given a PRF key $K = (k_0, k_1)$. Party P_1 is given an element t . Both parties are given a security parameter 1^n , a polynomial degree d and (\mathbb{G}, p, g) for a group description \mathbb{G} of prime order p and a generator g .
- **Convention:** Homomorphic operations on ciphertexts are computed componentwise.
- **The protocol:**
 1. **Distributed key generation.** P_0 and P_1 run protocols $\pi_{\text{KeyGen}}(1^n, 1^n)$ in order to generate additive El Gamal public key $\text{PK} = \langle \mathbb{G}, p, g, h \rangle$ for which the corresponding shares of the secret key SK are $(\text{SK}_0, \text{SK}_1)$. P_0 then sends P_1 encryptions of k_0 and k_1 , denoted by c_{k_0} and c_{k_1} , and proves their knowledge using π_{DL} .
 2. **Computing encryption of $k_1 t$.** Upon receiving ciphertexts c_{k_0} and c_{k_1} , P_1 sends P_0 an encryption of its input t , denoted by c_t . It further computes the encryption of $k_1 t$, denoted by $c_{k_1 t}$, and proves consistency relative to c_t and $c_{k_1 t}$ using the zero-knowledge proof π_{Eq} .
 3. **Computing encryptions of k_1^{d+1} and t^{d+1} .** P_0 computes the encryption of k_1^{d+1} , denoted by $c_{k_1^{d+1}}$, and proves consistency between g^{d+1} and $c_{k_1^{d+1}}$ using π_{Eq} . Similarly, P_0 computes the encryption of t^{d+1} , denoted by $c_{t^{d+1}}$, and proves correctness.
 4. **Computing encryption of $(k_1 t - 1)^{-1}$.** The parties compute the inverse of $(k_1 t - 1)$, by first computing the encryption of $k_1 t - 1$ given ciphertext $c_{k_1 t}$ from above, and then raising the result to the power of $p - 2$. Let c_{inv} denote the outcome.
 5. **Computing encryptions of $k_0(k_1 t - 1)^{-1}$ and $k_0 k_1^{d+1}(k_1 t - 1)^{-1}$.** Given ciphertexts c_{inv} , $c_{k_1^{d+1}}$ and c_{k_0} , P_0 computes the encryptions of $k_0(k_1 t - 1)^{-1}$ and $k_0 k_1^{d+1}(k_1 t - 1)^{-1}$ and proves consistency relative to c_{inv} , $c_{k_1^{d+1}}$ and c_{k_0} using π_{Eq} (where the proof of the later computation involves running π_{Eq} twice). Let c_0 and c'_0 denote the respective outcomes.
 6. **Computing encryption of $k_0 k_1^{d+1} t^{d+1}(k_1 t - 1)^{-1}$.** Given ciphertexts $c_{t^{d+1}}$ and c'_0 , P_1 computes the encryption of $k_0 k_1^{d+1} t^{d+1}(k_1 t - 1)^{-1}$ and proves consistency using π_{Eq} . Let c_1 denote the respective outcome.
 7. **Outcome.** Finally, the parties decrypt c_1/c_0 for P_1 that outputs the result.

Theorem 3.6 Assume $\Pi = (\pi_{\text{KeyGen}}, \text{Enc}, \pi_{\text{DEC}})$, π_{DL} and π_{Eq} are as above, then Protocol 3 securely realizes functionality $\mathcal{F}_{\text{MaskPoly}}$ with respect to \mathcal{PRF}_1 in the presence of malicious adversaries.

Proof Sketch: Informally, simulation follows since the parties are forced to prove their computations using ZK proofs which guarantees correctness. In addition, privacy is implied by the fact that all the computations are carried out on ciphertexts. More concretely, whenever P_0 is corrupted, the simulator extracts k_0 and k_1 from the ZK proof π_{DL} , and sends these values to the trusted party. It further plays the role of the honest P_1 using an arbitrary input and aborts if one of the proofs is not verified correctly. Security is easily reduced to the CPA-IND security of the distributed El Gamal PKE since P_0 does not learn any output. On the other hand, whenever P_1 is corrupted simulation follows similarly by having the simulator use an arbitrary PRF key and extract t from the proof that verifies the computation of $k_1 t$, except that the simulator also cheats in the mutual decryption in the final step and forces the result to be the output it received from the trusted party. Security in this case follows due to the CPA-IND security of the distributed El Gamal PKE and the security of protocol π_{DEC} . Finally, we note that the complexity of our protocol is constant. ■

The implementation of PRF \mathcal{PRF}_2 follows similarly. Namely, recall that the parties compute the value $g^{k_0(1+k_1x)(1+k_2x^2)\dots(1+k_mx^{2^m})}$ which can be carried out in $O(m)$ time as follows. First, P_0 commits to its key (k_0, k_1, \dots, k_m) , whereas P_1 commits to the elements (x, x^2, \dots, x^{2^m}) together with a ZK proof of consistency. Next, given the product $\tilde{g} = g^{k_0(1+k_1x)(1+k_2x^2)\dots(1+k'_mx^{2^{m'}})}$ for some integer $m' < m$, the parties mutually compute

$$\tilde{g} \cdot \tilde{g}^{k_{m'+1}x^{2^{(m'+1)}}} = \tilde{g}^{(1+k_{m'+1}x^{2^{(m'+1)}})} = g^{k_0(1+k_1x)(1+k_2x^2)\dots(1+k_{m'+1}x^{2^{(m'+1)}})}$$

where $\hat{g} = \tilde{g}^{k_{m'+1}}$ is carried out by P_0 and proven correct with respect the commitment of $g^{k_{m'+1}}$. This computation is followed by P_1 computing $\hat{g}^{x^{2^{(m'+1)}}}$ which is also verified against the commitment of $g^{x^{2^{(m'+1)}}}$ where the commitment is realized using El Gamal. See the ZK proof π_{Eq} for more details.

4 Secure Set-Intersection

One important application that benefits from our OPE construction is the set-intersection functionality which is defined by having each party's input consists of a *set* of elements from domain $\{0, 1\}^t$. Formally:

Definition 4.1 Let X and Y be subsets of a predetermined arbitrary domain $\{0, 1\}^t$ and m_X and m_Y the respective upper bounds on the sizes of X and Y .² Then functionality \mathcal{F}_\cap is defined by:

$$(X, Y) \mapsto (m_Y, (X \cap Y, m_X)).$$

To achieve a secure set-intersection protocol, we modify protocol π_{OPE} from Section 3.2.1 as follows. First, P_0 prepares a polynomial $Q(\cdot)$ with coefficients in \mathbb{Z}_p and the set of roots X . It then masks $Q(\cdot)$ as in Protocol 2 using the sequence of pseudorandom elements $\tilde{R}(\cdot)$. The parties then interact with a trusted party that computes functionality $\mathcal{F}_{\text{EqMask}}$, which is a slight variation of functionality $\mathcal{F}_{\text{MaskPoly}}$. Namely, instead of implementing $\mathcal{F}_{\text{MaskPoly}}$ the functionality checks for equality with respect to P_1 's polynomial evaluations of $g^{Q(\cdot)}\tilde{R}(\cdot)$ and $\tilde{R}(\cdot)$ on the set Y . This modification in the functionality's description is required due to the fact that we cannot let P_1 learn $Q(y)$ for arbitrary $y \in Y$ (even if P_1 is honest), since that would leak

²In order to deal with a proof technicality, where a corrupted party inputs less elements than its set size, prior constructions assume a super polynomial lower bound on the input domain sizes. Since we do not wish to restrict the input domains, we assume that the set sizes are not strict and may denote some upper bound on the actual numbers of elements.

information about X . More specifically, $\mathcal{F}_{\text{EqMask}}$ is defined by $(K, \{(y_i, T_i)\}_{y_i \in Y}) \mapsto (-, \{b_i\}_i)$, where $b_i = 1$ only if $T_i = g^{R(y_i)}$ and 0 otherwise, $g^{R(\cdot)} = (g^{r_0}, \dots, g^{r_{m_X}})$ and $g^{r_i} = \text{PRF}_K(i)$ for all i . Stated differently, $b_i = 1$ if and only if $Q(y_i) = 0$ (or $y_i \in X \cap Y$) with overwhelming probability. Finally, P_1 outputs the set of elements $Z \subseteq Y$ for which $b_i = 1$.

Our implementation for $\mathcal{F}_{\text{MaskPoly}}$ from Section 3.3 easily supports this functionality, since P_0 can run its zero-knowledge proofs with respect to a single set of ciphertexts encrypting its PRF key. In addition, in order to enable the extraction of the set X by the simulator we add zero-knowledge proofs of knowledge for the relation \mathcal{R}_{DL} , formally defined in Section 2.5. This technicality arises because P_0 sends elements in \mathbb{G} yet the polynomial $Q(\cdot) + R(\cdot)$ is evaluated *in the exponent*, implying that X and Y must be sampled from \mathbb{Z}_p as well. Note that P_0 may fix X and its masked polynomial in \mathbb{G} . Nevertheless, P_1 needs to know the discrete logarithms of Y with respect to some group generator g in order to evaluate the masked polynomial.

Formally, let $d = m_X - 1$, then define our set-intersection protocol as follows,

Protocol 4 (Protocol π_{\cap} with malicious security.)

