

Function-Private Subspace-Membership Encryption and Its Applications*

Dan Boneh¹, Ananth Raghunathan¹, and Gil Segev^{2,**}

¹ Stanford University

{dabo, ananthr}@cs.stanford.edu

² Hebrew University

segev@cs.huji.ac.il

Abstract. Boneh, Raghunathan, and Segev (CRYPTO '13) have recently put forward the notion of *function privacy* and applied it to identity-based encryption, motivated by the need for providing predicate privacy in public-key searchable encryption. Intuitively, their notion asks that decryption keys reveal essentially no information on their corresponding identities, beyond the absolute minimum necessary. While Boneh et al. showed how to construct function-private identity-based encryption (which implies predicate-private encrypted keyword search), searchable encryption typically requires a richer set of predicates. In this paper we significantly extend the function privacy framework. First, we consider the notion of *subspace-membership* encryption, a generalization of inner-product encryption, and formalize a meaningful and realistic notion for capturing its function privacy. Then, we present a generic construction of a *function-private* subspace-membership encryption scheme based on *any* inner-product encryption scheme. This is the first generic construction that yields a function-private encryption scheme based on a non-function-private one. Finally, we present various applications of function-private subspace-membership encryption. Among our applications, we significantly improve the function privacy of the identity-based encryption schemes of Boneh et al.: whereas their schemes are function private only for identities that are highly unpredictable (with min-entropy of at least $\lambda + \omega(\log \lambda)$ bits, where λ is the security parameter), we obtain function-private schemes assuming only the *minimal* required unpredictability (i.e., min-entropy of only $\omega(\log \lambda)$ bits). This improvement offers a much more realistic function privacy guarantee.

Keywords: Function privacy, functional encryption.

1 Introduction

Predicate encryption systems [13, 23] are public-key schemes where a single public encryption key has *many* corresponding secret keys: every secret key corresponds to a predicate $p : \Sigma \rightarrow \{0, 1\}$ where Σ is some pre-defined set of indices

* The full version is available as Cryptology ePrint Archive, Report 2013/403 [11].

** This work was done while the author was visiting Stanford University.

(or attributes). Plaintext messages are pairs (x, m) where $x \in \Sigma$ and m is in some message space. A secret key sk_p for a predicate p has the following semantics: if c is an encryption of the pair (x, m) then sk_p can be used to decrypt c only if the “index” x satisfies the predicate p . More precisely, attempting to decrypt c using sk_p will output m if $p(x) = 1$ and output \perp otherwise. A predicate encryption system is secure if it provides semantic security for the pair (x, m) even if the adversary has a few benign secret keys.

The simplest example of predicate encryption is a system supporting the set of equality predicates, that is, predicates $p_{\text{id}} : \Sigma \rightarrow \{0, 1\}$ defined as $p_{\text{id}}(x) = 1$ iff $x = \text{id}$. In such a system there is a secret key sk_{id} for every $\text{id} \in \Sigma$ and given the encryption c of a pair (x, m) the key sk_{id} can decrypt c and recover m only when $x = \text{id}$. It is easy to see that predicate encryption for the set of equality predicates is the same thing as (anonymous) identity-based encryption [8, 1].

Currently the most expressive collusion-resistant predicate encryption systems [23, 3] support the family of inner product predicates: for a vector space $\Sigma = \mathbb{F}_q^\ell$ this is the set of predicates $p_v : \Sigma \rightarrow \{0, 1\}$ where $v \in \Sigma$ and $p_v(x) = 1$ iff $x \perp v$. This family of predicates includes the set of equality predicates and others.

Searching on encrypted data. Predicate encryption systems provide a general framework for searching on encrypted data. Consider a mail gateway whose function is to route incoming user email based on characteristics of the email. For example, emails from “boss” that are marked “urgent” are routed to the user’s cell phone as are all emails from “spouse.” All other emails are routed to the user’s desktop. When the emails are transmitted in the clear the gateway’s job is straight forward. However, when the emails are encrypted with the user’s public key the gateway cannot see data needed for the routing decision. The simplest solution is to give the gateway the user’s secret key, but this enables the gateway to decrypt all emails and exposes more information than the gateway needs.

A better solution is to encrypt emails using predicate encryption. The email header functions as the index x and the the routing instructions are used as m . The gateway is given a secret key sk_p corresponding to the “route to cell phone” predicate. This secret key enables the gateway to learn the routing instructions for messages satisfying the predicate p , but learn nothing else about emails.

Function privacy. A limitation of many existing predicate encryption systems is that the secret key sk_p reveals information about the predicate p . As a result, the gateway, and anyone else who has access to sk_p , learns the predicate p . Since in many practical settings it is important to keep the predicate p secret, our goal is to provide *function privacy*: sk_p should reveal as little information about p as possible.

At first glance it seems that hiding p is impossible: given sk_p the gateway can itself encrypt messages (x, m) and then apply sk_p to the resulting ciphertext. In doing so the gateway learns if $p(x) = 1$ which reveals some information about p . Nevertheless, despite this inherent limitation, function privacy can still be achieved.

Towards a solution. In recent work Boneh, Raghunathan, and Segev [10] put forward a new notion of function privacy and applied it to identity-based encryption systems (i.e. to predicate encryption supporting equality predicates). They observe that if the identity id is chosen from a distribution with super-logarithmic min-entropy then the inherent limitation above is not a problem since the attacker cannot learn id from sk_{id} by a brute force search since there are too many potential identities to test. They define function privacy for IBE systems by requiring that when id has sufficient min-entropy then sk_{id} is indistinguishable from a secret key derived for an independently and uniformly distributed identity. This enables function private keyword searching on encrypted data. They then construct several IBE systems supporting function-private keyword searching.

While Boneh et al. [10] showed how to achieve function privacy for equality predicates, encrypted search typically requires a richer set of searching predicates, including conjunctions, disjunctions, and many others. The authors left open the important question of achieving function privacy for a larger family of predicates.

Our contributions. In this paper we extend the framework and techniques of Boneh et al. [10] for constructing function-private encryption schemes. We put forward a generalization of inner-product predicate encryption [23, 18, 3], which we denote subspace-membership encryption, and present a definitional framework for capturing its function privacy. Our framework identifies the minimal restrictions under which a strong and meaningful notion of function privacy can be obtained for subspace-membership encryption schemes.

Then, we present a generic construction of a *function-private* subspace membership encryption scheme based on any underlying inner-product encryption scheme (even when the underlying scheme is *not* function private). Our construction is efficient, and in addition to providing function privacy, it preserves the security properties of the underlying scheme. This is the first generic construction that yields a function-private encryption scheme based on a non-function-private one. Recall that even for the simpler case of identity-based encryption, Boneh et al. [10] were not able to provide a generic construction, and had to individually modify various existing schemes.

Finally, we present various applications of function-private subspace membership encryption (we refer the reader to Section 1.1 for an overview of these applications). Among our applications, we significantly improve the function privacy of the identity-based encryption schemes of Boneh et al. [10]. Specifically, whereas their schemes guarantee function privacy only for identity distributions that are highly unpredictable (with min-entropy of at least $\lambda + \omega(\log \lambda)$ bits, where λ is the security parameter), we construct schemes that guarantee function privacy assuming only *minimal* unpredictability (i.e., min-entropy of $\omega(\log \lambda)$ bits). This improvement presents a much more realistic function privacy guarantee.

1.1 Overview of Our Contributions

A subspace-membership encryption scheme is a predicate encryption scheme supporting subspace-membership predicates. That is, an encryption of a message is associated with an attribute $\mathbf{x} \in \mathbb{S}^\ell$, and secret keys are derived for subspaces defined by all vectors in \mathbb{S}^ℓ orthogonal to a matrix $\mathbf{W} \in \mathbb{S}^{m \times \ell}$ (for integers $m, \ell \in \mathbb{N}$ and an additive group \mathbb{S}).³ Decryption recovers the message iff $\mathbf{W} \cdot \mathbf{x} = \mathbf{0}$. We refer the reader to [11, Section 2.3] for the standard definitions of the functionality and data security of predicate encryption (following [23, 3]).

