Function Private Predicate Encryption for Low Min-Entropy Predicates

Sikhar Patranabis, Debdeep Mukhopadhyay, and Somindu C. Ramanna

Department of Computer Science and Engineering
Indian Institute of Technology Kharagpur, India

sikhar.patranabis@iitkgp.ac.in
{debdeep,somindu}@cse.iitkgp.ac.in

Abstract. In this work, we propose new constructions for zero inner-product encryption (ZIPE) and non-zero inner-product encryption (NIPE) from prime-order bilinear pairings, which are both attribute and function private in the public-key setting.

- Our ZIPE scheme is adaptively attribute private under the standard Matrix DDH assumption for unbounded collusions. It is additionally computationally function private under a min-entropy variant of the Matrix DDH assumption for predicates sampled from distributions with super-logarithmic min-entropy. Existing (statistically) function private ZIPE schemes due to Boneh et al. [Crypto'13, Asiacrypt'13] *necessarily* require predicate distributions with significantly larger min-entropy in the public-key setting.
- Our NIPE scheme is adaptively attribute private under the standard Matrix DDH assumption, albeit for bounded collusions. In addition, it achieves computational function privacy under a min-entropy variant of the Matrix DDH assumption for predicates sampled from distributions with super-logarithmic min-entropy. To the best of our knowledge, existing NIPE schemes from bilinear pairings were neither attribute private nor function private.

Our constructions are inspired by the linear FE constructions of Agrawal et al. [Crypto'16] and the simulation secure ZIPE of Wee [TCC'17]. In our ZIPE scheme, we show a novel way of embedding two different hard problem instances in a single secret key - one for unbounded collusion-resistance and the other for function privacy. For NIPE, we introduce new techniques for simultaneously achieving attribute and function privacy. We further show that the two constructions naturally generalize to a wider class of predicate encryption schemes such as subspace membership, subspace non-membership and hidden-vector encryption.

1 Introduction

Predicate encryption (PE) [14, 5, 30] is a modern public-key primitive that enables fine-grained role-based access control on encrypted data, which makes it desirable for a number of real-life applications. In a PE scheme, a single master secret key msk is used to derive several secret keys of the form sk_f, where f is

a Boolean function over Σ. A ciphertext corresponds to an attribute-message pair $(I, M) \in \Sigma \times \mathcal{M}$, where Σ is a pre-defined set of attributes and $\mathcal{M}$ is a set of payload messages. Decryption of a ciphertext corresponding to (I, M) by sk_f reveals M if and only if $f(I) = 1$. Based on the security notion achieved, a PE scheme may be classified into one or more of the categories described below.

Public Attribute PE. In a public attribute PE system, a ciphertext ct on (I, M) leaks no information about the message M to an adversary possessing secret-keys that do not decrypt ct (i.e., sk_f such that $f(I) = 0$). The attribute I, on the other hand, is public. Such schemes are often nomenclatured as *attribute-based encryption* (ABE). Concrete ABE schemes have been proposed for a wide range of Boolean predicates, including equality/identity testing (IBE) [10, 24], keyword search [9, 1], Boolean formulae [29], regular languages [36], general polynomial-size circuits [22, 11, 27], and even Turing machines [25].

Attribute Private PE. In an attribute private PE, the ciphertext ct leaks no information about either the attribute I or the message M to an adversary possessing secret-keys that do not decrypt ct. Concrete instantiations of private attribute PE have been achieved for hidden vector encryption (HVE) [14] that supports, in addition to equality, conjunctive, range and subset predicates, and also for zero-inner-product encryption (ZIPE) [30, 33]. ZIPE has been realized using bilinear maps [30, 33] and also from lattice-based techniques [5, 2, 6].