- **Input:** Party P_0 is given a set X of size m_X . Party P_1 is given a set Y of size m_Y . Both parties are given a security parameter 1^n .
- **The protocol:**
 1. **Masking the input polynomial.** P_0 defines an d -degree polynomial $Q(\cdot) = (q_0, \dots, q_d)$ with coefficients in \mathbb{Z}_p and the set of roots X , for $d = m_X - 1$. It then invokes $(K, \text{param}) \leftarrow \text{KeyGen}(1^n, d)$ where param includes a group description \mathbb{G} of prime order p and a generator g , and defines a new d -degree polynomial $\tilde{R}(\cdot) = (\tilde{r}_0, \dots, \tilde{r}_d)$ over \mathbb{G} , where r_i is defined by $\text{PRF}_K(i)$ for all i .
 P_0 sends P_1 param and the masked polynomial $C(\cdot) = (g^{q_0} \tilde{r}_0, \dots, g^{q_d} \tilde{r}_d)$, where multiplication is implemented in \mathbb{G} . P_0 further proves the knowledge of the discrete logarithm of $c_i = g^{q_i} \tilde{r}_i$ for all i with respect to a generator g , by invoking an ideal execution of \mathcal{F}_{DL} on input $\{(g, c_i), \log_g c_i\}_{i \in [0, d]}$.³ The input of P_1 for \mathcal{F}_{DL} is $\{(g, c_i)\}_{i \in [0, d]}$.
 2. **Unmasking the result.** Upon receiving the masked polynomial $C(\cdot) = (c_0, \dots, c_d)$ and upon receiving from \mathcal{F}_{DL} the value 1, denoting “accept” for all i , party P_1 computes the polynomial evaluation $C(y) = \prod_{i=0}^d (c_i)^{y^i}$ for all $y \in Y$ (picked in a random order). I.e., $C(\cdot)$ is evaluated in the exponent.
Next, the parties invoke an ideal execution of $\mathcal{F}_{\text{EqMask}}$, where the input of P_0 is K and the input of P_1 is the set $\{(y, C(y))\}_{y \in Y}$. P_1 outputs y if and only if the output from $\mathcal{F}_{\text{EqMask}}$ on $(y, C(y))$ is 1.

Correctness follows easily since P_1 outputs only elements in Y that zeros polynomial $Q(\cdot)$, whom its roots are the set X . Next, we prove the following theorem.

Theorem 4.2 *Assume $\mathcal{PRF} = \langle \text{KeyGen}, F, \text{CFEval} \rangle$ is an algebraic PRF, then Protocol 4 securely realizes functionality \mathcal{F}_{\cap} in the presence of malicious adversaries in the $\{\mathcal{F}_{\text{DL}}, \mathcal{F}_{\text{EqMask}}\}$ -hybrid model.*

The proof follows similarly to the proof of Protocol 2. Formally,

Proof: We prove security for each corruption case separately.

P_0 is corrupted. Let \mathcal{A} be a PPT adversary corrupting party P_0 , we design a PPT simulator \mathcal{SIM} that simulates the view \mathcal{A} , playing the role of the honest P_1 while extracting \mathcal{A} 's input set X , details follow.

1. Given input $(1^n, X, z)$, \mathcal{SIM} invokes \mathcal{A} on this input and receives \mathcal{A} 's first message, (\mathbb{G}, p, g) and a d -degree polynomial $C(\cdot) = (c_0, \dots, c_d)$.

³We implicitly assume that P_0 knows the discrete logarithms of the r_i 's by its knowledge of K . This is the case for all PRF implementations presented in [BGV11].

2. SLM emulates the ideal calls of \mathcal{F}_{DL} by playing the role of the trusted party that receives from \mathcal{A} tuples $\{(g, c_i), c'_i\}_{i \in [0, d]}$ and records these values. SLM verifies whether $c_i = g^{c'_i}$ for all i and records 1 only if these conditions are met, and 0 otherwise. In case SLM records 0 it aborts and outputs whatever \mathcal{A} does.
3. SLM defines the input set X' as follows. For every i let $\tilde{r}_i = \text{PRF}_K(i)$ and $r_i = \log_g \tilde{r}_i$ and let $q'_i = c'_i - r_i$.⁴ SLM fixes polynomial $Q'(\cdot) = (q'_0, \dots, q'_d)$ and defines X' to be the set of roots of $Q'(\cdot)$. SLM computes X' by factoring $Q'(\cdot)$ over \mathbb{Z}_p and sends the set X' to the trusted party, receiving back m_Y .
4. SLM emulates the ideal call of $\mathcal{F}_{\text{MaskPoly}}$ by playing the role of the trusted party that receives from \mathcal{A} a PRF key K .
5. SLM outputs whatever \mathcal{A} does.

Note that the adversary's view is identical to its view in the hybrid execution since it does not get any output from the internal ideal calls as well as from \mathcal{F}_\cap . We now claim that P_1 's output is identical with overwhelming probability in both executions due to the following. In the hybrid execution the correctness of the ideal call for $\mathcal{F}_{\text{EqMask}}$ ensures that P_1 obtains the correct equality bit for every $y \in Y$. Namely, if $C(y) \neq \tilde{R}(y)$ then the honest P_1 obtains 0 from $\mathcal{F}_{\text{EqMask}}$ and does not output y . On the other hand, if $C(y) = \tilde{R}(y)$ then P_0 receives 1 and returns y . Stating differently, P_1 returns $y \in Y$ only if $C(y)/\tilde{R}(y) = 1$ where division is computed component-wise. Next, in the simulation SLM defines the input set X' of the adversary as the set of roots with respect to the unmasked polynomial $C(\cdot)/\tilde{R}(\cdot)$ (computed component-wise), where the masking is defined by the PRF key K input by the adversary to $\mathcal{F}_{\text{EqMask}}$. Therefore the intersection is computed with respect to the same set X' .

P_1 is corrupted. Let \mathcal{A} be a PPT adversary corrupting party P_1 , we design a PPT simulator SLM that generates the view of \mathcal{A} as follows. SLM first sends a random polynomial $\tilde{S}(\cdot)$. Next, upon receiving the adversary's set of elements Y' to $\mathcal{F}_{\text{MaskPoly}}$, SLM forwards it to the trusted party for \mathcal{F}_\cap . Let Z' denotes the output returned by the trusted party, then SLM completes the simulation by forcing the output of \mathcal{A} within $\mathcal{F}_{\text{EqMask}}$ to be consistent with the set Z . More formally,

1. Given input $(1^n, Y, z)$, SLM invokes \mathcal{A} on this input and sends it (\mathbb{G}, p, g) .
2. SLM picks a random d -degree polynomial $\tilde{S}(\cdot) = (\tilde{s}_0, \dots, \tilde{s}_d) = (g^{s_0}, \dots, g^{s_d})$ with coefficients in \mathbb{G} and sends it to \mathcal{A} . (We assume that the simulator knows m_X as part of its auxiliary information. This can also be assured by modifying the definition of the functionality, given m_X to P_1 as part of its input).
3. SLM emulates the ideal calls of \mathcal{F}_{DL} by playing the role of the trusted party that receives from \mathcal{A} tuples $\{(g, \tilde{s}_i)\}_{i \in [0, d]}$ and sends \mathcal{A} the value 1 for all i (denoting accept calls).
4. SLM then emulates the ideal call of $\mathcal{F}_{\text{EqMask}}$ by playing the role of the trusted party that receives from \mathcal{A} the set $\{(y'_j, T_{y'_j})\}_{j \in [m_Y]}$. SLM sends the set $Y' = \{y'_j\}_{j \in [m_Y]}$ to the trusted party, receiving back the intersection $Z = X \cap Y'$.

For all $y'_j \in Z$, SLM emulates the ideal response of $\mathcal{F}_{\text{EqMask}}$ as follows. If $T_{y'_j} = g^{S(y'_j)}$ then SLM sends \mathcal{A} the value 1. Otherwise it sends 0. For all $y'_j \notin Z$, SLM always replies with 0.

⁴See Footnote 3.

5. \mathcal{SIM} outputs whatever \mathcal{A} does.

Note that the protocol never verifies that \mathcal{A} 's inputs to $\mathcal{F}_{\text{EqMask}}$ are consistent pairs $\{(y'_j, T_{y'_j})\}_j$ of which $T_{y'_j} = g^{S(y'_j)}$ for all $j \in [m_Y]$. We prove that this is not required. Specifically, the differences between the hybrid and simulated executions are as follows. First, \mathcal{SIM} sends in the simulation a random polynomial instead of a real masked polynomial. In addition, \mathcal{SIM} fixes the output of $\mathcal{F}_{\text{EqMask}}$ based on the correctness of \mathcal{A} 's computations which deviates from the way this functionality is defined. Consider a hybrid game Hyb where the simulator $\mathcal{SIM}_{\text{Hyb}}$ uses the real input X of P_0 to define polynomial $Q(\cdot)$, but decides on the output of $\mathcal{F}_{\text{EqMask}}$ according to the strategy specified in the simulation. Namely for every pair $(y'_j, T_{y'_j})$, $\mathcal{SIM}_{\text{Hyb}}$ verifies first whether $T_{y'_j} = C(y'_j)$ and returns 1 if equality holds. Clearly, the views induced in Hyb and in the simulation are computationally indistinguishable due to the pseudorandomness of F . This argument is similar to the argument presented in the proof of Protocol 2. Next, we claim that the distributions induced by the views of the hybrid execution and game Hyb are statistically close.

Formally, for every y'_j consider two cases. (i) $y'_j \notin X$ which implies that y'_j is not in the intersection and that $b_j = 0$ in the simulation of Hyb. Next, define a Bad event in which \mathcal{A} receives $b_j = 1$ from the trusted party for $\mathcal{F}_{\text{EqMask}}$ in the hybrid execution. Clearly, this event holds only if $T_{y'_j} = \text{CFEval}(y'_j, K) = g^{R(y'_j)}$ for K the PRF key entered by the honest P_0 , which implies that \mathcal{A} must correctly guess $\text{CFEval}(y'_j, K)$. We claim that the probability this event occurs is negligible due to the pseudorandomness of F and CFEval (in Section 5 we discuss the pseudorandomness of CFEval). Specifically, any successful guess with a non-negligible probability implies an attack on the PRF. Thus, the probability that Bad occurs is negligible. It therefore holds that the adversary's views are statistically close condition on the event that y'_j is not in the intersection. (ii) $y'_j \in X$ which implies that y'_j is in the intersection. Nevertheless, here there is no analogue bad event. This is because $b_j = 1$ only when $T_{y'_j} = C(y'_j) = \text{CFEval}(y'_j, K)$, which implies that $b_j = 1$ in both executions due to correctness of $\mathcal{F}_{\text{EqMask}}$.

This concludes the proof. ■

Efficiency. As in Protocol 2, the efficiency of Protocol 4 is dominated by the implementation of functionality $\mathcal{F}_{\text{EqMask}}$. Our protocols from Section 3.3 can be easily modified to support this functionality without significantly effecting their overhead, since the parties can first compute the encryption of the closed form efficiency of the PRF and then compare it with the input of P_1 . Therefore, the overall communication complexity is $O(m_X)$ group elements for sending the first message and $O(m_Y)$ (resp. $O(m_Y \log m_Y)$) group elements for the second phase of implementing $\mathcal{F}_{\text{EqMask}}$ for each $y \in Y$, depending on the underlying PRF. In particular, the number of modular exponentiations implies multiplicative costs in the sets sizes since P_1 evaluates its masked polynomial for each element in Y . Next, we demonstrate how to reduce this cost using hash functions.