Function privacy for subspace-membership encryption. Our goal is to design subspace-membership encryption schemes in which a secret key, $\text{sk}_{\mathbf{W}}$, does not reveal any information, beyond the absolute minimum necessary, on the matrix \mathbf{W} . Formalizing a realistic notion of function privacy, however, is not straightforward due to the actual functionality of subspace-membership encryption. Specifically, assuming that an adversary who is given a secret key $\text{sk}_{\mathbf{W}}$ has some a-priori information that the matrix \mathbf{W} belongs to a small set of matrices (e.g., $\{\mathbf{W}_0, \mathbf{W}_1\}$), then the adversary may be able to fully recover \mathbf{W} : The adversary simply needs to encrypt a (possibly random) message m for some attribute \mathbf{x} that is orthogonal to \mathbf{W}_0 but not to \mathbf{W}_1 , and then run the decryption algorithm on the given secret key $\text{sk}_{\mathbf{W}}$ and the resulting ciphertext to identify the one that decrypts correctly. In fact, as in [10], as long as the adversary has some a-priori information according to which the matrix \mathbf{W} is sampled from a distribution whose min-entropy is at most logarithmic in the security parameter, there is a non-negligible probability for a full recovery.

In the setting of subspace-membership encryption (unlike that of identity-based encryption [10]), however, the requirement that \mathbf{W} is sampled from a source of high min-entropy does not suffice for obtaining a meaningful notion of function privacy. In Section 2 we show that even if \mathbf{W} has nearly full min-entropy, but two of its columns may be correlated, then a meaningful notion of function privacy is not within reach.

In this light, our notion of function privacy for subspace-encryption schemes focuses on secret key $\text{sk}_{\mathbf{W}}$ for which the columns of \mathbf{W} form a block source. That is, each column of \mathbf{W} should have a reasonable amount of min-entropy even given all previous columns. Our notion of function privacy requires that such a secret key $\text{sk}_{\mathbf{W}}$ (where \mathbf{W} is sampled from an *adversarially*-chosen distribution) be indistinguishable from a secret key for a subspace chosen uniformly at random.

A function-private construction from inner-product encryption. Given any underlying inner-product encryption scheme we construction a function-private subspace-membership encryption scheme quite naturally. We modify the key-generation algorithm as follows: for generating a secret key for a subspace described by \mathbf{W} , we first sample a uniform $\mathbf{s} \leftarrow \mathbb{S}^m$ and use the key-generation algorithm of the underlying scheme for generating a secret key for the vector $\mathbf{v} = \mathbf{W}^\top \mathbf{s}$. Observe that as long as the columns of \mathbf{W} form a block source, then

³ Note that by setting $m = 1$ one obtains the notion of an inner-product encryption scheme [23, 18, 3].

the leftover hash lemma for block sources guarantees that \mathbf{v} is statistically close to uniform. In particular, essentially no information on \mathbf{W} is revealed.

We also observe that extracting from the columns of \mathbf{W} using the same seed for the extractor $\langle \mathbf{s}, \cdot \rangle$ interacts nicely with the subspace-membership functionality. Indeed, if $\mathbf{W} \cdot \mathbf{x} = \mathbf{0}$, it holds that $\mathbf{v}^\top \mathbf{x} = 0$ and vice-versa with high probability. We note that the case where the attribute set is small requires some additional refinement that we omit from this overview, and we refer the reader to Section 3 for more details.

Application 1: Function privacy when encrypting to roots of polynomials. We consider predicate encryption schemes supporting polynomial evaluation where secret keys correspond to polynomials $p \in \mathbb{S}[\mathbb{X}]$ and messages are encrypted to an attribute $x \in \mathbb{S}$. Given a secret key sk_p and a ciphertext with an attribute x , decryption recovers the message iff $p(x)$ evaluates to 0. Our work constructs such schemes from any underlying subspace-membership scheme.

We also explore the notion of function privacy for such polynomial encryption schemes. We require that secret keys for degree- d polynomials $p(x)$ with coefficients $(p_0, p_1, \dots, p_d) \in \mathbb{S}^{d+1}$ coming from a sufficiently unpredictable adversarially chosen (joint) distribution be indistinguishable from secret keys for degree- d polynomials where each coefficient is sampled uniformly from the underlying set. Unlike the case of subspace membership, we do not restrict our security to those distributions whose unpredictability holds even when conditioned on all previous (i.e., here we obtain security for any min-entropy source and not only for block sources).

Our function-private construction maps attributes x to Vandermonde vectors $\mathbf{x} = (1, x, x^2, \dots)$ and a polynomial $p(x)$ to a subspace \mathbf{W} as follows. We sample $d + 1$ polynomials $r_1(x), \dots, r_{d+1}(x)$ in a particular manner (as a product of d uniformly random linear polynomials) and construct the subspace \mathbf{W} whose i^{th} row comprises the coefficients of $p(x) \cdot r_i(x)$. In section 4.1, we elaborate on the details and prove that our choice of randomizing polynomials allows us to show that for polynomials whose coefficients come from an unpredictable distribution, \mathbf{W} 's columns have conditional unpredictability. And similarly, for polynomials with uniformly distributed coefficients, \mathbf{W} 's columns are uniformly distributed. This allows us to infer the function privacy of the polynomial encryption scheme from the function privacy of the underlying subspace-membership encryption scheme.

Application 2: Function-private IBE with minimal unpredictability.

As another interesting application of predicate encryption supporting polynomial evaluation, we consider the question of constructing function-private IBE schemes whose function privacy requires only the minimal necessary unpredictability assumption. It is easy to see (and as was shown in [10]) that if the adversary has some a-priori information according to which identities are sampled from a distribution with only logarithmic bits of entropy, then a simple adversary recovers id from sk_{id} with non-negligible probability by simply encrypting a message to a guessed id and checking if decryption recovers the message successfully.

Their constructions use a technique of preprocessing the id with a randomness extractor to recover id_{Ext} that is statistically close to uniform and thus hides any information about the underlying distribution of identities. As the extracted identity is roughly λ bits long, the distribution of identities must have min-entropy at least $\lambda + \omega(\log \lambda)$ bits to guarantee that extraction works. The identity space is much larger and this is still a meaningful notion of function privacy but the question of designing schemes that require the minimal min-entropy of $\omega(\log \lambda)$ bits was left open.

Starting from encryption schemes supporting polynomial evaluation (for our construction, linear polynomials suffice), this work shows how to construct function-private IBE schemes with the only restriction on identities being that they are unpredictable. We consider identities in a set \mathbb{S} and consider a polynomial $p_{\text{id}}(x) = (x - \text{id})$. By first randomizing the polynomial with uniformly chosen r in \mathbb{S} , we observe that if id has the minimal super-logarithmic unpredictability, then the coefficients of the polynomial $r \cdot (x - \text{id})$ have sufficient unpredictability. Thus, considering polynomial encryption schemes where secret keys correspond to such polynomials and attributes correspond to $x = \text{id}$, we construct IBE schemes that are function private against distributions that only have the minimum necessary unpredictability.

1.2 Related Work

As discussed above, the notion of function privacy was recently put forward by Boneh, Raghunathan, and Segev [10]. One of the main motivations of Boneh et al. was that of designing public-key searchable encryption schemes [8, 20, 1, 13, 28, 23, 5, 14, 2, 3] that are keyword private. That is, public-key searchable encryption schemes in which search tokens hide, as much as possible, their corresponding predicates. They presented a framework for modeling function privacy, and constructed various function-private anonymous identity-based encryption schemes (which, in particular, imply public-key keyword-private searchable encryption schemes).

More generally, the work of Boneh et al. initiated the study of function privacy in functional encryption [12, 26, 6, 21, 4, 19], where a functional secret key sk_f corresponding to a function f enables to compute $f(m)$ given an encryption $c = \text{Enc}_{\text{pk}}(m)$. Intuitively, in this setting function privacy guarantees that a functional secret key sk_f does not reveal information about f beyond what is already known and what can be obtained by running the decryption algorithm on test ciphertexts. In [10], the authors also discuss connections of function privacy to program obfuscation.

Our notion of subspace-membership encryption generalizes that of inner-product encryption introduced by Katz, Sahai, and Waters [23]. They defined and constructed predicate encryption schemes for predicates corresponding to inner products over \mathbb{Z}_N (for some large N). Informally, this class of predicates corresponds to functions $f_{\mathbf{v}}$ where $f_{\mathbf{v}}(\mathbf{x}) = 1$ if and only if $\langle \mathbf{v}, \mathbf{x} \rangle = 0$. Subsequently, Freeman [18] modified their construction to inner products over groups of prime order p , and Agrawal, Freeman, and Vaikuntanathan [3] constructed an

inner-product encryption scheme over \mathbb{Z}_p for a small prime p . Other results on inner product encryption study adaptive security [25], delegation in the context of hierarchies [24], and generalized IBE [9].

Finally, we note that function privacy in the symmetric-key setting, where the encryptor and decryptor have a shared secret key, was studied by Shen, Shi, and Waters [27]. They designed a function-private inner-product encryption scheme. As noted by Boneh et al. [10], achieving function privacy in the public-key setting is a more subtle task due to the inherent conflict between privacy and functionality.