In a more recent work [37], Wee demonstrated many new techniques for achieving selectively simulation-secure attribute private PE from prime-order bilinear groups under the standard Matrix DDH assumption. The main result of this work is a partially hiding predicate encryption scheme for functions that compute an arithmetic branching program on public attributes, followed by an inner product predicate on private attributes. In the realm of lattices, Gorbunov et al. [28] showed how to construct attribute private PE for all circuits from the learning with errors (LWE) assumption.

Although attribute privacy has been realized for many different predicates from bilinear pairings, it remains open to construct pairing-based attribute private PE for certain simple predicates such as non-zero inner-product encryption (NIPE) [7] and its natural generalization to a broader class of subspace non-membership encryption (SNME) predicates.

Function Private PE. In a function private PE, a secret-key sk_f reveal no information beyond the absolute minimum about the underlying predicate f. Note that the notions of attribute and function privacy for a PE are mutually exclusive in the sense that one does not necessarily imply the other. In the setting of private-key PE, there already exist function private constructions from pairings for predicates such as ZIPE [8, 20]. In fact, using techniques introduced by Brakerski et al. [15], any private-key PE can be made function private in a generic manner. However, in the setting of public-key PE, formalizing a realistic notion of function privacy is significantly more challenging [12, 13].

Consider, for example, an adversary against an IBE scheme who is given a secret-key $\mathsf{sk}_{\mathsf{id}}$ corresponding to an identity id. As long as the adversary has some apriori information that id belongs to a set $\mathcal{S}$ such that $|\mathcal{S}|$ is at most polynomial in the security parameter λ, it can fully recover id from $\mathsf{sk}_{\mathsf{id}}$: it can simply resort to encrypting a random message M under each identity in $\mathcal{S}$, and decrypting using $\mathsf{sk}_{\mathsf{id}}$ to check for a correct recovery.

Hence, in the setting of public-key PE, function privacy can only hold under the minimal assumption that each predicate is sampled from a distribution with min-entropy at least super logarithmic in the security parameter λ [12, 13]. Under similar assumptions, function private public-key constructions have been reported for IBE [12], ZIPE [3] and subspace membership encryption (SME) [13], which is essentially a generalization of ZIPE. These works throw open several interesting questions. We discuss them below.

1. The PE schemes proposed in [12, 13] are inherently restricted to satisfying a *statistical* notion of function privacy. For a vast majority of applications, a relaxed *computational* notion of function privacy suffices. It is currently open to design public-key PE schemes with function privacy in this relaxed computational setting.
2. The function private PE schemes in [12, 13] *necessarily* assume predicate distributions with min-entropy $k \geq \lambda$ (where λ is the security parameter). [1] This is a rather stringent assumption in the context of real-world predicates. An interesting question is whether a public-key PE scheme can be function private for predicate distributions with only super-logarithmic min-entropy.

There are several real-world applications that warrant the study of PE schemes which are simultaneously attribute and function private. These include searching on encrypted data, secure information retrieval, secure mail gateways and payment gateways, and many others. The reader is referred to [12] for an elaborate discussion of these applications.

1.1 Our Contributions

We focus on the following questions discussed in the previous section:

Is it possible to design attribute private PE from bilinear maps for the non-zero inner product functionality?

What is a meaningful definition of function privacy against resource-bounded adversaries?

Can the min-entropy requirements on the underlying predicate distributions be restricted to a bare minimum while defining function privacy?

[1] The PE schemes in [12, 13] are not function private, even in the weaker computational setting, if the min-entropy requirements are relaxed any further.

Are there constructions for public-key PE that are provably function private, with respect to the relaxed definition, under standard computational assumptions?

In this paper, we answer these questions in the affirmative by first presenting a relaxed definition of function privacy taking into account resource bounded adversaries and restricting the min-entropy requirements of the underlying predicate distributions to $\omega(\log \lambda)$. We then present new pairing-based constructions in the public key setting for subspace membership encryption (SME) and subspace non-membership encryption (SNME) that generalize ZIPE and NIPE respectively. Our constructions are adaptively attribute private and computationally function private in tandem, under variants of the well-known matrix Diffie-Hellman (MDDH) assumption.