4.1 Improved Constructions Using Hash Functions

We now show how to reduce the computational overhead using hash functions by splitting the set elements into smaller bins. Our protocol is applicable for different hash functions such as: simple hashing, balanced allocations [ABKU99] and Cuckoo hashing [KMW08]. For simplicity, we first describe our protocol for the simple hashing case; see Section 4.1.3 for a discussion about extensions to the other two hashing. Informally, the parties first agree on a hash function that is picked from a family of hash functions and induces a set of bins with some upper bound on the number of elements in each bin. Next, P_0 maps its elements into these bins and generates a polynomial for each such bin, which is computed as in Protocol 4 but with a smaller degree. Finally, P_0 masks all the polynomials and sends them to P_1 . Upon receiving the masked

polynomials, P_1 maps its elements into the same set of bins and evaluates the masked polynomials for these mapped bins. In the last step, the parties unmask these evaluations. To be precise, we need to specify how the masking procedure works and ensure that the parties do not deviate from the instructions of the protocol.

We fix some notations first. We denote by h the hash function picked by the parties, by \mathcal{B} the number of bins and by \mathcal{M} the maximum number of elements allocated to any single bin (where \mathcal{B} and \mathcal{M} are parameters specified by the concrete hash function in use and further depend on m_X). Note that the potential number of allocated elements is bounded by $\mathcal{B}\mathcal{M}$ which may be higher than the exact number m_X . This implies that the protocol must ensure that P_0 does not take advantage of that and introduce more set elements into the protocol execution. In addition, it must be ensured that a corrupted P_0 does not mask the zero polynomial, which would imply that P_1 accepts any value it substitutes in the masked polynomial. On the other hand, the protocol must ensure that a corrupted P_1 does not gain any information by entering incorrect values. Verifying that a polynomial is not all zeros can be easily done by substitution a random element in it and checking that the result is different than zero. In Section 4.1.1 we demonstrate how to enforce P_0 's correct behaviour by designing a new proof that exploits the algebraic properties of the underlying PRF. The verification procedure for P_1 is even simpler as demonstrated below.

Next, we explain how the masking procedure is computed. Denote by $Q_j(\cdot)$ the polynomial associated with the j th bin. If the degree of $Q_j(\cdot)$ is smaller than $\mathcal{M}-1$ then P_0 fixes the values of the $\mathcal{M}-\deg(Q_j(\cdot))$ leading coefficients to be zeros. It then masks the i th coefficient of $Q_j(\cdot)$ by multiplying it with $\text{PRF}_K((j-1) \cdot \mathcal{M} + i)$ for $i \in [0, \mathcal{M}-1]$. Furthermore, unmasking is computed by comparing the evaluation of the j th polynomial to the following computation

$$\prod_{i=0}^{j\mathcal{M}-1} \text{PRF}_K(i)^{x^i} / \prod_{i=0}^{(j-1)\mathcal{M}-1} \text{PRF}_K(i)^{x^i} = \text{PRF}_K((j-1)\mathcal{M})^{x^{(j-1)\mathcal{M}}} \cdot \dots \cdot \text{PRF}_K(j\mathcal{M}-1)^{x^{j\mathcal{M}-1}},$$

which is exactly the set of PRF values that mask polynomial $Q_j(\cdot)$.

More formally, our protocol uses two functionalities in order to ensure correctness. First, the parties call functionality $\mathcal{F}_{\text{Bins}}$ for proving that the masked polynomials sent by P_0 are correctly defined. Namely, $\mathcal{F}_{\text{Bins}} : (K, \{C_j(\cdot) = (c_0^j, \dots, c_{\mathcal{M}-1}^j)\}_{j \in [\mathcal{B}]}) \mapsto (-, b)$ and $b = 1$ only if none of the unmasked polynomials $\{Q_j(\cdot)\}_j$ is the zero polynomial and the overall degrees of these polynomials $\{Q_j(\cdot)\}_j$ is bounded by m_X . In addition, the parties call functionality $\mathcal{F}_{\text{EqMaskHash}}$ in order to correctly unmask polynomial evaluations $\{C_{h(y)}(y)\}_{y \in Y}$ for P_1 . We continue with the detailed description of our set-intersection protocol in the hybrid model. In Sections 4.1.1 and 4.1.2 we discuss how to securely implement these functionalities.

Protocol 5 (Protocol π_{\cap} with malicious security and hash functions.)

- **Input:** Party P_0 is given a set X of size m_X . Party P_1 is given a set Y of size m_Y . Both parties are given a security parameter 1^n .
- **The protocol:**
 1. **Fixing the parameters of the hash function.** The parties fix the parameters \mathcal{B} and \mathcal{M} of the hash function and picks a hash function $h : \{0, 1\}^t \mapsto [\mathcal{B}]$. P_0 invokes $(K, \text{param}) \leftarrow \text{KeyGen}(1^n, \mathcal{M}-1)$ where param includes a group description \mathbb{G} of prime order p and a generator g .
 2. **Masking the input polynomial.** For every $x \in X$, P_0 maps x into bin $h(x)$. Let \mathcal{B}_j denote the set of elements mapped into bin j . Next, P_0 constructs a polynomial $Q_j(\cdot) = (q_0^j, \dots, q_d^j)$ with coefficients in \mathbb{Z}_p and the set of roots \mathcal{B}_j . If $|\mathcal{B}_j| < \mathcal{M}$, P_0 fixes the leading $\mathcal{M} - |\mathcal{B}_j| - 1$ coefficients to zero. For each bin $j \in [\mathcal{B}]$, P_0 defines a new $(\mathcal{M}-1)$ -degree polynomial $\tilde{R}_j(\cdot) = (\tilde{r}_0^j, \dots, \tilde{r}_{\mathcal{M}-1}^j)$ over \mathbb{G} , where \tilde{r}_i^j is defined by $\text{PRF}_K((j-1)\mathcal{M} + i)$ for all $i \in [0, \mathcal{M}-1]$. P_0 sends P_1 param and the masked polynomials $\{C_j(\cdot)\}_j = \{g^{q_0^j} \tilde{r}_0^j, \dots, \dots, g^{q_{\mathcal{M}-1}^j} \tilde{r}_{\mathcal{M}-1}^j\}_j$, where multiplication is implemented in \mathbb{G} .

P_0 further proves the knowledge of the discrete logarithm of $c_i^j = g^{a_i} \tilde{r}_i^j$ for all i and j with respect to a generator g , by invoking an ideal execution of \mathcal{F}_{DL} on input $\{((g, c_i^j), \log_g c_i^j)\}_{i \in [0, \mathcal{M}-1], j \in [\mathcal{B}]}$.⁵ The input of P_1 for \mathcal{F}_{DL} is $\{(g, c_i^j)\}_{i \in [0, \mathcal{M}-1], j \in [\mathcal{B}]}$.

Finally, P_0 proves correctness using $\mathcal{F}_{\text{Bins}}$ where P_0 enters K and P_1 enters the masked polynomials.

3. **Unmasking the result.** Upon receiving the masked polynomials $\{C_j(\cdot) = (c_0^j, \dots, c_{\mathcal{M}-1}^j)\}_{j \in [\mathcal{B}]}$ and upon receiving accepting messages from $\mathcal{F}_{\text{DL}}, \mathcal{F}_{\text{Bins}}$, party P_1 computes the following for every $y \in Y$ (picked in a random order). It first maps y into bin $h(y)$ and then computes the polynomial evaluation $C_{h(y)}(y) = \prod_{i=(h(y)-1)\mathcal{M}}^{h(y)\mathcal{M}-1} (c_i^{h(y)})^{y^i}$. I.e., $C_{h(y)}(\cdot)$ is evaluated in the exponent.

Next, the parties invoke an ideal execution of $\mathcal{F}_{\text{EqMaskHash}}$, where the input of P_0 is K and the input of P_1 is the set $\{(y, h(y), C_{h(y)}(y))\}_{y \in Y}$.

P_1 outputs y only if the output from $\mathcal{F}_{\text{EqMaskHash}}$ on $(y, h(y), C_{h(y)}(y))$ is 1.

Theorem 4.3 *Protocol 5 securely realizes functionality \mathcal{F}_{\cap} in the presence of malicious adversaries in the $\{\mathcal{F}_{\text{DL}}, \mathcal{F}_{\text{Bins}}, \mathcal{F}_{\text{EqMaskHash}}\}$ -hybrid model.*

Security follows easily from the secure implementations of $\mathcal{F}_{\text{Bins}}$ and $\mathcal{F}_{\text{EqMaskHash}}$ and the proof of Protocol 4. We discuss these protocols next. We stress that P_1 needs to ensure in Protocol 5 that P_0 indeed uses the same PRF key for both sub-protocols (for instance by ensuring that P_0 enters the same commitment of K).

4.1.1 A Secure Protocol for $\mathcal{F}_{\text{Bins}}$

In this section we design a protocol π_{Bins} for securely implementing functionality $\mathcal{F}_{\text{Bins}} : (K, \{C_j(\cdot)\}_{j \in [\mathcal{B}]}) \mapsto (-, b)$ for which $b = 1$ only if none of the unmasked polynomials $\{Q_j(\cdot)\}_j$ is the zero polynomial and the overall degrees of all polynomials $\{Q_j(\cdot)\}_j$ is bounded by m_X . To prove that none of the polynomials is the all zeros polynomial we evaluate each masked polynomial on a random element and then verify that the result is different than zero. In particular, for each j the parties first agree on a random element z_j and then compute the polynomial evaluation $C_j(z_j)$. Next, the parties verify whether $C_j(z_j) = \tilde{R}_j(z_j)$ where $\tilde{R}_j(\cdot)$ is the masking polynomial of $C_j(\cdot)$. Note that if $Q_j(\cdot)$ is not the all zeros polynomial then $C_j(z_j) \neq \tilde{R}_j(z_j)$ with overwhelming probability over the choice of z_j . This is because there exists a coefficient $q_{i,j} \neq 0$ which implies that for $C_j(z_j) = Q_j(z_j) \cdot \tilde{R}_j(z_j)$. Now since $Q_j(z_j) \neq 0$ it holds that $C_j(z_j) \neq \tilde{R}_j(z_j)$. On the other hand, in case $Q_j(\cdot)$ is the zero polynomial then it holds that $C_j(z_j) = \tilde{R}_j(z_j)$ for all z_j . This is because $Q_j(z_j) = 0$ as all its coefficients equal zero.