1.3 Notation

For an integer $n \in \mathbb{N}$ we denote by $[n]$ the set $\{1, \dots, n\}$, and by U_n the uniform distribution over the set $\{0, 1\}^n$. For a random variable X we denote by $x \leftarrow X$ the process of sampling a value x according to the distribution of X . Similarly, for a finite set S we denote by $x \leftarrow S$ the process of sampling a value x according to the uniform distribution over S . We denote by \mathbf{x} (and sometimes \mathbf{x}) a vector $(x_1, \dots, x_{|\mathbf{x}|})$. We denote by $\mathbf{X} = (X_1, \dots, X_T)$ a joint distribution of T random variables. A non-negative function $f : \mathbb{N} \rightarrow \mathbb{R}$ is *negligible* if it vanishes faster than any inverse polynomial. A non-negative function $f : \mathbb{N} \rightarrow \mathbb{R}$ is *super-polynomial* if it grows faster than any polynomial.

The *min-entropy* of a random variable X is $\mathbf{H}_\infty(X) = -\log(\max_x \Pr[X = x])$. A *k-source* is a random variable X with $\mathbf{H}_\infty(X) \geq k$. A (T, k) -*block source* is a random variable $\mathbf{X} = (X_1, \dots, X_T)$ where for every $i \in [T]$ and x_1, \dots, x_{i-1} it holds that $\mathbf{H}_\infty(X_i | X_1 = x_1, \dots, X_{i-1} = x_{i-1}) \geq k$. The *statistical distance* between two random variables X and Y over a finite domain Ω is $\mathbf{SD}(X, Y) = \frac{1}{2} \sum_{\omega \in \Omega} |\Pr[X = \omega] - \Pr[Y = \omega]|$. Two random variables X and Y are *δ -close* if $\mathbf{SD}(X, Y) \leq \delta$. Two distribution ensembles $\{X_\lambda\}_{\lambda \in \mathbb{N}}$ and $\{Y_\lambda\}_{\lambda \in \mathbb{N}}$ are *statistically indistinguishable* if it holds that $\mathbf{SD}(X_\lambda, Y_\lambda)$ is negligible in λ . They are *computationally indistinguishable* if for every probabilistic polynomial-time algorithm \mathcal{A} it holds that $|\Pr[\mathcal{A}(1^\lambda, x) = 1] - \Pr[\mathcal{A}(1^\lambda, y) = 1]|$ is negligible in λ , where $x \leftarrow X_\lambda$ and $y \leftarrow Y_\lambda$.

1.4 Paper Organization

The remainder of this paper is organized as follows. Due to space constraints, we refer the reader to the full version [11, Section 2] for standard definitions and tools. In Section 2 we introduce the notions of subspace-membership encryption and function privacy for subspace-membership encryption. In Section 3 we present a generic construction of a function-private subspace-membership encryption scheme based on any inner-product encryption scheme. In Section 4 we present various applications of function-private subspace-membership encryption. In Section 5 we discuss several open problems that arise from this work.

2 Subspace-Membership Encryption and Its Function Privacy

In this section we formalize the notion of subspace-membership encryption and its function privacy within the framework of Boneh, Raghunathan and Segev [10]. A subspace-membership encryption scheme is a predicate encryption scheme [13, 23] supporting the class of predicates \mathcal{F} , over an attribute space $\Sigma = \mathbb{S}^\ell$, defined as

$$\mathcal{F} = \{f_{\mathbf{W}} : \mathbf{W} \in \mathbb{S}^{m \times \ell}\} \quad \text{with} \quad f_{\mathbf{W}}(\mathbf{x}) = \begin{cases} 1 & \mathbf{W} \cdot \mathbf{x} = \mathbf{0} \in \mathbb{S}^m \\ 0 & \text{otherwise} \end{cases}$$

for integers $m, \ell \in \mathbb{N}$, and an additive group \mathbb{S} . Informally, in a subspace-membership encryption, an encryption of a message is associated with an attribute $\mathbf{x} \in \mathbb{S}^\ell$, and secret keys are derived for subspaces defined by all vectors in \mathbb{S}^ℓ orthogonal to a matrix $\mathbf{W} \in \mathbb{S}^{m \times \ell}$. Decryption recovers the message if and only if $\mathbf{W} \cdot \mathbf{x} = \mathbf{0}$. (See [11, Section 2.3] for the standard definitions of the functionality and data security of predicate encryption.) Subspace-membership encryption with delegation was also studied in [24, 25]. Here we do not need the delegation property.

Based on the framework introduced by Boneh, Raghunathan, and Segev [10], our notion of function privacy for subspace-membership encryption considers adversaries that are given the public parameters of the scheme and can interact with a “real-or-random” function-privacy oracle RoR^{FP} defined as follows, and with a key-generation oracle.

Definition 2.1 (Real-or-random function-privacy oracle). *The real-or-random function-privacy oracle RoR^{FP} takes as input triplets of the form $(\text{mode}, \text{msk}, V)$, where $\text{mode} \in \{\text{real}, \text{rand}\}$, msk is a master secret key, and $V = (V_1, \dots, V_\ell) \in \mathbb{S}^{m \times \ell}$ is a circuit representing a joint distribution over $\mathbb{S}^{m \times \ell}$ (i.e., each V_i is a distribution over \mathbb{S}^m). If $\text{mode} = \text{real}$ then the oracle samples $\mathbf{W} \leftarrow V$ and if $\text{mode} = \text{rand}$ then the oracle samples $\mathbf{W} \leftarrow \mathbb{S}^{m \times \ell}$ uniformly. It then invokes the algorithm $\text{KeyGen}(\text{msk}, \cdot)$ on \mathbf{W} for outputting a secret key $\text{sk}_{\mathbf{W}}$.*

Definition 2.2 (Function-privacy adversary). *An (ℓ, k) -block-source function-privacy adversary \mathcal{A} is an algorithm that is given as input a pair $(1^\lambda, \text{pp})$ and oracle access to $\text{RoR}^{\text{FP}}(\text{mode}, \text{msk}, \cdot)$ for some $\text{mode} \in \{\text{real}, \text{rand}\}$, and to $\text{KeyGen}(\text{msk}, \cdot)$. It is required that each of \mathcal{A} 's queries to RoR^{FP} be an (ℓ, k) -block-source.*

Definition 2.3 (Function-private subspace-membership encryption). *A subspace-membership encryption scheme $\Pi = (\text{Setup}, \text{KeyGen}, \text{Enc}, \text{Dec})$ is (ℓ, k) -block-source function private if for any probabilistic polynomial-time (ℓ, k) -block-source function-privacy adversary \mathcal{A} , there exists a negligible function $\nu(\lambda)$ such that*

$$\text{Adv}_{\Pi, \mathcal{A}}^{\text{FP}}(\lambda) \stackrel{\text{def}}{=} \left| \Pr \left[\text{Expt}_{\text{FP}, \Pi, \mathcal{A}}^{\text{real}}(\lambda) = 1 \right] - \Pr \left[\text{Expt}_{\text{FP}, \Pi, \mathcal{A}}^{\text{rand}}(\lambda) = 1 \right] \right| \leq \nu(\lambda),$$

where for each $\text{mode} \in \{\text{real}, \text{rand}\}$ and $\lambda \in \mathbb{N}$ the experiment $\text{Expt}_{\text{FP}, \Pi, \mathcal{A}}^{\text{mode}}(\lambda)$ is defined as follows:

1. $(\text{pp}, \text{msk}) \leftarrow \text{Setup}(1^\lambda)$.
2. $b \leftarrow \mathcal{A}^{\text{RoR}^{\text{FP}}(\text{mode}, \text{msk}, \cdot), \text{KeyGen}(\text{msk}, \cdot)}(1^\lambda, \text{pp})$.
3. Output b .

In addition, such a scheme is statistically (ℓ, k) -block-source function private if the above holds for all computationally-unbounded (ℓ, k) -block-source function-privacy adversary making a polynomial number of queries to the RoR^{FP} oracle.

We note here that a security model that allows the adversary to receive the master secret key msk in place of the oracle $\text{KeyGen}(\text{msk}, \cdot)$ leads to a seemingly stronger notion of function privacy. However, such a notion is subsumed by *statistical* function privacy and the schemes constructed in this paper actually satisfy this stronger notion.

Multi-shot vs. single-shot adversaries. Note that Definition 2.3 considers adversaries that query the function-privacy oracle for any polynomial number of times. In fact, as adversaries are also given access to the key-generation oracle, this “multi-shot” definition is polynomially equivalent to its “single-shot” variant in which adversaries query the real-or-random function-privacy oracle RoR^{FP} at most once. This is proved via a straightforward hybrid argument, where the hybrids are constructed such that only one query is forwarded to the function-privacy oracle, and all other queries are answered using the key-generation oracle.