Our ZIPE scheme is the first to achieve computational function privacy for predicates with super-logarithmic min-entropy. As already mentioned, existing (statistically) function private ZIPE schemes due to Boneh et al. [13] *necessarily* require predicate distributions with significantly larger min-entropy in the public-key setting. Our NIPE scheme is first to achieve both attribute and function privacy under group-theoretic assumptions, albeit in the bounded collusion setting. Existing constructions for NIPE based on group-theoretic assumptions [7, 16] were neither attribute nor function private, even in the bounded collusion setting.

Our key technical contributions may be summarized as follows.

- Relaxing function privacy definition to account for resource-bounded adversraries and underlying predicates sampled from distributions with min-entropy $k = \omega(\log \lambda)$ (λ being the security parameter).
- Introduction of a min-entropy variant of MDDH assumption where the matrix provided in the instance does not have the uniform distribution but guaranteed to have $\omega(\log \lambda)$ min-entropy.
- Simple and efficient constructions for ZIPE and NIPE from prime-order asymmetric bilinear pairings, that are simultaneously attribute and function private under the presumed hardness of matrix DDH and its min-entropy variant, respectively, so long as the predicates are sampled from distributions with super-logarithmic min-entropy.
- Generalizations of the aforementioned constructions to a broader class of predicates, namely SME and SNME.

Our constructions are inspired by the linear FE constructions of Agrawal et al. [6] and the simulation secure ZIPE of Wee [37]. In our SME (and hence ZIPE) scheme, we show a novel way of embedding two different hard problem instances in a single secret key - one for unbounded collusion-resistance and the other for function privacy. With respect to SNME (and hence NIPE), we introduce new techniques for simultaneously achieving attribute and function privacy, albeit in the bounded collusion setting.

1.2 Overview of Results and Techniques

In this section, we briefly explain the core ideas of our attribute private and function private SME/SNME in terms of the simplest cases, namely, ZIPE/NIPE. The security of our constructions follow from different variants of the Matrix DDH assumption over both source groups of a bilinear pairing.

The Matrix DDH assumption in a group $\mathbb{G}$ of prime order q given by a generator g requires distinguishing between two distributions – $(g^{\mathbf{A}}, g^{\mathbf{Ar}})$ and $(g^{\mathbf{A}}, g^{\mathbf{u}})$ – where $\mathbf{A} \in \mathbb{Z}_q^{(k+1)\times k}$, $\mathbf{r} \in \mathbb{Z}_q^k$ and $\mathbf{u} \in \mathbb{Z}_q^{k+1}$ are sampled uniformly and independently from their respective domains (here $k \geq 1$ and it is assumed that $\mathbf{A}$ has full rank with overwhelming probability). For the function privacy proofs we rely on a special form of the MDDH assumption parameterized by (m, n) – an instance (with respect to a group $\mathbb{G} = \langle g \rangle$) consists of $g^{\mathbf{W}}, g^{\mathbf{u}}$ where $\mathbf{W} \xleftarrow{R} \mathbf{V}^*$ for some source distribution $\mathbf{V}^*$ over $\mathbb{Z}_q^{m\times n}$ of min-entropy $\omega(\log \lambda)$ and the task is to determine if $\mathbf{u} = \mathbf{W}^{\mathbf{T}} \cdot \mathbf{y}$ for $\mathbf{y} \xleftarrow{R} \mathbb{Z}_q^m$ or $\mathbf{u}$ is randomly distributed in $\mathbb{Z}_q^n$.

Denote an asymmetric pairing by the 7-tuple $\mathcal{G} = (\mathbb{G}_1, \mathbb{G}_2, \mathbb{G}_T, q, g_1, g_2, e)$ where $|\mathbb{G}_1| = |\mathbb{G}_2| = |\mathbb{G}_T| = q$, g_1, g_2 respectively generate $\mathbb{G}_1, \mathbb{G}_2$ and $e : \mathbb{G}_1 \times \mathbb{G}_2 \to \mathbb{G}_T$ is a non-degenerate, efficiently computable bilinear map. Call $\mathcal{G}$ Matrix DDH-hard if the Matrix DDH assumption holds in both $\mathbb{G}_1$ and $\mathbb{G}_2$.