The more challenging part is to prove that the overall degrees of all polynomials $\{Q_j(\cdot)\}_j$ is bounded by $m_X + \mathcal{B}$.⁶ Our proof ensures that as follows. First, P_0 picks a PRF key K and forwards P_1 a commitment of K together with encryptions of $f = (f_0 = \text{PRF}_K(0), \dots, f_{\mathcal{B}\mathcal{M}-1} = \text{PRF}_K(\mathcal{B}\mathcal{M}-1))$ (that are encrypted using the El Gamal encryption scheme; see Section 2.4). Next, P_0 proves that it computed the sequence f correctly. This can be achieved by exploiting the closed form efficiency property of the PRF. Namely, the parties mutually compute the encryption of $\prod_{i=0}^{\mathcal{B}\mathcal{M}-1} \text{PRF}_K(i)^{z^i}$ for some random z , and then compare it with the encryption of $\prod_{i=0}^{\mathcal{B}\mathcal{M}-1} f_i^{z^i}$. In particular, the latter computation is carried out on the ciphertexts that encrypt the corresponding values from f by utilizing the homomorphic property of El Gamal. Then, equality is verified such that P_0 proves that the two ciphertexts encrypt the same value. Finally, the parties divide the vector of ciphertexts f with the polynomials coefficients $\{C_j(\cdot)\}_{j \in [\mathcal{B}]}$ component-wise (note that both vectors have the same length). P_0 then proves that the overall degrees of the polynomials is as required using a sequence of zero-knowledge proofs. The last part of our proof borrows ideas from [HN12]. We continue with the formal description of our protocol.

⁵See Footnote 3.

⁶For technical reasons, we require that in case of an empty bin, P_0 fixes the polynomial that is associated with this bin to be 1.

Protocol 6 (Protocol π_{Bins} with malicious security.)

- **Input:** Party P_0 is given a PRF key K for function PRF. Both parties are given a security parameter 1^n , masked polynomials $\{C_j(\cdot) = (c_0^j, \dots, c_{\mathcal{M}-1}^j)\}_{j \in [\mathcal{B}]}$, (\mathbb{G}, p, g) for a group description \mathbb{G} of prime order p and a generator g , and an integer m_X .
- **The protocol:**
 1. **Setup.** P_0 generates $(\text{PK}, \text{SK}) \leftarrow \text{Gen}(1^n)$ for the El Gamal encryption scheme for group \mathbb{G} . It then computes the set $f = (f_0 = \text{PRF}_K(0), \dots, f_{\mathcal{B}\mathcal{M}-1} = \text{PRF}_K(\mathcal{B}\mathcal{M}-1))$ and sends to P_1 their encryptions under PK, denoted by $(e_0, \dots, e_{\mathcal{B}\mathcal{M}-1})$, as well as PK.
 2. **Proving the correctness of f .** The parties pick $z \leftarrow \mathbb{Z}_p$ at random and then compute $e_f = \prod_{i=0}^{\mathcal{B}\mathcal{M}-1} e_i^{z^i}$. Next, the parties compute the encryption of the product $\prod_{i=0}^{\mathcal{B}\mathcal{M}-1} \text{PRF}_K(i)^{z^i}$, denoted by e_{PRF} , which corresponds to the closed form efficiency function of PRF. Finally, P_0 proves that the two ciphertexts encrypt the same plaintext by proving that e_f / e_{PRF} is a Diffie-Hellman tuple using π_{DL} (see Section 2.5).
 3. **Proving a bound m_X on the overall degrees.** If π_{DL} is verified correctly, the parties compute the differences with respect to the masked polynomials $\{C_j(\cdot)\}_j$ and plaintexts f , component-wise. Namely, for all $j \in [\mathcal{B}]$ and $i \in [0, \mathcal{M}-1]$ the parties compute the encryption of $c_i^j / f_{(j-1)\mathcal{M}+i}$. We denote the result vector of ciphertexts by c_{Diff} .
 P_0 then sets $Z_{i,j} = 1$ for $0 \leq i \leq \deg(Q_j(\cdot))$, and otherwise $Z_{i,j} = 0$. P_0 computes $z_{i,j} = \text{Enc}_{\text{PK}}(Z_{i,j})$ and sends $\{z_{i,j}\}_{i,j}$ to P_1 . P_0 proves that $Z_{0,j}, Z_{1,j}, \dots, Z_{\mathcal{M}-1,j}$ is monotonically non-increasing. For that, P_0 and P_1 compute encryptions of $Z_{i,j} - Z_{i+1,j}$ and $Z_{i,j} - Z_{i+1,j} - 1$, and P_0 proves that $Z_{i,j} - Z_{i+1,j} \in \{0, 1\}$ by showing that one of these encryptions denotes a Diffie-Hellman tuple using π_{DDH} .
 P_0 completes the proof that the values $Z_{i,j}$ were constructed correctly by proving for all i, j that one of the encryptions $\{e_{(j-1)\mathcal{M}+i}, z'_{i,j}\}$ is an encryption of zero, where $z'_{i,j}$ is an encryption of $1 - Z_{i,j}$.⁷
Finally, to prove that the sum of degrees of the polynomials $\{Q_j(\cdot)\}$ equals m_X , both parties compute an encryption τ of $T = \sum_{i,j} Z_{i,j} - \mathcal{B} - m_X$. Then P proves that $(\text{PK}, \text{Enc}_{\text{PK}}(T))$ is a Diffie-Hellman tuple using π_{DDH} .
 4. **Checking zero polynomials.** If all the proofs are verified correctly, then for any $j \in [\mathcal{B}]$ the parties compute $C_j(z_j)$ where $z_j \leftarrow \mathbb{Z}_p$. The parties call $\mathcal{F}_{\text{EqMaskHash}}$ with inputs $(K, \{z_j, j, C_j(z_j)\}_{j \in [\mathcal{B}]})$. Let $\{b_j\}_{j \in [\mathcal{B}]}$ be P_1 's output from this ideal call.⁸
 5. P_1 outputs $b = 1$ only if $b_j = 0$ for all j .

Theorem 4.4 Assume that El Gamal is IND-CPA, then Protocol 6 securely realizes functionality $\mathcal{F}_{\text{Bins}}$ in the presence of malicious adversaries in the $\{\mathcal{F}_{\text{DL}}, \mathcal{F}_{\text{DDH}}, \mathcal{F}_{\text{EqMaskHash}}\}$ -hybrid model.

Proof Sketch: Security is easily stated due to the security of the zero-knowledge proofs and functionality $\mathcal{F}_{\text{EqMaskHash}}$. Specifically, the protocol enables P_1 to verify the following two properties on the masked polynomials $\{C_j(\cdot) = (c_0^j, \dots, c_{\mathcal{M}-1}^j)\}_{j \in [\mathcal{B}]}$. First, that the overall number of polynomial degrees is $m_X + \mathcal{B}$, which implies that P_0 used at most m_X set elements when creating these polynomials. This is verified by ensuring that f was correctly constructed and by the soundness of the zero knowledge proofs π_{DL} and π_{DDH} . Specifically, if one of the elements in f is incorrect then P_0 will be caught with all but negligible probability. Denote by c_i the ciphertext that encrypts f_i and say that $f_\tau \neq \text{PRF}_K(\tau)$ for some $\tau \in [0, \dots, \mathcal{B}\mathcal{M}-1]$. Then the probability that

$$\prod_{i=0}^{\mathcal{B}\mathcal{M}-1} \text{PRF}_K(i)^{z^i} = \prod_{i=0}^{\mathcal{B}\mathcal{M}-1} f_i^{z^i}$$

⁷We wish to avoid the case where $e_{(j-1)\mathcal{M}+i}$ is an encryption of a non-zero value while $z'_{i,j}$ encrypts zero.

⁸Note that z_j may be an element that is not mapped to the j th bin.

implies the probability that

$$f_\tau = \left(\prod_{i=0}^{\mathcal{B}\mathcal{M}-1} \text{PRF}_K(i) z^i - \prod_{i=0, i \neq \tau}^{\mathcal{B}\mathcal{M}-1} f_i^{z^i} \right) (z^\tau)^{-1}$$

which holds with negligible probability due to the random choice of z . Next, in Step 3, P_0 proves that it constructed the sequence of ciphertexts $z_{i,j}$ correctly by demonstrating that non-zero values are always mapped to ciphertexts that encrypt 1 and that this sequence is a monotonically non-increasing.

Second, P_0 proves that none of the masked polynomials is the zero polynomial. This is shown by substituting a random element z_j in $C_j(\cdot)$. Note that z_j is a root of $C_j(\cdot)$ only with negligible probability. Therefore, the probability that $C_j(z_j) = 0$ condition on the even that z_j is not a root of $C_j(\cdot)$ is also negligible due to the random choice of z_j . ■

The efficiency of our protocol is dominated by Steps 2 and 4, where in the former step the parties compute the closed form efficiency relative to the set f in time $O(\mathcal{B}\mathcal{M}) = O(m_X)$ and in the latter step the parties substitute a random element in every polynomial C_j . Overall, the overhead of this step relative to PRF \mathcal{PRF}_1 implies $O(\mathcal{B}) = O(m_X)$ group elements and modular exponentiations. For PRF \mathcal{PRF}_2 this step implies $O(\mathcal{B} \log m_X) = O((m_X \setminus \log \log m_X) \cdot \log m_X)$ cost; see a discussion below.

4.1.2 A Secure Protocol for $\mathcal{F}_{\text{EqMaskHash}}$

The next protocol is designed in order to compare the result of P_1 's polynomial evaluations on the set Y with the masking polynomials. Basically, for every $y \in Y$, P_1 computes first $C_{h(y)}(y)$. The parties then run a protocol for comparing $\{C_{h(y)}(y)\}_{y \in Y}$ with $\{\tilde{R}_{h(y)}(y)\}_{y \in Y}$. To do so, P_1 must also input the value $h(y)$ which determines the bin's name. Nevertheless, we do not require from the parties to mutually compute $h(y)$ since that would imply a far less efficient protocol. Instead, we demonstrate that P_1 cannot learn additional information by entering an inconsistent bin number. Finally, for every j , P_1 outputs 1 only if equality holds.