The block-source requirement on the columns of \mathbf{W} . Our definition of function privacy for subspace-membership encryption requires that a secret key $\text{sk}_{\mathbf{W}}$ reveals no unnecessary information about \mathbf{W} as long as the columns of \mathbf{W} form a block source (i.e., each column is unpredictable even given the previous columns). One might consider a stronger definition, in which the columns of \mathbf{W} may be arbitrarily correlated, as long as each column of \mathbf{W} is sufficiently unpredictable. Such a definition, however, is impossible to satisfy.

Specifically, consider the special case of inner-product encryption (i.e., $m = 1$), and an adversary that queries the real-or-random oracle with a distribution over vectors $\mathbf{w} \in \mathbb{S}^\ell$ defined as follows: sample $\ell - 1$ independent and uniform values $u_1, \dots, u_{\ell-1} \leftarrow \mathbb{S}$ and output $\mathbf{w} = (u_1, 2u_1, u_2, \dots, u_{\ell-1})$. Such a distribution clearly has high min-entropy (specifically, $(\ell - 1) \log |\mathbb{S}|$ bits), and each coordinate of \mathbf{w} has min-entropy $\log |\mathbb{S}|$ bits. However, secret keys for vectors drawn from this distribution can be easily distinguished from secret keys for vectors drawn from the uniform distribution over \mathbb{S}^ℓ : encrypt a message \mathbf{M} to the attribute $\mathbf{x} = (-2, 1, 0, \dots, 0) \in \mathbb{S}^\ell$ and check to see if decryption succeeds in recovering \mathbf{M} . For a random vector $\mathbf{w} \in \mathbb{S}^\ell$ the decryption succeeds only with probability $1/|\mathbb{S}|$ giving the adversary an overwhelming advantage.

Therefore, restricting function privacy adversaries to query the RoR^{FP} oracle only with sources whose columns form block sources is essential for achieving a meaningful notion of function privacy.

On correlated RoR^{FP} queries. In Definition 2.2 we consider adversaries that receives only a single secret key $\text{sk}_{\mathbf{W}}$ for each query to the RoR^{FP} oracle. Our definition easily generalizes to include adversaries that are allowed to query

the RoR^{FP} oracle with *correlated* queries. More specifically, an adversary can receive secret keys $\text{sk}_{\mathbf{W}_1}, \dots, \text{sk}_{\mathbf{W}_T}$ for any parameter T that is polynomial in the security parameter. The RoR^{FP} oracle samples subspaces $\mathbf{W}_1, \dots, \mathbf{W}_T$ from an adversarially chosen joint distribution over $(\mathbb{S}^{m \times \ell})^T$ with the restriction that for every $1 \leq i \leq T$, the columns of \mathbf{W}_i come from a (ℓ, k) -block-source even conditioned on any fixed values for $\mathbf{W}_1, \dots, \mathbf{W}_{i-1}$.⁴

Function privacy of existing inner-product encryption schemes. The inner-product predicate encryption scheme from lattices [3] is trivially not function private as the secret key includes the corresponding function $f_{\mathbf{v}}$ as part of it (this is necessary for the decryption algorithm to work correctly). The scheme constructed from bilinear groups with composite order [23] however presents no such obvious attack, but we were not able to prove its function privacy based on any standard cryptographic assumption.

3 A Generic Construction Based on Inner-Product Encryption

In this section we present a generic construction of a function-private subspace-membership encryption scheme starting from any inner-product encryption scheme. Due to space constraints, we deal with a large attribute space \mathbb{S} of size super-polynomial in the security parameter λ here, and explain our idea of extending our construction to the case when $|\mathbb{S}|$ is small (see [11, Section 4.2] for the details).

Our construction. Let $\mathcal{IP} = (\text{IP.Setup}, \text{IP.KeyGen}, \text{IP.Enc}, \text{IP.Dec})$ be an inner-product encryption scheme with attribute set $\Sigma = \mathbb{S}^\ell$. We construct a subspace-membership encryption scheme $\mathcal{SM} = (\text{SM.Setup}, \text{SM.KeyGen}, \text{SM.Enc}, \text{SM.Dec})$ as follows.

- **Setup:** SM.Setup is identical to IP.Setup . On input the security parameter it outputs public parameters pp and the master secret key msk by running $\text{IP.Setup}(1^\lambda)$.
- **Key generation:** SM.KeyGen takes as input the master secret key msk and a function $f_{\mathbf{W}}$ where $\mathbf{W} \in \mathbb{S}^{m \times \ell}$ and proceeds as follows. It samples uniform $\mathbf{s} \leftarrow \mathbb{S}^m$ and computes $\mathbf{v} = \mathbf{W}^\top \mathbf{s} \in \mathbb{S}^\ell$. Next, it samples a secret key $\text{sk}_{\mathbf{v}} \leftarrow \text{IP.KeyGen}(\text{msk}, \mathbf{v})$ and outputs $\text{sk}_{\mathbf{W}} \stackrel{\text{def}}{=} \text{sk}_{\mathbf{v}}$.
- **Encryption:** SM.Enc is identical to IP.Enc . On input the public parameters, an attribute $\mathbf{x} \in \mathbb{S}^\ell$, and a message M , the algorithm outputs a ciphertext $c \leftarrow \text{IP.Enc}(\text{pp}, \mathbf{x}, M)$.
- **Decryption:** SM.Dec is identical to IP.Dec . On input the public parameters pp , a secret key $\text{sk}_{\mathbf{W}}$, and a ciphertext c , it outputs $M \leftarrow \text{IP.Dec}(\text{pp}, \text{sk}_{\mathbf{W}}, c)$.

⁴ Or equivalently, the columns of $[\mathbf{W}_1 \mid \mathbf{W}_2 \mid \dots \mid \mathbf{W}_T]$ are distributed according to a $(T\ell, k)$ -block-source.

Correctness. Correctness of the construction follows from the correctness of the underlying inner-product encryption scheme. For every $\mathbf{W} \in \mathbb{S}^{m \times \ell}$ and every $\mathbf{x} \in \mathbb{S}^\ell$, it suffices to show the following:

- If $f(I) = 1$, then it holds that $\mathbf{W} \cdot \mathbf{x} = \mathbf{0}$. This implies $\mathbf{x}^\top \mathbf{v} = \mathbf{x}^\top (\mathbf{W}^\top \mathbf{s}) = 0$ and therefore SM.Dec correctly outputs M as required.
- If $f(I) = 0$, then it holds that $\mathbf{e} \stackrel{\text{def}}{=} \mathbf{W} \cdot \mathbf{x} \neq \mathbf{0} \in \mathbb{S}^m$. As $\mathbf{x}^\top \mathbf{v} = \mathbf{x}^\top (\mathbf{W}^\top \mathbf{s}) = \mathbf{e}^\top \mathbf{s}$, for any $\mathbf{e} \neq \mathbf{0}$ the quantity $\mathbf{x}^\top \mathbf{v}$ is zero with probability $1/|\mathbb{S}|$ over choices of \mathbf{s} . As $1/|\mathbb{S}|$ is negligible in λ whenever $|\mathbb{S}|$ is super-polynomial in λ , the proof of correctness follows.

Security. We state the following theorem about the security of our construction.

Theorem 3.1. *If \mathcal{IP} is an attribute hiding (resp. weakly attribute hiding) inner-product encryption scheme for an attribute set \mathbb{S} of size super-polynomial in the security parameter, then it holds that:*

1. *The scheme \mathcal{SM} is an attribute hiding (resp. weakly attribute hiding) subspace-membership encryption scheme under the same assumption as the security of the underlying inner-product encryption scheme.*
2. *The scheme \mathcal{SM} when $m \geq 2$ is statistically function private for (ℓ, k) -block-sources for any $\ell = \text{poly}(\lambda)$ and $k \geq \log |\mathbb{S}| + \omega(\log \lambda)$.*

Proof. We first prove the attribute-hiding property of the scheme, and then prove its function privacy.

Attribute hiding. Attribute-hiding property of \mathcal{SM} follows from the attribute-hiding property of \mathcal{IP} in a rather straightforward manner. Given a challenger for the attribute-hiding property of \mathcal{IP} , an \mathcal{SM} adversary \mathcal{A} can be simulated by algorithm \mathcal{B} as follows: \mathcal{A} 's challenge attributes are forwarded to the \mathcal{IP} -challenger and the resulting public parameters are published. Secret key queries can be simulated by first sampling uniform $\mathbf{s} \leftarrow \mathbb{S}^m$, then computing $\mathbf{v} = \mathbf{W}^\top \mathbf{s}$ and forwarding \mathbf{v} to the \mathcal{IP} key generation oracle. Similarly, the challenge messages from the adversary are answered by forwarding them to the challenger. In the full version [11, Section 4.1], we elaborate on the details and show that if Q denotes the number of secret key queries by \mathcal{A} , it holds that

$$\mathbf{Adv}_{\mathcal{IP}, \mathcal{B}}(\lambda) \geq \mathbf{Adv}_{\mathcal{SM}, \mathcal{A}}(\lambda) - 2Q/|\mathbb{S}|, \quad (1)$$

thus completing the proof of the attribute hiding property of \mathcal{SM} .