Zero Inner-Product Encryption (ZIPE). Our attribute and function private ZIPE construction, named Π^{ZIPE}, is inspired by the simulation secure ZIPE scheme of Wee [37]. The public parameters and the master secret key in Π^{ZIPE} are given by

$$\begin{aligned} \mathsf{pp} &= \left(g_1, g_1^{\mathbf{A}}, g_1^{\mathbf{S}_0\cdot\mathbf{A}}, g_1^{\mathbf{S}_1\cdot\mathbf{A}}, \ldots, g_1^{\mathbf{S}_n\cdot\mathbf{A}}, e(g_1, g_2)^{\mathbf{K}\cdot\mathbf{A}}\right), \\ \mathsf{msk} &= (g_2, \mathbf{S}_0, \mathbf{S}_1, \ldots, \mathbf{S}_n, \mathbf{K}, \mathbf{B}_0)\,, \end{aligned}$$

where $\mathbf{A} \xleftarrow{R} \mathbb{Z}_q^{(k+1)\times k}$, $\mathbf{S}_0, \mathbf{S}_1, \ldots, \mathbf{S}_n \xleftarrow{R} \mathbb{Z}_q^{(2k+1)\times(k+1)}$, $\mathbf{K} \xleftarrow{R} \mathbb{Z}_q^{1\times(k+1)}$ and $\mathbf{B}_0 \xleftarrow{R} \mathbb{Z}_q^{(2k+1)\times k}$ are sampled uniformly. A ciphertext ct on attribute vector $\mathbf{x} = (x_1, \ldots, x_n) \in \mathbb{Z}_q^n$ and message M is given by

$$\mathsf{ct} = \left(c_0, \{c_j\}_{j=1}^n, c_{n+1}\right) = \left(g_1^{(\mathbf{A}\cdot\mathbf{r})^{\mathbf{T}}}, \left\{g_1^{((x_j\cdot\mathbf{S}_0+\mathbf{S}_j)\cdot\mathbf{A}\cdot\mathbf{r})^{\mathbf{T}}}\right\}_{j=1}^n, M\cdot e(g_1, g_2)^{(\mathbf{K}\cdot\mathbf{A}\cdot\mathbf{r})^{\mathbf{T}}}\right),$$

for $\mathbf{r} \xleftarrow{R} \mathbb{Z}_q^k$. The secret key $\mathsf{sk}_{\mathbf{w}}$ on a vector $\mathbf{w} = (w_1, \ldots, w_n) \in \mathbb{Z}_q^n$ is defined as

$$\mathsf{sk}_{\mathbf{w}} = \left(h_0, \{h_j\}_{j=1}^n\right) = \left(g_2^{\mathbf{K}+y\sum_{j=1}^n w_j\cdot\mathbf{t}\cdot\mathbf{S}_j}, \left\{g_2^{yw_j\mathbf{t}}\right\}_{j=1}^n\right),$$

where $y \xleftarrow{R} \mathbb{Z}_q$ and $\mathbf{t} = (\mathbf{B}_0 \cdot \mathbf{s})^{\mathbf{T}}$ for $\mathbf{s} \xleftarrow{R} \mathbb{Z}_q^k$.

For correctness, we restrict the message space $\mathcal{M}$ to an exponentially smaller subset of $\mathbb{G}_T$. The decryption algorithm computes

$$M = c_{n+1} \cdot \left(\prod_{j=1}^{n} e(c_j, h_j) \right) \Big/ e(c_0, h_0),$$

which returns the correct message if $\langle \mathbf{x}, \mathbf{w} \rangle = 0$. When $\langle \mathbf{x}, \mathbf{w} \rangle \neq 0$ the message thus computed is uniformly distributed in $\mathbb{G}_T$ and with high probability will be outside $\mathcal{M}$. In such a case, the decryption algorithm may return a symbol $\perp$ indicating failure.