More formally, we define functionality $\mathcal{F}_{\text{EqMaskHash}}$ by $(K, \{y, h(y), C_{h(y)}(y)\}_{y \in Y}) \mapsto (-, \{b_j\}_j)$, where $b_j = 1$ if $C_{h(y)}(y) = \prod_{i=0}^{h(y)\mathcal{M}-1} \text{PRF}_K(i) y^i / \prod_{i=0}^{(h(y)-1)\mathcal{M}-1} \text{PRF}_K(i) y^i$. The actual implementation of this functionality depends on the underlying PRF. We consider two different implementations here. First, considering our protocol from Section 3.3 designed for \mathcal{PRF}_1 , an analogue protocol for our purposes can be similarly designed with the modification that the parties now compare $C_{h(y)}(y)$ against the result of the following formula evaluation,

$$g^{\frac{k_0 \binom{(k_1 x)^{(h(y)+1)\mathcal{M}-1} - (k_1 x)^{h(y)\mathcal{M}-1}}{k_1 x - 1}}{k_1 x - 1}}$$

where $h(y)$ is only known to P_1 . Note that our protocol from Section 3.3 does not need to rely on the fact that both parties know the polynomial degree d for computing this formula. It is sufficient to prove that the computation of some ciphertext c to the power of $h(y)$ is consistent with a ciphertext encrypting $g^{h(y)}$, where such a ciphertext can be provided by P_1 . See this protocol from Section 3.3 and the ZK proof π_{Eq} for more details. The overall overhead of the modified protocol is also constant.

Next, considering the unmasking protocol for function \mathcal{PRF}_2 , the parties compute the following formula that corresponds to the masking of the polynomial that is associated with bin $h(y)$,

$$g^{k_0 (1+k_1 x) (1+k_2 x^2) \dots (1+k_m x^{2^{\log((h(y)\mathcal{M}-1)})})} / g^{k_0 (1+k_1 x) (1+k_2 x^2) \dots (1+k_m x^{2^{\log((h(y)-1)\mathcal{M}-1)})}}$$

Note that computing this formula requires $O(\log m_X)$ exponentiations on the worst case if the bin number implies a high value so that $h(y)\mathcal{M}$, which determines the polynomial degree, is $O(m_X)$.

Security is stated as follows. If P_0 is corrupted then security follows similarly to the security proof of the protocols implementing $\mathcal{F}_{\text{MaskPoly}}$ (Section 3.3) since P_0 enters the same input for both functionalities and runs the same computations with respect to its PRF key. The interesting and less trivial corruption case is of P_1 . We consider two bad events here: (1) A corrupted P_1 enters y, h' for which $h' \neq h(y)$. This implies that the parties will not compute the correct unmasking. (2) A corrupted P_1 enters consistent $y, h(y)$, but an incorrect value $C_{h(y)}(y)$. Note that upon extracting P_1 's input to the protocol execution, the simulator can always tell whether this input corresponds to the first or the second case, or neither.

Specifically, in the first case the parties compute the unmasking on y for which element y is not allocated to the specified bin h' . This implies that P_1 would always obtain 0 from the protocol execution unless it correctly guesses $\tilde{R}_{h'}(y)$, which only occurs with a negligible probability due to the security of the PRF. Therefore we can successfully simulate this case by always returning zero. We further note that the security argument of the later case boils down to the security presented in the proof for a single polynomial shown in the proof of Theorem 4.2, since in this case P_1 enters $h(y)$ that is consistent with y so the parties compute the correct masking for y .

4.1.3 Using More than One Hash Function

In some cases, such as for balanced allocation hash function [ABKU99], better performance are obtained by using a pair of hash functions h_1, h_2 , which allocate elements into two distinct bins. That is, the input to the functionality are defined by $(K, \{y, h_1(y), h_2(y), C_{h(y)}(y)\}_{y \in Y}) \mapsto (-, \{b_j\}_j)$. This poses a problem in our setting since a corrupted P_1 may deviate from the protocol by substituting a different element with respect to each hash function, and learn some information about P_0 's input. Specifically, if P_1 learns that some element $y \in X$ was not allocated to $h_1(y)$ it can conclude that P_0 has \mathcal{M} additional elements that are already mapped into bin $h_1(y)$. Note that this leaked information cannot be simulated since it depends on the real input X . In this case we need to verify that P_1 indeed maps the same element into both bins correctly. A simple observation shows that if this is not the case then the simulation fails only for elements that are in the intersection. Meaning, there exists a bin for which the membership result is positive (since otherwise the adversary anyway learns 0, and it cannot distinct the cases of non-membership and incorrect behaviour). We thus define the polynomials slightly different, forcing correct behaviour.

Specifically, the polynomial $Q_j(\cdot)$ that is associated with the set of elements \mathcal{B}_j (namely, the elements that are mapped to the j th bin) is defined as follows. For each $x \in \mathcal{B}_j$, set $Q_j(x) = g^{h_1(x)+h_2(x)}$ where $h_1(x)$ and $h_2(x)$ are viewed as elements in \mathbb{Z}_p . Next, in the unmasking phase, for any tuple (y, h_1, h_2, C_y) entered by P_1 , the parties compare C_y with both $(\prod_{i=0}^{h_1 \mathcal{M}-1} \text{PRF}_K(i)^{y^i} / \prod_{i=0}^{(h_1-1)\mathcal{M}-1} \text{PRF}_K(i)^{y^i}) \cdot g^{h_1 \cdot h_2}$ and $(\prod_{i=0}^{h_2 \mathcal{M}-1} \text{PRF}_K(i)^{y^i} / \prod_{i=0}^{(h_2-1)\mathcal{M}-1} \text{PRF}_K(i)^{y^i}) \cdot g^{h_1 \cdot h_2}$ such that the functionality returns 1 to P_1 if equality holds with respect to one of the comparisons. Therefore, P_1 will learn that an element $y \in X$ only if it entered h_1 and h_2 such that $h_1 + h_2 = h_1(y) + h_2(y)$. Note that this implies that if one of the h_1, h_2 values is inconsistent with $h_1(y), h_2(y)$ yet equality holds, then the other value is also inconsistent with high probability. In this case, P_1 's output will always be 0 since the incorrect polynomials will be unmasked.

We further need to claim that for any $y \notin X$ the protocol returns 0 with overwhelming probability. Specifically, we need to claim that the probability that either $Q_{h_1(y)}(y) = g^{h_1(y)+h_2(y)}$ or $Q_{h_2(y)}(y) = g^{h_1(y)+h_2(y)}$, is negligible. In order to simplify our argument, we slightly modify our construction and fix $Q_j(x) = \text{PRF}_K(g^{h_1(x)+h_2(x)})$ for any $x \in \mathcal{B}_j$ using a PRF K key that is mutually picked by both parties. In this case, we can easily claim that the probability that the protocol returns 1 for $y \notin X$ is negligible since that implies that either $Q_{h_1(y)}(y)$ or $Q_{h_2(y)}(y)$ equal the pseudorandom value $\text{PRF}_K(g^{h_1(y)+h_2(y)})$ for $y \notin X$. We stress that the PRF key for this purpose can be publicly known since pseudorandomness is still maintained as long as the algorithm for generating the bin polynomials does not use this key. We further note

that both algebraic PRFs that we consider in this paper can be easily evaluated over an encrypted plaintext given the PRF key since it only require linear operations.

Finally, a similar solution can be easily adapted for Cuckoo hashing with a stash [KMW08] (by treating the stash as a third polynomial). Nevertheless, Cuckoo hashing using a stash suffers from the following drawback. It has been proven in [KMW08] that for any constant s , using a stash of size s implies an overflow with probability $O(n^s)$ (taken over the choice of the hash functions). Specifically, if the algorithm aborts whenever the original choice of hash functions results in more than s items being moved to the stash, then this means that the algorithm aborts with probability of at most $O(n^s)$. Consequently, P_1 can identify with that probability whether a specific potential input of P_0 does not agree with the hash functions h_1 and h_2 . This probability is low but not negligible. On the other hand, Broder and Mitzenmacher [BM01] have shown for balanced allocations hash function that asymptotically, when mapping n items into n bins, the number of bins with i or more items falls approximately like $2^{2.6i}$. This means that if $\mathcal{M} = \omega(\log \log n)$ then except with negligible probability no bin will be of size greater than \mathcal{M} . Nevertheless, (ignoring the abort probability), Cuckoo hashing performs better than balanced allocation hash functions, and by tuning the parameters accordingly this abort probability can be ignored for most practical applications.

4.1.4 Efficiency

The efficiency of our protocol depends on the parameters $\mathcal{B} = O(m_X / \log \log m_X)$ and $\mathcal{M} = O(\log \log m_X)$ that are specified by the underlying hash function, as well as the PRF implementation that induce the overhead of the implementations of $\mathcal{F}_{\text{Bins}}$ and $\mathcal{F}_{\text{EqMaskHash}}$. Concretely, when implementing the algebraic PRF with \mathcal{PRF}_1 the number of exponentiations computed by the parties is $O(\mathcal{B}\mathcal{M} + m_Y\mathcal{M}) = O(m_X + m_Y \log \log m_X)$, whereas the number of transmitted group elements is $O(\mathcal{B}\mathcal{M} + m_Y) = O(m_X + m_Y)$. Moreover, implementing the algebraic PRF using \mathcal{PRF}_2 implies the overhead of $O(m_X + m_Y \log \log m_X)$ exponentiations and the communication is as above.

5 Committed Oblivious PRF Evaluation

The oblivious PRF evaluation functionality \mathcal{F}_{PRF} is an important functionality that is defined by $(K, x) \mapsto (-, \text{PRF}_K(x))$. One known example for a protocol that implements \mathcal{F}_{PRF} is the instantiation based on the Naor-Reingold pseudorandom function [NR97] (specified in Section 3.1.2), that is implemented by the protocol presented in [FIPR05] (and proven secure in the malicious setting in [HL10a]). This protocol involves executing an oblivious transfer for every bit of the input x . Nevertheless, it has major drawback since it does not enforce the usage of *the same* key for multiple evaluations, which is often required. In this section, we observe first that the algebraic closed form efficiency of PRFs \mathcal{PRF}_1 and \mathcal{PRF}_2 , specified in Section 3.1, are PRFs as well. Moreover, the protocols for securely evaluating these functions induce efficient implementations for the *committed* oblivious PRF evaluation functionality with respect to these new PRFs in the presence of adaptive inputs. This is because the PRF evaluations protocols are implemented with respect to the same set of key commitments. We formally define this functionality in Figure 1.