Function privacy. Let \mathcal{A} be a computationally unbounded (ℓ, k) -block-source function-privacy adversary that makes a polynomial number $Q = Q(\lambda)$ of queries to the RoR^{FP} oracle. We prove that the distribution of \mathcal{A} 's view in the experiment $\text{Expt}_{\text{FP}, \mathcal{SM}, \mathcal{A}}^{\text{real}}$ is statistically close to the distribution of \mathcal{A} 's view in the experiment $\text{Expt}_{\text{FP}, \mathcal{SM}, \mathcal{A}}^{\text{rand}}$ (we refer the reader to Definition 2.3 for the descriptions of these experiments). We denote these two distributions by $\text{View}_{\text{real}}$ and $\text{View}_{\text{rand}}$, respectively.

As the adversary \mathcal{A} is computationally unbounded, we assume without loss of generality that \mathcal{A} does not query the $\text{KeyGen}(\text{msk}, \cdot)$ oracle—such queries can be internally simulated by \mathcal{A} . Moreover, as discussed in Section 2, it suffices to focus on adversaries \mathcal{A} that query the RoR^{FP} oracle exactly once. From this point on we fix the public parameters pp chosen by the setup algorithm, and show that the two distributions $\text{View}_{\text{real}}$ and $\text{View}_{\text{rand}}$ are statistically close for any such pp .

Denote by $V = (V_1, \dots, V_\ell)$ the random variable corresponding to the (ℓ, k) -source with which \mathcal{A} queries the RoR^{FP} oracle. For each $i \in [\ell]$, let $(w_{i,1}, \dots, w_{i,m})$ denote a sample from V_i . Also, let $\mathbf{s} = (s_1, \dots, s_m) \in \mathbb{S}^m$. As \mathcal{A} is computationally unbounded, and having fixed the public parameters, we can in fact assume that

$$\text{View}_{\text{mode}} = \left(\left(\sum_{i=1}^m s_i \cdot w_{i,1} \right), \dots, \left(\sum_{i=1}^m s_i \cdot w_{i,\ell} \right) \right) \quad (2)$$

for $\text{mode} \in \{\text{real}, \text{rand}\}$, where $\mathbf{W} = \{w_{i,j}\}_{i \in [m], j \in [\ell]}$ is drawn from V for $\text{mode} = \text{real}$, \mathbf{W} is uniformly distributed over $\mathbb{S}^{m \times \ell}$ for $\text{mode} = \text{rand}$, and $s_i \leftarrow \mathbb{S}$ for every $i \in [m]$. For $\text{mode} \in \{\text{real}, \text{rand}\}$ we prove that the distribution $\text{View}_{\text{mode}}$ is statistically close to a uniform distribution over \mathbb{S}^m .

Note that the collection of functions $\{g_{s_1, \dots, s_m} : \mathbb{S}^m \rightarrow \mathbb{S}\}_{s_1, \dots, s_m \in \mathbb{S}}$ defined by $g_{s_1, \dots, s_m}(w_1, \dots, w_m) = \sum_{j=1}^m s_j \cdot w_j$ is universal. This enables us to directly apply the Leftover Hash Lemma for block-sources [16, 22, 29, 17] implying that for our choice of parameters m, ℓ and k the statistical distance between $\text{View}_{\text{real}}$ and the uniform distribution is negligible in λ .⁵ The same clearly holds also for $\text{View}_{\text{rand}}$, as the uniform distribution over $\mathbb{S}^{m \times \ell}$ is, in particular, a (ℓ, k) -block-source. This completes the proof of function privacy. \blacksquare

Theorem 3.1 for correlated RoR^{FP} queries. Recollect that the definition of function privacy for subspace membership (Definition 2.3) extends to adversaries that query the RoR^{FP} oracle with secret keys for T correlated subspaces $\mathbf{W}_1, \dots, \mathbf{W}_T$ for any $T = \text{poly}(\lambda)$. If the columns of the jointly sampled subspaces $[\mathbf{W}_1 | \mathbf{W}_2 | \dots | \mathbf{W}_T]$ form a block source, we can extend the proof of function privacy to consider such correlated queries. The adversary's view comprises T terms as in Equation (2) with randomly sampled vectors $\mathbf{s}_1, \dots, \mathbf{s}_T$ in place of \mathbf{s} . The collection of functions g remains universal and a simple variant of the Leftover Hash Lemma implies that for our choice of parameters, the statistical distance between $\text{View}_{\text{real}}$ and the uniform distribution is negligible in λ (and similarly for $\text{View}_{\text{rand}}$).

Dealing with small attribute spaces. We also consider constructing subspace-membership encryption schemes where we do not place any restrictions on the size of the underlying attribute space \mathbb{S} . In our generic construction, observe that correctness requires that $1/|\mathbb{S}|$ be negligible in λ . If $|\mathbb{S}|$ is not super-polynomial in

⁵ We note here that a weaker version of the Leftover Hash Lemma will suffice as the adversary's view does not include (s_1, \dots, s_m) .

the security parameter, then correctness fails with a non-negligible probability. Additionally, this breaks the proof of attribute-hiding security in Theorem 3.1: In Equation (1), if the quantity $2Q/|\mathbb{S}|$ is non-negligible, then a non-negligible advantage of an adversary \mathcal{A} *does not* translate to a non-negligible advantage for the reduction algorithm \mathcal{B} against the inner-product encryption scheme.

To overcome this difficulty, we refine the construction as follows using a parameter $\tau = \tau(\lambda) \in \mathbb{N}$. We split the message into τ secret shares and apply parallel repetition of τ copies of the underlying inner-product encryption scheme, where each copy uses independent public parameters and master secret keys. For the proof of security, it suffices to have τ such that the quantity $\tau/|\mathbb{S}|^\tau$ is negligible in λ . Due to space constraints, a formal description of the scheme and a statement of its security is deferred to [11, Section 4.2].

4 Applications of Function-Private Subspace-Membership Encryption

4.1 Roots of a Polynomial Equation

We can construct a predicate encryption scheme for predicates corresponding to polynomial evaluation. Let $\Phi_{<d}^{\text{poly}} \stackrel{\text{def}}{=} \{f_p : p \in \mathbb{S}[X], \deg(p) < d\}$, where

$$f_p(x) = \begin{cases} 1 & \text{if } p(x) = 0 \in \mathbb{S} \\ 0 & \text{otherwise} \end{cases} \quad \text{for } x \in \mathbb{S}.$$

Correctness and attribute hiding properties of the predicate encryption scheme for the class of predicates $\Phi_{<d}^{\text{poly}}$ are defined as in the case of a generic predicate encryption scheme in a natural manner (see [11, Definition 2.3]).

Function-private polynomial encryption. For the class $\Phi_{<d}^{\text{poly}}$, consider a real-or-random function privacy oracle $\text{RoR}^{\text{FP-}\Phi}$ (along the lines of Definition 2.1) that takes as input triplets of the form $(\text{mode}, \text{msk}, \mathbf{P})$, where $\text{mode} \in \{\text{real}, \text{rand}\}$, msk is a master secret key, and $\mathbf{P} = (P_0, \dots, P_{d-1}) \in \mathbb{S}^d$ is a circuit representing a joint distribution over coefficients of polynomials p with $\deg(p) < d$. If $\text{mode} = \text{real}$ then the oracle samples $p \leftarrow \mathbf{P}$ and if $\text{mode} = \text{rand}$ then the oracle samples $p \leftarrow \mathbb{S}^d$ uniformly. It then invokes the algorithm $\text{KeyGen}(\text{msk}, \cdot)$ on p and outputs secret key sk_p .

Along the lines of Definition 2.2, we consider a k -source $\Phi_{<d}^{\text{poly}}$ function-privacy adversary \mathcal{A} . Such an adversary is given inputs $(1^\lambda, \text{pp})$ and oracle access to $\text{RoR}^{\text{FP-}\Phi}$ and each query to the oracle is a k -source (over the coefficients of the polynomial).

Definition 4.1 ($\Phi_{<d}^{\text{poly}}$ **Function privacy**). *A predicate encryption scheme for the class of predicates $\Phi_{<d}^{\text{poly}}$ denoted $\Pi = (\text{Setup}, \text{KeyGen}, \text{Enc}, \text{Dec})$ is k -source function-private if for any probabilistic polynomial-time k -source $\Phi_{<d}^{\text{poly}}$ function-privacy adversary \mathcal{A} , there exists a negligible function $\nu(\lambda)$ such that*

$$\text{Adv}_{\Pi, \mathcal{A}}^{\text{FP-}\Phi}(\lambda) \stackrel{\text{def}}{=} \left| \Pr \left[\text{Expt}_{\text{FP-}\Phi, \Pi, \mathcal{A}}^{\text{real}}(\lambda) = 1 \right] - \Pr \left[\text{Expt}_{\text{FP-}\Phi, \Pi, \mathcal{A}}^{\text{rand}}(\lambda) = 1 \right] \right| \leq \nu(\lambda),$$

where for each $\text{mode} \in \{\text{real}, \text{rand}\}$ and $\lambda \in \mathbb{N}$ the experiment $\text{Expt}_{\text{FP}, \Phi, \Pi, \mathcal{A}}^{\text{mode}}(\lambda)$ is defined as follows:

1. $(\text{pp}, \text{msk}) \leftarrow \text{Setup}(1^\lambda)$.
2. $b \leftarrow \mathcal{A}^{\text{RoR}^{\text{FP}, \Phi}(\text{mode}, \text{msk}, \cdot), \text{KeyGen}(\text{msk}, \cdot)}(1^\lambda, \text{pp})$.
3. Output b .

In addition, such a scheme is statistically k -source function private if the above holds for any computationally-unbounded k -source $\Phi_{<d}^{\text{poly}}$ function privacy adversary making a polynomial number of queries to the $\text{RoR}^{\text{FP}, \Phi}$ oracle.

Correlated $\text{RoR}^{\text{FP}, \Phi}$ queries. Definition 4.1 extends to adversaries that query the $\text{RoR}^{\text{FP}, \Phi}$ oracle on T correlated queries. A scheme Π is said to be (T, k) -source (resp. (T, k) -block-source) function private if each query $(\mathbf{P}_1, \dots, \mathbf{P}_T)$ of a joint distribution over T polynomials is a (T, k) -source (resp. (T, k) -block-source).

Constructing function-private predicate encryption schemes supporting polynomial evaluation. Given a subspace membership encryption scheme $(\text{Setup}, \text{KeyGen}, \text{Enc}, \text{Dec})$ with parameters $m = d$ and $\ell = 2d - 1$, we can construct a predicate encryption scheme for $\Phi_{<d}^{\text{poly}}$ as follows (for simplicity, we consider the instructive case $d = 3$ and subsequently explain how our technique generalizes):

- **Setup:** The **Setup** algorithm remains unchanged.
- **Encryption:** To encrypt a message M for the attribute $x \in \mathbb{S}$, the encryption algorithm sets $\mathbf{x} = (x^4, x^3, x^2, x, 1)^\top$ and outputs the ciphertext $\text{Enc}(\text{pp}, \mathbf{x}, M)$.
- **Key generation:** To generate a secret key corresponding to the polynomial $p = p_2 \cdot x^2 + p_1 \cdot x + p_0$, the key-generation algorithm constructs a vector $\mathbf{p} = (p_2, p_1, p_0)^\top \in \mathbb{S}^3$. Next, it “blinds” the polynomial $p(x)$ with two linear polynomials $r(x) = r_1 \cdot x + r_0$ and $s(x) = s_1 \cdot x + s_0$ and computes the coefficients of the polynomial $p(x) \cdot r(x) \cdot s(x)$. The coefficients r_1, r_0, s_1, s_0 are sampled independently and uniformly at random from \mathbb{S} . The key generation algorithm repeats this step with two more sets of polynomials (we refer to them as “randomizing” polynomials) $r'(x), s'(x)$ and $r''(x), s''(x)$ whose coefficients are also sampled uniformly at random. It constructs

$$\mathbf{W} = \begin{bmatrix} \text{--- coefficients of } p(x) \cdot r(x) \cdot s(x) \text{ ---} \\ \text{--- coefficients of } p(x) \cdot r'(x) \cdot s'(x) \text{ ---} \\ \text{--- coefficients of } p(x) \cdot r''(x) \cdot s''(x) \text{ ---} \end{bmatrix} \in \mathbb{S}^{3 \times 5}. \quad (3)$$

$$= \begin{bmatrix} p_2 r_1 s_1 & p_2 r_1 s_0 + p_2 r_0 s_1 & p_2 r_0 s_0 + p_1 r_1 s_0 & p_1 r_0 s_0 + p_0 r_0 s_1 & p_0 r_0 s_0 \\ & + p_1 r_1 s_1 & + p_1 r_0 s_1 + p_0 r_1 s_1 & + p_0 r_1 s_0 & \\ p_2 r'_1 s'_1 & p_2 r'_1 s'_0 + p_2 r'_0 s'_1 & p_2 r'_0 s'_0 + p_1 r'_1 s'_0 & p_1 r'_0 s'_0 + p_0 r'_0 s'_1 & p_0 r'_0 s'_0 \\ & + p_1 r'_1 s'_1 & + p_1 r'_0 s'_1 + p_0 r'_1 s'_1 & + p_0 r'_1 s'_0 & \\ p_2 r''_1 s''_1 & p_2 r''_1 s''_0 + p_2 r''_0 s''_1 & p_2 r''_0 s''_0 + p_1 r''_1 s''_0 & p_1 r''_0 s''_0 + p_0 r''_0 s''_1 & p_0 r''_0 s''_0 \\ & + p_1 r''_1 s''_1 & + p_1 r''_0 s''_1 + p_0 r''_1 s''_1 & + p_0 r''_1 s''_0 & \end{bmatrix}.$$

- The algorithm then runs $\text{KeyGen}(\text{msk}, \mathbf{W})$ and outputs $\text{sk}_{\mathbf{W}}$.
- **Decryption:** The decryption algorithm remains unchanged.

Correctness and attribute hiding. Given a ciphertext c for attribute x and a secret key for polynomial p , if $p(x) = 0$ then it follows that $\mathbf{W} \cdot \mathbf{x} = \mathbf{0}$. If $\mathbf{W} \cdot \mathbf{x} = \mathbf{0}$, then x is a root of polynomials $p \cdot r \cdot s$, $p \cdot r' \cdot s'$, and $p \cdot r'' \cdot s''$ which implies that x is a root of $p(x)$ with overwhelming probability over the choices of polynomials $r, r', r'', s, s', s'' \in \mathbb{S}[\mathbb{X}]$.⁶ The attribute hiding property of the scheme follows in a fairly straightforward manner from the attribute hiding property of the subspace membership encryption scheme.

Function privacy. We show that with overwhelming probability over the choices of the randomizing polynomials: (a) if the coefficients of p , namely (p_2, p_1, p_0) are sampled from a k -source, then \mathbf{W} is distributed according to a $(5, k)$ -block source, and (b) if the coefficients of p are sampled uniformly at random from \mathbb{S}^3 , then \mathbf{W} is distributed uniformly over $\mathbb{S}^{3 \times 5}$. Given the above two claims, a straightforward reduction allows us to simulate a $\text{RoR}^{\text{FP}-\Phi}$ oracle given access to a RoR oracle for the subspace membership predicate with parameters $m = 3$ and $\ell = 5$. Thus, we can state the following theorem.

Theorem 4.2. *If \mathcal{SM} is a subspace membership encryption scheme with parameters $m = 3$ and $\ell = 5$ that satisfies function privacy against $(5, k)$ -block-source adversaries, then the predicate encryption scheme for the class of predicates $\Phi_{<3}^{\text{poly}}$ constructed above is statistically function private against k -source adversaries.*

Applying Theorem 3.1 for adversaries that query the RoR^{FP} oracle with T correlated queries immediately gives us the following corollary.