More formally, let PRF be an algebraic PRF from a domain $\{0, 1\}^m$ into a group \mathbb{G} . Then, define the new function $\text{PRF}' : \mathbb{Z}_p \mapsto \mathbb{G}$ by $\text{PRF}'_K(x) = \prod_{i=0}^l [\text{PRF}_K(i)]^{x^i}$. Note that the domain size of PRF' is bounded by $l + 1$, since upon observing $l + 1$ evaluations of PRF' it is possible to interpolate the coefficients of the polynomial $\{\text{PRF}_K(i)\}_i$ (in the exponent). On the other hand, it is easy to verify that if $l + 1 \leq 2^m$ then PRF' is a PRF.

Theorem 5.1 *Assume that $F : \{0, 1\}^m \mapsto \mathbb{G}$ is a PRF, then PRF' is a PRF for $(l + 1) \leq 2^m$.*

Functionality $\mathcal{F}_{\text{CPRF}}$

Functionality $\mathcal{F}_{\text{CPRF}}$ communicates with parties P_0 and P_1 , and adversary \mathcal{SIM} .

1. Upon receiving a message (key, K) from P_0 , send message key to \mathcal{SIM} and record K .
2. Upon receiving (input, x) from P_1 , send message input to adversary \mathcal{SIM} . Upon receiving an approve message, send $\text{PRF}_K(x)$ to P_1 . Otherwise, send \perp to P_1 and abort.

Figure 1: The committed oblivious PRF evaluation functionality.

Proof: Observe first that the $\text{PRF}_K(i)$ values within PRF' can be replaced with truly random values in \mathbb{G} due to the pseudorandom property of PRF. Namely, consider a hybrid game and a function \tilde{F}' such that instead of evaluating $\text{PRF}_K(i)$ for a new input i , a random $r_i \leftarrow \mathbb{G}$ is picked. This implies that $\tilde{\text{PRF}}'_K(x) = \prod_{i=0}^l r_i^{x^i}$. Note that for $(l+1) \leq 2^m$ the outputs of PRF' and \tilde{F}' are indistinguishable. Thus, the proof boils down to a probabilistic argument claiming that the probability an adversary distinguishes the hybrid function from a truly random function is negligible since the distributions of PRF' with random r_i 's and a truly random function are identically distributed. ■

We implement PRF' using the two PRFs from Section 3.1.1 and obtain two new PRF constructions under the strong-DDH and DDH assumptions. Let $K = (k_0, k_1)$ be the key for the PRF \mathcal{PRF}_1 with the strong-DDH based security, and recall that the closed form efficiency for this function is defined by

$$\text{PRF}'_K(x) = \text{CFEval}_h(x, K) = g^{\frac{k_0(k_1^{d+1}x^{d+1}-1)}{k_1x-1}}.$$

This implies that PRF' only requires a constant number of modular exponentiations. See Section 3.3 for secure implementations of obviously evaluating PRF' . Next, let $K = (k_0, \dots, k_m)$ be the key for the Naor-Reingold PRF, and recall that the closed form efficiency of this function is defined by

$$\text{PRF}'_K(x) = \text{CFEval}_{h,z}(x, K) = g^{k_0(1+k_1x)(1+k_2x^2)\dots(1+k_mx^{2^m})}$$

which requires $O(\log l) = O(m)$ operations, namely, a logarithmic number of operations in the domain size where x is an m -bits string. This is the same order of overhead induced by the [FIPR05] implementation that requires an OT for each input bit. Nevertheless, our construction has the advantage that it also achieves easily the property of a committed key since multiple evaluations can be computed with respect to the same committed PRF key. Plugging-in our protocol inside the protocols for keyword search, OT with adaptive queries [FIPR05] and set-intersection [HL10a] implies security against malicious adversaries fairly immediately. It is further useful for search functionalities as demonstrated below.

5.1 The Set-Intersection Protocol

We continue with describing our set-intersection protocol. Informally, P_0 generates a PRF key for PRF and evaluates this function on its set X . It then sends the evaluation results to P_1 and the parties engage in a committed oblivious PRF protocol that evaluates PRF on the set Y . P_1 then concludes the intersection. In order to handle a technicality in the security proof, we require that P_0 must generate its PRF key *independently* of its input X , since otherwise it may maliciously pick a secret key that implies collisions on elements from X and Y , causing the simulation to fail. We ensure key independence by asking the parties to mutually generate the PRF key after P_0 has committed to its input. Then upon choosing the PRF key, the parties invoke two variations of functionality $\mathcal{F}_{\text{CPRF}}$, denoted by $\mathcal{F}_{\text{CPRF}}^0$ and $\mathcal{F}_{\text{CPRF}}^1$. Formally, we define functionality $\mathcal{F}_{\text{CPRF}}^0$

as follows: $((K, (x_1, \dots, x_{m_X}), R), (c_{\text{KEY}}, (c_1, \dots, c_{m_X}), \text{PK})) \mapsto (-, (\text{PRF}_K(x_1), \dots, \text{PRF}_K(x_{m_X})))$ only if c_i encrypts x_i for all i and c_{KEY} is a commitment of K where verification is carried out using randomness R . In the final step, the parties call functionality $\mathcal{F}_{\text{CPRF}}^1$ in order to evaluate the PRF on the set Y and is defined by $((K, R), (c_{\text{KEY}}, (y_1, \dots, y_{m_Y}))) \mapsto (-, (\text{PRF}_K(y_1), \dots, \text{PRF}_K(y_{m_Y})))$ only if c_{KEY} is a commitment of K where verification is carried out using randomness R . In both executions the output is given to P_1 that computes the intersection of the results.

Implementing $\mathcal{F}_{\text{CPRF}}^0$ and $\mathcal{F}_{\text{CPRF}}^1$. Implantation-wise, there is not much of a difference between the protocols for the two functionalities, which mainly differ due to the identity of the party that enters the input values to the PRF (where the same committed key is used for both protocol executions). We note that the realization of $\mathcal{F}_{\text{CPRF}}^0$ and $\mathcal{F}_{\text{CPRF}}^1$ can be carried out securely based on the implementations of the closed form efficiency functions shown in Section 3.3, since our committed PRFs are based on these functions. More concretely, the difference with respect to functionality $\mathcal{F}_{\text{CPRF}}^0$ is that now when P_0 is corrupted the simulator needs to extract the randomness used for committing to the PRF key and the x_i 's elements which can be achieved using the proof of knowledge π_{DL} since the parties use the El Gamal PKE. Specifically, P_0 proves the knowledge of the discrete logarithm of (c_1, \dots, c_{m_X}) with respect to a generator g , by invoking an ideal execution of \mathcal{F}_{DL} on input $\{((g, c_i), \log_g c_i)\}_{i \in [m_X]}$.⁹ The input of P_1 for \mathcal{F}_{DL} is $\{(g, c_i)\}_{i \in [m_X]}$. In case P_1 does not receive an “accept” message from \mathcal{F}_{DL} it aborts. Next, the parties continue with the PRF evaluations where the ZK proofs are carried out with respect to the same key commitment. We note that extracting the PRF key and the set (x_1, \dots, x_{m_X}) is already implied by the protocols from Section 3.3 due to the ZK proofs of knowledge. Finally, the implementation of $\mathcal{F}_{\text{CPRF}}^1$ follows similarly but without the additional proof we added above for $\mathcal{F}_{\text{CPRF}}^0$ in order to extract the randomness of the committed input.

Next, we describe our set-intersection protocol using committed oblivious PRF.

Protocol 7 (Protocol π_{\cap} with malicious security from committed oblivious PRF.)

- **Input:** Party P_0 is given a set X of size m_X . Party P_1 is given a set Y of size m_Y . Both parties are given a security parameter 1^n .
- **The protocol:**
 1. **Distributed key generation.** P_0 and P_1 run protocol $\pi_{\text{KeyGen}}(1^n, 1^n)$ in order to generate additive El Gamal public key $\text{PK} = \langle \mathbb{G}, p, g, h \rangle$ where the corresponding shares of the secret key SK are $(\text{SK}_0, \text{SK}_1)$.
 2. **Input commitment and PRF key generation.** P_0 sends encryptions of its input X under PK ; denote this set of ciphertexts by $C = (c_1, \dots, c_{m_X})$.
 P_0 invokes $(K, \text{param}) \leftarrow \text{KeyGen}(1^n, d = \log(m_X + m_Y))$ where param includes a group description \mathbb{G} of prime order p and a generator g , and sends P_1 param and a ciphertext $\text{Enc}_{\text{PK}}(K; R)$.
 P_1 picks a new key $(K', \text{param}) \leftarrow \text{KeyGen}(1^n, d = \log(m_X + m_Y))$ and sends it to P_0 . The parties then compute the encryption c_{KEY} of $\tilde{K} = KK'$, relying on the homomorphic property of El Gamal.
 3. **PRF evaluations on X.** The parties call functionality $\mathcal{F}_{\text{CPRF}}^0$ where P_0 enters the set X , key \tilde{K} and randomness R and P_1 enters C, c_{KEY} and PK . Denote by $\text{PRF}_X = \{\text{PRF}'_{\tilde{K}}(x)\}_{x \in X}$ the output of P_1 from this ideal call only if C is a vector of ciphertexts that encrypts X and c_{KEY} is a commitment of \tilde{K} , where verification is computed using randomness R .
 4. **Oblivious PRF evaluations on Y.** The parties call functionality $\mathcal{F}_{\text{CPRF}}^1$ where P_0 enters the key \tilde{K} and randomness R and P_1 enters the commitment c_{KEY} and the set Y . Denote by $\text{PRF}_Y = \{f_y\}_{y \in Y}$ the output of P_1 from this ideal call only if c_{KEY} is a commitment of \tilde{K} where verification is computed using randomness R .
 P_1 outputs all $y \in Y$ for which $f_y \in \text{PRF}_X$.

⁹We abuse notation and write $\log c$ to denote the discrete logarithm of the two group elements in ciphertext c .