Corollary 4.3. *Given any large attribute space inner-product encryption scheme with $\ell = 3$, there exists a predicate encryption scheme for the class of predicates $\Phi_{<3}^{\text{poly}}$ that is statistically function-private against (T, k) -block-sources for any $T = \text{poly}(\lambda)$ and $k \geq \log |\mathbb{S}| + \omega(\log \lambda)$.*

Proof of claims (a) and (b). Consider column $\mathbf{w}_1 = (p_2 r_1 s_1, p_2 r'_1 s'_1, p_2 r''_1 s''_1)^\top$. We observe that over choices of $s_1, s'_1,$ and s''_1 , the column \mathbf{w}_1 is distributed uniformly over \mathbb{S}^3 . The second column \mathbf{w}_2 is also distributed uniformly at random by noting that the elements $p_2 r_1 s_0, p_2 r'_1 s'_0,$ and $p_2 r''_1 s''_0$ are distributed uniformly in \mathbb{S}^3 over choices of $r_1, r'_1,$ and r''_1 (which are themselves information theoretically hidden in \mathbf{w}_1). An identical argument shows that over choices of $r_0, r'_0,$ and $r''_0,$ and $s_0, s'_0,$ and s''_0 , the fourth and fifth columns, \mathbf{w}_4 and \mathbf{w}_5 , are distributed uniformly in \mathbb{S}^3 . This is true even conditioned on all the other columns. It suffices to show that conditioned on $\mathbf{w}_1, \mathbf{w}_2, \mathbf{w}_4,$ and \mathbf{w}_5 , column \mathbf{w}_3 has entropy at least $\log |\mathbb{S}| + \omega(\log \lambda)$.

⁶ From a simple union bound over the events where three linear polynomials share a root, this probability works out to be $\geq 1 - 8/|\mathbb{S}|^2$ which is indeed overwhelming.

We re-write \mathbf{w}_3 as $\mathbf{R} \cdot \mathbf{p}$ where

$$\mathbf{R} = \begin{bmatrix} r_0 s_0 & r_1 s_0 + r_0 s_1 & r_1 s_1 \\ r'_0 s'_0 & r'_1 s'_0 + r'_0 s'_1 & r'_1 s'_1 \\ r''_0 s''_0 & r''_1 s''_0 + r''_0 s''_1 & r''_1 s''_1 \end{bmatrix} \in \mathbb{S}^{3 \times 3}. \quad (4)$$

With overwhelming probability over random choices of all the coefficients in the polynomials $r, s, r', s', r'',$ and s'' , the matrix \mathbf{R} is full-rank over \mathbb{S} . Therefore, the distribution of \mathbf{w}_3 has a one-one correspondence with the distribution of \mathbf{p} . Therefore, \mathbf{w}_3 has entropy at least k even given \mathbf{R} if p is sampled from a k -source and \mathbf{w}_3 is uniform over \mathbb{S}^3 even given \mathbf{R} if p is sampled uniformly from \mathbb{S}^3 . This concludes the proof of claims (a) and (b). \blacksquare

A general technique for $\Phi_{<d}^{\text{poly}}$. As stated earlier, we can construct predicate encryption for the class of predicates $\Phi_{<d}^{\text{poly}}$ starting with a subspace membership encryption scheme with parameters $m = d$ and $\ell = 2d - 1$. The main idea in extending beyond $d = 3$ is to construct d randomized “blindings” of $p(x)$. For $i \in [d]$, the i^{th} row of \mathbf{W} now comprises coefficients of a polynomial $p(x) \cdot r_{i,1}(x) \cdots r_{i,d-1}(x)$ where each of the $r_{i,j}(x)$ ’s are random linear polynomials sampled as $r(x)$ and $s(x)$ are sampled in the $d = 3$ construction. The details of our construction are as follows. Due to space constraints the details about the construction are deferred to the full version [11, Section 5.1].

Comparing entropy requirements. In Definition 4.1 and Corollary 4.3 it suffices to consider function-privacy adversaries that query the “real-or-random” oracle with polynomials whose coefficients come from a k -source. We *do not* require the sources have conditional min-entropy in contrast to subspace membership function privacy (see Definition 2.3 and the discussion in Section 2). The reason this weaker restriction on $\Phi_{<d}^{\text{poly}}$ function-privacy adversaries suffices when it does not suffice against subspace membership function-privacy adversaries is that the class of predicates $\Phi_{<d}^{\text{poly}}$ offers a weaker functionality than is offered by subspace membership. In particular, if the adversary evaluates ciphertexts with attributes corresponding to “ill-formed” non-Vandermonde vectors, i.e., vectors not of the form $(1, x, x^2, \dots)$, correctness of decryption is not guaranteed and the particular attack outlined in Section 2 fails. It is easy to see this in our construction as well—the randomizing polynomials ensure correctness only holds when the subspace membership predicate is evaluated on Vandermonde vectors.

4.2 Function-Private IBE with Minimal Unpredictability

As discussed in Section 1.1, the IBE schemes of Boneh et al. [10] are function private only for identity distributions with min-entropy at least $\lambda + \omega(\log \lambda)$. However, the only inherent restriction required for a meaningful notion of security is that identity distributions have min-entropy $\omega(\log \lambda)$. In this section, starting with predicate encryption schemes for polynomial evaluation constructed in Section 4.1, we construct an IBE scheme satisfying function privacy with only a super-logarithmic min-entropy restriction on identity distributions.

Scheme. Consider a predicate encryption scheme for the class of *linear* predicates $\Phi_{<2}^{\text{poly}}$ comprising algorithms (Setup, KeyGen, Enc, Dec). From Section 4.1, such a predicate encryption scheme can be built from any underlying subspace membership scheme for parameters $m = 2$ and $\ell = 3$. Given such a scheme, we construct an IBE scheme $\mathcal{IBE}^{\text{OPT}}$ for the space of identities \mathbb{S} as follows.

- **Setup:** On input 1^λ , the IBE setup algorithm runs $\text{Setup}(1^\lambda)$ to receive (pp, msk) and publishes pp .
- **Key generation:** On input msk and an identity $\text{id} \in \mathbb{S}$, the key generation algorithm constructs a (randomized) polynomial $p_{\text{id}}(x)$ such that $p_{\text{id}}(x) = 0$ if and only if $x = \text{id}$ as follows. The algorithm samples uniform $r \leftarrow \mathbb{S}$ and computes $p_{\text{id}}(x) = r(x - \text{id})$. It then runs the underlying KeyGen algorithm to output $\text{sk}_{\text{id}} \leftarrow \text{KeyGen}(\text{msk}, p_{\text{id}})$.
- **Encryption:** On input pp , an identity id , and a message M , the encryption algorithm computes $\text{Enc}(\text{pp}, \text{id}, M)$.
- **Decryption:** On input pp , a ciphertext c , and a secret key sk , the decryption algorithm simply computes the underlying decryption algorithm to output $M \leftarrow \text{Dec}(\text{pp}, \text{sk}, c)$.

Correctness of the IBE scheme follows from the correctness of the underlying $\Phi_{<2}^{\text{poly}}$ -predicate encryption scheme. Data privacy and anonymity of the IBE scheme (see [11, Definition 2.5]) follows directly from the attribute hiding property of the underlying $\Phi_{<2}^{\text{poly}}$ -predicate encryption scheme. In the theorem that follows, we prove that $\mathcal{IBE}^{\text{OPT}}$ is function-private against minimally unpredictable sources.

Theorem 4.4. *Given any large attribute space inner-product encryption scheme for dimension $\ell = 3$, there exists an IBE scheme function private against (T, k) -block-sources for any $T = \text{poly}(\lambda)$ and $k \geq \omega(\log \lambda)$.*

Proof outline. For simplicity, consider adversaries that query the real-or-random oracle with k -sources (i.e., $T = 1$). As outlined in Section 4.1 we first construct a predicate encryption scheme for $\Phi_{<2}^{\text{poly}}$ that is function private against k' -sources for $k' \geq \log |\mathbb{S}| + \omega(\log \lambda)$. We instantiate $\mathcal{IBE}^{\text{OPT}}$ described above with this predicate encryption scheme.

The proof proceeds by showing that $\text{RoR}^{\text{FP-IBE}}$ queries (see [11, Definition 2.6]) ID can be compiled to distributions over coefficients of linear polynomials $\mathbf{P} = (P_1, P_0)$ such that if $\mathbf{H}_\infty(ID) = k$, then $\mathbf{H}_\infty(\mathbf{P}) = k + \log |\mathbb{S}|$. This allows us to simulate a $\text{RoR}^{\text{FP-IBE}}$ oracle given an oracle $\text{RoR}^{\text{FP-}\Phi}$ for linear polynomials thus showing that $\mathcal{IBE}^{\text{OPT}}$ is function-private against k -sources if the encryption scheme for $\Phi_{<2}^{\text{poly}}$ is function-private against k' -sources. Due to space constraints, the reader is referred to the full version for details [11, Section 5.1].

Fully-secure function-private IBE. Current constructions of inner-product encryption schemes [23, 3] satisfy a selective notion of security where the challenge attributes are chosen by the adversary before seeing the public parameters. Our transformation of inner-product encryption schemes to function-private IBE

schemes with minimal unpredictability is not limited to selective security. Starting from an inner-product encryption scheme satisfying an adaptive version of attribute hiding, we can construct fully-secure IBE schemes. We also note that the standard complexity leveraging approach (see [7, Section 7.1]) gives a generic transformation from selectively-secure IBE to fully-secure IBE. This approach does not modify the key generation algorithm and therefore preserves function privacy.