Theorem 5.2 Assume $\text{PRF}'_K(\cdot)$ is a PRF defined as above and that El Gamal is IND-CPA, then Protocol 7 securely realizes functionality \mathcal{F}_\cap in the presence of malicious adversaries in the $\{\mathcal{F}_{\text{DL}}, \mathcal{F}_{\text{CPRF}}^0, \mathcal{F}_{\text{CPRF}}^1\}$ -hybrid model.

Proof: We prove security for each corruption case separately. We assume that the simulator is given m_X and m_Y as part of its auxiliary input.

P_0 is corrupted. Let \mathcal{A} be a PPT adversary corrupting party P_0 , we design a PPT simulator \mathcal{SIM} that generates the view of \mathcal{A} as follows.

1. Given $(1^n, X, z)$, \mathcal{SIM} engages in an execution of $\pi_{\text{KeyGen}}(1^n, 1^n)$ with \mathcal{A} . Denote the outcome by PK.
2. Upon receiving from \mathcal{A} its commitment for the PRF key $K \leftarrow \text{KeyGen}(1^n, d = \log(m_X + m_Y))$, \mathcal{SIM} picks a new key share K' and sends it to \mathcal{A} using PK. Denote the combined key by $\tilde{K} = KK'$.
3. \mathcal{SIM} extracts the adversary's input X' from the input to the ideal call $\mathcal{F}_{\text{CPRF}}^0$. It then sends X' to the trusted party and completes the execution as would the honest P_1 do on an arbitrary set.

In the hybrid setting, computational indistinguishability between the hybrid and simulated executions is trivially claimed since the adversary does not receive any message from P_1 that depends on Y . An important observation here is that the probability of the event for which there exists $y \in Y$ such that $y \notin X'$ and yet $\text{PRF}_{\tilde{K}}(y) \in \text{PRF}_{X'}$ is negligible, since the key \tilde{K} is picked independently of the set X' . This argument follows from similarly to the proof in [HL10a] and implies that P_1 's output in both executions is identical condition that the above event does not occur.

P_1 is corrupted. Let \mathcal{A} be a PPT adversary corrupting party P_1 , we design a PPT simulator \mathcal{SIM} that generates the view of \mathcal{A} as follows.

1. Given $(1^n, Y, z)$, \mathcal{SIM} engages in an execution of $\pi_{\text{KeyGen}}(1^n, 1^n)$ with \mathcal{A} . Denote the outcome by PK.
2. \mathcal{SIM} picks a PRF key share $K \leftarrow \text{KeyGen}(1^n, d = \log(m_X + m_Y))$ and sends its encryption to \mathcal{A} using PK. Upon receiving \mathcal{A} 's key share K' the simulator sets the combined key by $\tilde{K} = KK'$.
3. \mathcal{SIM} picks a set of m_X arbitrary elements $X_{\mathcal{SIM}}$ from \mathbb{Z}_p . It then emulates the ideal call $\mathcal{F}_{\text{CPRF}}^0$ and hands the adversary a random set U of size m_X and proper length.
4. Finally, the simulator extracts the adversary's input Y' to the ideal call $\mathcal{F}_{\text{CPRF}}^1$ and sends this set to the trusted party, receiving back $Z = X \cap Y'$. The simulator completes the execution as follows. For each element $y' \in Y' \cap Z$ it programs the ideal answer of $\mathcal{F}_{\text{CPRF}}^1$ to be $r \in U$ where r is picked from the remaining elements from the set U that were not picked thus far. Otherwise, the simulator returns a fresh random element from \mathbb{Z}_p .

Security here follows from the IND-CPA security of the El Gamal PKE and the security of the PRF. That is, the simulated view is different from the hybrid view relative to the encrypted input of P_0 and the fact that the simulator uses a random function to evaluate the sets in $X'_{\mathcal{SIM}}$ and Y' . Therefore, the proof can be shown by defining a hybrid game where in the first game the simulator encrypts P_0 's real input X but completes the simulation as in the original simulation. Indistinguishability between the simulation and the hybrid game follows easily by a reduction to the IND-CPA security of El Gamal since the simulator never uses the secret key of the encryption scheme. Indistinguishability between the hybrid game and the hybrid execution follows by a reduction to the pseudorandomness of the PRF. ■

Efficiency. The overhead of protocol 7 depends on the implementations of $\mathcal{F}_{\text{CPRF}}^0$ and $\mathcal{F}_{\text{CPRF}}^1$ discussed above. Our protocol obtains $O(m_X + m_Y)$ communication and computation overheads under the strong-DDH assumption and $O((m_X + m_Y) \log(m_X + m_Y))$ overheads under the DDH assumption, where the former analysis matches the [JL09] analysis (such that both constructions rely on dynamic assumptions).

5.2 Search Functionalities

In search functionalities a receiver searches in a sender’s database, retrieving the appropriate record(s) according to some search query. The database for search functionalities can be described by pairs of queries/records $\{(q_i, T_i)\}_i$ such that the answer to a query q_i is a record T_i .¹⁰ In a private setting we need to ensure that nothing beyond these records leaks to the receiver, while the sender does not learn anything about the receiver’s search queries. Committed oblivious PRF evaluation is a useful tool for securely implementing various search functionalities [FIPR05]. First, in the setup phase the database is encoded and handed to the receiver. That is, for each query q_i the sender defines the pair $(\text{PRF}_K(q_i\|1), \text{PRF}_K(q_i\|2) \oplus T_i)$. Next, in the query phase the parties run a committed oblivious PRF evaluation protocol twice such that the sender inputs K and the receiver inputs a query q . The receiver’s output are the values $\text{PRF}_K(q\|1)$ and $\text{PRF}_K(q\|2)$, where the first outcome is used to find the encrypted record while the second outcome is used to extract the record. (Alternative implementations involve a single invocation of PRF by splitting $\text{PRF}_K(q)$ into two parts). Examples for such functionalities are keyword search, oblivious transfer with adaptive queries and pattern matching (and all its variants). The functionality of committed oblivious PRF is important in this context since the sender must be enforced to use the same PRF key.

References

- [ABKU99] Yossi Azar, Andrei Z. Broder, Anna R. Karlin, and Eli Upfal. Balanced allocations. *SIAM J. Comput.*, 29(1):180–200, 1999.
- [ACT10] Giuseppe Ateniese, Emiliano De Cristofaro, and Gene Tsudik. (if) size matters: Size-hiding private set intersection. *IACR Cryptology ePrint Archive*, 2010:220, 2010.
- [ADDV12] Giuseppe Ateniese, Özgür Dagdelen, Ivan Damgård, and Daniele Venturi. Entangled cloud storage. *IACR Cryptology ePrint Archive*, 2012:511, 2012.
- [AMP04] Gagan Aggarwal, Nina Mishra, and Benny Pinkas. Secure computation of the k th-ranked element. In *EUROCRYPT*, pages 40–55, 2004.
- [BDOZ11] Rikke Bendlin, Ivan Damgård, Claudio Orlandi, and Sarah Zakarias. Semi-homomorphic encryption and multiparty computation. In *EUROCRYPT*, pages 169–188, 2011.
- [Bea91] Donald Beaver. Foundations of secure interactive computing. In *CRYPTO*, pages 377–391, 1991.
- [BG12] Stephanie Bayer and Jens Groth. Efficient zero-knowledge argument for correctness of a shuffle. In *EUROCRYPT*, pages 263–280, 2012.
- [BGV11] Siavosh Benabbas, Rosario Gennaro, and Yevgeniy Vahlis. Verifiable delegation of computation over large datasets. In *CRYPTO*, pages 111–131, 2011.
- [BM01] Andrei Z. Broder and Michael Mitzenmacher. Using multiple hash functions to improve IP lookups. In *INFOCOM*, pages 1454–1463, 2001.
- [Can00] Ran Canetti. Security and composition of multi-party cryptographic protocols. *Journal of Cryptology*, 13:143–202, 2000.

¹⁰This may not be the most concise description of the database but it is the simplest. In particular, it will do for our purposes.

- [CKT10] Emiliano De Cristofaro, Jihye Kim, and Gene Tsudik. Linear-complexity private set intersection protocols secure in malicious model. In *ASIACRYPT*, pages 213–231, 2010.
- [CL05] Yan-Cheng Chang and Chi-Jen Lu. Oblivious polynomial evaluation and oblivious neural learning. *Theor. Comput. Sci.*, 341(1-3):39–54, 2005.
- [CP92] David Chaum and Torben P. Pedersen. Wallet databases with observers. In *CRYPTO*, pages 89–105, 1992.
- [CT10] Emiliano De Cristofaro and Gene Tsudik. Practical private set intersection protocols with linear complexity. In *Financial Cryptography*, pages 143–159, 2010.
- [DCW13] Changyu Dong, Liqun Chen, and Zikai Wen. When private set intersection meets big data: An efficient and scalable protocol. *IACR Cryptology ePrint Archive*, 2013:515, 2013.
- [DJN10] Ivan Damgård, Mads Jurik, and Jesper Buus Nielsen. A generalization of paillier’s public-key system with applications to electronic voting. *Int. J. Inf. Sec.*, 9(6):371–385, 2010.
- [DKL⁺13] Ivan Damgård, Marcel Keller, Enrique Larraia, Valerio Pastro, Peter Scholl, and Nigel P. Smart. Practical covertly secure MPC for dishonest majority - or: Breaking the SPDZ limits. In *Computer Security - ESORICS 2013 - 18th European Symposium on Research in Computer Security, Egham, UK, September 9-13, 2013. Proceedings*, pages 1–18, 2013.
- [DPSZ12] Ivan Damgård, Valerio Pastro, Nigel P. Smart, and Sarah Zakarias. Multiparty computation from somewhat homomorphic encryption. In *CRYPTO*, pages 643–662, 2012.
- [DSMRY09] Dana Dachman-Soled, Tal Malkin, Mariana Raykova, and Moti Yung. Efficient robust private set intersection. In *ACNS*, pages 125–142, 2009.
- [DY05] Yevgeniy Dodis and Aleksandr Yampolskiy. A verifiable random function with short proofs and keys. In *Public Key Cryptography - PKC 2005, 8th International Workshop on Theory and Practice in Public Key Cryptography, Les Diablerets, Switzerland, January 23-26, 2005, Proceedings*, pages 416–431, 2005.
- [FHV13] Sebastian Faust, Carmit Hazay, and Daniele Venturi. Outsourced pattern matching. In *ICALP*, 2013.
- [FIPR05] Michael J. Freedman, Yuval Ishai, Benny Pinkas, and Omer Reingold. Keyword search and oblivious pseudorandom functions. In *TCC*, pages 303–324, 2005.
- [FNP04] Michael J. Freedman, Kobbi Nissim, and Benny Pinkas. Efficient private matching and set intersection. In *EUROCRYPT*, pages 1–19, 2004.
- [Gam85] Taher El Gamal. A public key cryptosystem and a signature scheme based on discrete logarithms. *IEEE Transactions on Information Theory*, 31(4):469–472, 1985.
- [Gil99] Niv Gilboa. Two party rsa key generation. In *CRYPTO*, pages 116–129, 1999.
- [GMW87] Oded Goldreich, Silvio Micali, and Avi Wigderson. How to play any mental game or a completeness theorem for protocols with honest majority. In *STOC*, pages 218–229, 1987.
- [Gol04] Oded Goldreich. *Foundations of Cryptography: Volume 2, Basic Applications*. Cambridge University Press, New York, NY, USA, 2004.
- [HL09] Carmit Hazay and Yehuda Lindell. Efficient oblivious polynomial evaluation with simulation-based security. *IACR Cryptology ePrint Archive*, 2009:459, 2009.
- [HL10a] Carmit Hazay and Yehuda Lindell. Efficient protocols for set intersection and pattern matching with security against malicious and covert adversaries. *J. Cryptology*, 23(3):422–456, 2010.
- [HL10b] Carmit Hazay and Yehuda Lindell. *Efficient Secure Two-Party Protocols – Techniques and Constructions*. Springer-Verlag, 2010.
- [HMRT12] 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.