5 Conclusions and Open Problems

Our work proposes subspace-membership encryption and constructs the first such function-private schemes from any inner-product encryption scheme. We also show its application to constructing function-private polynomial encryption schemes and function-private IBE schemes with minimal unpredictability. In this section, we discuss a few extensions and open problems that arise from this work.

Function privacy from computational assumptions. In this work we construct subspace-membership schemes that are *statistically* function private. Although the construction of inner-product encryption schemes from lattices [3] presents an immediate function-privacy attack, we were unable to find such attacks for the construction from composite-order groups [23] (or its prime order variant [18]). We conjecture that suitable “min-entropy” variants of the decisional Diffie-Hellman assumption [15] have a potential for yielding a proof of computational function privacy for these schemes.

Other predicates. A pre-cursor to the work on predicate encryption supporting inner-products was work on predicate encryption supporting comparison and range queries by Boneh and Waters [13]. They achieve this by constructing predicate encryption supporting an interesting primitive, denoted Hidden-Vector Encryption (HVE). Briefly, in HVE, attributes correspond to vectors over an alphabet Σ and secret keys correspond to vectors over the *augmented* alphabet $\Sigma \cup \{\star\}$. Decryption works if the attributes and secret key match for every coordinate that is not a \star .

HVE can be implemented using inner-product encryption schemes [23] but it breaks function privacy in a rather trivial manner. Formalizing function privacy for HVE does not immediately follow from the notion of function privacy for inner-products because of the role played by \star . The questions of formalizing function privacy (which in turn will imply realistic notions also for encryption supporting range and comparison queries) and designing function-private HVE schemes are left as open problems. It is also open to formalize security and design function-private encryption schemes that support multivariate polynomial evaluation.

Enhanced function privacy. A stronger notion of function privacy, denoted enhanced function privacy [10], asks that an adversary learn nothing more than the minimum necessary from a secret key even given corresponding ciphertexts with attributes that allow successful decryption. Constructing enhanced

function-private schemes for subspace membership and inner products is an interesting line of research that may require new ideas and techniques.

Acknowledgments

This work was supported by NSF, the DARPA PROCEED program, an AFOSR MURI award, a grant from ONR, an IARPA project provided via DoI/NBC, and by Samsung. Opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of DARPA or IARPA.

References

1. Michel Abdalla, Mihir Bellare, Dario Catalano, Eike Kiltz, Tadayoshi Kohno, Tanja Lange, John Malone-Lee, Gregory Neven, Pascal Paillier, and Haixia Shi. Searchable encryption revisited: Consistency properties, relation to anonymous IBE, and extensions. *Journal of Cryptology*, 21(3):350–391, 2008.
2. Michel Abdalla, Mihir Bellare, and Gregory Neven. Robust encryption. In *Proceedings of the 7th Theory of Cryptography Conference*, pages 480–497, 2010.
3. Shweta Agrawal, David Mandell Freeman, and Vinod Vaikuntanathan. Functional encryption for inner product predicates from learning with errors. In *Advances in Cryptology – ASIACRYPT ’11*, pages 21–40, 2011.
4. Shweta Agrawal, Sergey Gorbunov, Vinod Vaikuntanathan, and Hoeteck Wee. Functional encryption: New perspectives and lower bounds. In *Advances in Cryptology – CRYPTO ’13*, pages 500–518, 2013.
5. Joonsang Baek, Reihaneh Safavi-Naini, and Willy Susilo. Public key encryption with keyword search revisited. In *Proceedings on the International Conference Computational Science and Its Applications*, pages 1249–1259, 2008.
6. Mihir Bellare and Adam O’Neill. Semantically-secure functional encryption: Possibility results, impossibility results and the quest for a general definition. Cryptology ePrint Archive, Report 2012/515, 2012.
7. Dan Boneh and Xavier Boyen. Efficient selective identity-based encryption without random oracles. *Journal of Cryptology*, 24(4):659–693, 2011.
8. Dan Boneh, Giovanni Di Crescenzo, Rafail Ostrovsky, and Giuseppe Persiano. Public key encryption with keyword search. In *Advances in Cryptology – EUROCRYPT ’04*, pages 506–522, 2004.
9. Dan Boneh and Michael Hamburg. Generalized identity based and broadcast encryption schemes. In *Advances in Cryptology – ASIACRYPT ’08*, pages 455–470, 2008.
10. Dan Boneh, Ananth Raghunathan, and Gil Segev. Function-private identity-based encryption: Hiding the function in functional encryption. In *Advances in Cryptology – CRYPTO ’13*, pages 461–478, 2013.
11. Dan Boneh, Ananth Raghunathan, and Gil Segev. Function-private subspace-membership encryption and its applications. Cryptology ePrint Archive, Report 2013/403, 2013.
12. Dan Boneh, Amit Sahai, and Brent Waters. Functional encryption: Definitions and challenges. In *Proceedings of the 8th Theory of Cryptography Conference*, pages 253–273, 2011.

13. Dan Boneh and Brent Waters. Conjunctive, subset, and range queries on encrypted data. In *Proceedings of the 4th Theory of Cryptography Conference*, pages 535–554, 2007.
14. Jan Camenisch, Markulf Kohlweiss, Alfredo Rial, and Caroline Sheedy. Blind and anonymous identity-based encryption and authorised private searches on public key encrypted data. In *Proceedings of the 12th International Conference on Practice and Theory in Public Key Cryptography*, pages 196–214, 2009.
15. Ran Canetti. Towards realizing random oracles: Hash functions that hide all partial information. In *Advances in Cryptology – CRYPTO ’97*, pages 455–469, 1997.
16. Benny Chor and Oded Goldreich. Unbiased bits from sources of weak randomness and probabilistic communication complexity. *SIAM Journal on Computing*, 17(2):230–261, 1988.
17. Kai-Min Chung and Salil P. Vadhan. Tight bounds for hashing block sources. In *APPROX-RANDOM*, pages 357–370, 2008.
18. David Mandell Freeman. Converting pairing-based cryptosystems from composite-order groups to prime-order groups. In *Advances in Cryptology – EUROCRYPT ’10*, pages 44–61, 2010.
19. Shafi Goldwasser, Yael Tauman Kalai, Raluca A. Popa, Vinod Vaikuntanathan, and Nikolai Zeldovich. Reusable garbled circuits and succinct functional encryption. In *Proceedings of the 45th Annual ACM Symposium on the Theory of Computing*, pages 555–564, 2013.
20. Philippe Golle, Jessica Staddon, and Brent R. Waters. Secure conjunctive keyword search over encrypted data. In *Proceedings of the 2nd International Conference on Applied Cryptography and Network Security*, pages 31–45, 2004.
21. Sergey Gorbunov, Vinod Vaikuntanathan, and Hoeteck Wee. Functional encryption with bounded collusions via multi-party computation. In *Advances in Cryptology – CRYPTO ’12*, pages 162–179, 2012.
22. Johan Håstad, Russell Impagliazzo, Leonid A. Levin, and Michael Luby. A pseudorandom generator from any one-way function. *SIAM Journal on Computing*, 28(4):1364–1396, 1999.
23. Jonathan Katz, Amit Sahai, and Brent Waters. Predicate encryption supporting disjunctions, polynomial equations, and inner products. In *Advances in Cryptology – EUROCRYPT ’08*, pages 146–162, 2008.
24. Tatsuaki Okamoto and Katsuyuki Takashima. Hierarchical predicate encryption for inner-products. In *Advances in Cryptology – ASIACRYPT ’09*, pages 214–231, 2009.
25. Tatsuaki Okamoto and Katsuyuki Takashima. Adaptively attribute-hiding (hierarchical) inner product encryption. In *Advances in Cryptology – EUROCRYPT ’12*, pages 591–608, 2012.
26. Adam O’Neill. Definitional issues in functional encryption. IACR Cryptology ePrint Archive, Report 2010/556, 2010.
27. Emily Shen, Elaine Shi, and Brent Waters. Predicate privacy in encryption systems. In *Proceedings of the 6th Theory of Cryptography Conference*, pages 457–473, 2009.
28. Elaine Shi, John Bethencourt, Hubert T.-H. Chan, Dawn Song, and Adrian Perig. Multi-dimensional range query over encrypted data. In *IEEE Symposium on Security and Privacy*, pages 350–364, 2007.
29. David Zuckerman. Simulating BPP using a general weak random source. *Algorithmica*, 16(4/5):367–391, 1996.