- [HN12] Carmit Hazay and Kobbi Nissim. Efficient set operations in the presence of malicious adversaries. *J. Cryptology*, 25(3):383–433, 2012.
- [HT10] Carmit Hazay and Tomas Toft. Computationally secure pattern matching in the presence of malicious adversaries. In *ASIACRYPT*, pages 195–212, 2010.
- [IPS08] Yuval Ishai, Manoj Prabhakaran, and Amit Sahai. Founding cryptography on oblivious transfer - efficiently. In *CRYPTO*, pages 572–591, 2008.
- [JL09] Stanislaw Jarecki and Xiaomin Liu. Efficient oblivious pseudorandom function with applications to adaptive ot and secure computation of set intersection. In *TCC*, pages 577–594, 2009.
- [JL10] Stanislaw Jarecki and Xiaomin Liu. Fast secure computation of set intersection. In *SCN*, pages 418–435, 2010.
- [KMW08] Adam Kirsch, Michael Mitzenmacher, and Udi Wieder. More robust hashing: Cuckoo hashing with a stash. In *Algorithms - ESA 2008, 16th Annual European Symposium, Karlsruhe, Germany, September 15-17, 2008. Proceedings*, pages 611–622, 2008.
- [KS05] Lea Kissner and Dawn Xiaodong Song. Privacy-preserving set operations. In *CRYPTO*, pages 241–257, 2005.
- [Lin13] Yehuda Lindell. Fast cut-and-choose based protocols for malicious and covert adversaries. In *CRYPTO* (2), pages 1–17, 2013.
- [LOP11] Yehuda Lindell, Eli Oxman, and Benny Pinkas. The ips compiler: Optimizations, variants and concrete efficiency. In *CRYPTO*, pages 259–276, 2011.
- [LP02] Yehuda Lindell and Benny Pinkas. Privacy preserving data mining. *J. Cryptology*, 15(3):177–206, 2002.
- [LP11] Yehuda Lindell and Benny Pinkas. Secure two-party computation via cut-and-choose oblivious transfer. In *TCC*, pages 329–346, 2011.
- [MR91] Silvio Micali and Phillip Rogaway. Secure computation (abstract). In *CRYPTO*, pages 392–404, 1991.
- [NNOB12] Jesper Buus Nielsen, Peter Sebastian Nordholt, Claudio Orlandi, and Sai Sheshank Burra. A new approach to practical active-secure two-party computation. In *CRYPTO*, pages 681–700, 2012.
- [NO09] Jesper Buus Nielsen and Claudio Orlandi. Lego for two-party secure computation. In *TCC*, pages 368–386, 2009.
- [NP99] Moni Naor and Benny Pinkas. Oblivious transfer and polynomial evaluation. In *STOC*, pages 245–254, 1999.
- [NP06] Moni Naor and Benny Pinkas. Oblivious polynomial evaluation. *SIAM J. Comput.*, 35(5):1254–1281, 2006.
- [NR97] Moni Naor and Omer Reingold. Number-theoretic constructions of efficient pseudo-random functions. In *FOCS*, pages 458–467, 1997.
- [PSZ14] Benny Pinkas, Thomas Schneider, and Michael Zohner. Faster private set intersection based on OT extension. In *Proceedings of the 23rd USENIX Security Symposium, San Diego, CA, USA, August 20-22, 2014.*, pages 797–812, 2014.
- [Sch89] Claus-Peter Schnorr. Efficient identification and signatures for smart cards. In *CRYPTO*, pages 239–252, 1989.
- [ST04] Berry Schoenmakers and Pim Tuijls. Practical two-party computation based on the conditional gate. In *ASIACRYPT*, pages 119–136, 2004.
- [Ver11] Damien Vergnaud. Efficient and secure generalized pattern matching via fast fourier transform. In *AFRICACRYPT*, pages 41–58, 2011.
- [Yao86] Andrew Chi-Chih Yao. How to generate and exchange secrets (extended abstract). In *FOCS*, pages 162–167, 1986.
- [ZB05] Huafei Zhu and Feng Bao. Augmented oblivious polynomial evaluation protocol and its applications. In *ESORICS*, pages 222–230, 2005.

A Secure Two-Party Computation

We briefly present the standard definition for secure multiparty computation and refer to [Gol04, Chapter 7] for more details and motivating discussions. 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 $f = (f_1, f_2)$. That is, for every pair of inputs (x, y) , the output-vector is a random variable $(f_1(x, y), f_2(x, y))$ 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 $(x, y) \mapsto (f_1(x, y), f_2(x, y))$ to describe a functionality.

We prove the security of our protocols in the settings of *semi-honest* and *malicious* computationally bounded adversaries. 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.

A.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. Let $f = (f_1, f_2)$ be a two-party functionality and let π be a two-party protocol for computing f . The *view* of the first party in an execution of π on inputs (x, y) is

$$\mathbf{View}_{\pi,1}(x, y) = (x, r_1, m_1, \dots, m_t),$$

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 π on (x, y) is denoted $\mathbf{Output}_{\pi,1}(x, y)$ and can be computed from $\mathbf{View}_{\pi,1}(x, y)$. Similarly, $\mathbf{View}_{\pi,2}(x, y) = (y, r_2, m_1, \dots, m_t)$ 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 $\mathbf{Output}_{\pi,2}(x, y)$.

Definition A.1 Let f and π be as above. Protocol π is said to securely compute f in the presence of semi-honest adversaries if there exist probabilistic polynomial-time algorithms \mathcal{SIM}_1 and \mathcal{SIM}_2 such that

$$(\mathcal{SIM}_1(x, f_1(n, x, y)), f_2(n, x, y))_{n \in \mathbb{N}, x, y \in \{0,1\}^*} \stackrel{c}{\approx} \{(\mathbf{View}_{\pi,1}(n, x, y), \mathbf{Output}_{\pi,2}(n, x, y))\}_{n \in \mathbb{N}, x, y \in \{0,1\}^*}$$

$$(f_1(n, x, y), \mathcal{SIM}_2(y, f_2(n, x, y)))_{n \in \mathbb{N}, x, y \in \{0,1\}^*} \stackrel{c}{\approx} \{(\mathbf{Output}_{\pi,1}(n, x, y), \mathbf{View}_{\pi,2}(n, x, y))\}_{n \in \mathbb{N}, x, y \in \{0,1\}^*}$$

where n is the security parameter.

A.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 $f = (f_1, f_2)$, 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 n , denoted $\mathbf{IDEAL}_{f, \mathcal{A}(z), I}(n, x, y)$, 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 π .

Let f be as above and let π 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 π* on inputs (x, y) , auxiliary input z to \mathcal{A} and security parameter n , denoted $\mathbf{REAL}_{\pi, \mathcal{A}(z), I}(n, x, y)$, is defined as the output vector of the honest parties and the adversary \mathcal{A} from the real execution of π .

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 A.2 *Let f and π be as above. Protocol π 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 SIM for the ideal model, such that for every $I \subset [2]$,*

$$\{\mathbf{IDEAL}_{f, SIM(z), I}(n, x, y)\}_{n \in \mathbb{N}, x, y, z \in \{0, 1\}^*} \stackrel{c}{\approx} \{\mathbf{REAL}_{\pi, \mathcal{A}(z), I}(n, x, y)\}_{n \in \mathbb{N}, x, y, z \in \{0, 1\}^*}$$

where n is the security parameter.

The \mathcal{F} -hybrid model. In order to construct some of our protocols, we will use secure two-party protocols as subprotocols. The standard way of doing this is to work in a “*hybrid model*” where parties both interact with each other (as in the real model) and use trusted help (as in the ideal model). Specifically, when constructing a protocol π that uses a subprotocol for securely computing some functionality \mathcal{F} , we consider the case that the parties run π and use “ideal calls” to a trusted party for computing \mathcal{F} . Upon receiving the inputs from the parties, the trusted party computes \mathcal{F} and sends all parties their output. Then, after receiving these outputs back from the trusted party the protocol π continues.

Let \mathcal{F} be a functionality and let π be a two-party protocol that uses ideal calls to a trusted party computing \mathcal{F} . Furthermore, let \mathcal{A} be a non-uniform probabilistic polynomial-time machine. Then, the *\mathcal{F} -hybrid execution of π* on inputs (x, y) , auxiliary input z to \mathcal{A} and security parameter n , denoted $\mathbf{HYBRID}_{\pi, \mathcal{F}, \mathcal{A}(z)}(n, x, y)$, is defined as the output vector of the honest parties and the adversary \mathcal{A} from the hybrid execution of π with a trusted party computing \mathcal{F} . By the composition theorem of [Can00] any protocol that securely implements \mathcal{F} can replace the ideal calls to \mathcal{F} .