We prove that Π^{ZIPE} is adaptively attribute private assuming the hardness of the decisional MDDH problems in $\mathbb{G}_1$ and $\mathbb{G}_2$. The attribute privacy game asks an adversary to distinguish between encryptions to attribute vectors $\mathbf{x}_0$ and $\mathbf{x}_1$. Or in other words, the adversary is given a challenge ciphertext for $\mathbf{x}_b$ where $b \xleftarrow{R} \{0, 1\}$ and its task is to guess b. Essentially, we need to argue that the components $\{c_j\}_{j=1}^{n}$ in the challenge ciphertext hide the attribute $\mathbf{x}$.

The proof relies on the dual system proof methodology and proceeds through a sequence of games, each changing the distribution of challenge ciphertext and keys. The key steps in the proof are listed below.

1. The reduction first embeds an instance of MDDH in $\mathbb{G}_1$ in the challenge ciphertext to make it *semi-functional*. At this stage, the exponent of ciphertext component c_0 is no longer correlated to $\mathbf{A}$ and this is consistent with the other components.
2. In a series of subsequent games, we turn each secret key provided to the adversary upon a key extract query to *semi-functional* form by embedding MDDH instances in the group $\mathbb{G}_2$. This step is crucial for unbounded collusion resistance.
3. Once the distribution of all keys are modified, we apply a "change of basis" to the challenge ciphertext, and argue that $\mathbf{x}_b$ is information theoretically hidden from the adversary.

We prove the indistinguishability of each pair of consecutive games by resorting to a set of techniques involving dual bases in prime-order bilinear groups (similar techniques have been used in prior works, notably [23, 26, 17]). The reader may refer to Section 4 and the full version [**?**] for details of the proof.

For showing function privacy of Π^{ZIPE}, we rely on the min-entropy variant of the MDDH assumption. In the function privacy experiment, the adversary picks two vector distributions, each component of which is an $\omega(\log \lambda)$-source over $\mathbb{Z}_q$. The challenger samples a vector $\mathbf{w}$ according to one of the distributions, computes a secret key $\mathsf{sk}_{\mathbf{w}}$ for vector $\mathbf{w}$ and gives it to the adversary. The adversary's task is to determine the distribution of $\mathbf{w}$ looking at $\mathsf{sk}_{\mathbf{w}}$. To prove that the secret key hides the distribution from which $\mathbf{w}$ was sampled, we embed an instance of the min-entropy variant of the MDDH assumption in the challenge secret key provided to the adversary. If the instance is sampled from the correct distribution, the secret key is well-formed. On the other hand, if the instance is

uniformly random, the secret key perfectly hides the distribution from which $\mathbf{w}$ was sampled. The reader may refer to Section 4 for the detailed proof.

Note that our ZIPE scheme essentially embeds two different problem instances in the same secret key - an MDDH problem instance over $\mathbb{G}_2$ that is exploited to achieve unbounded collusion-resistance in the attribute privacy experiment, and a min-entropy MDDH instance over $\mathbb{G}_2$, which is the basis for the proof of function privacy. We believe that this "simultaneous embedding" strategy is of independent interest, and may be useful in other applications.

Non-Zero Inner-Product Encryption (NIPE). Our NIPE scheme is inspired by the linear FE construction of Agrawal, Libert and Stehlé [6] referred to as LinFE in what follows. A LinFE ciphertext ct is created by encrypting a vector $\mathbf{x}$ of length n. Decryption of ct by a secret key, generated for a linear function (given by a length n vector $\mathbf{w}$), returns the value of the inner-product $\langle \mathbf{x}, \mathbf{w} \rangle$.

In a NIPE scheme, a ciphertext is associated with a payload message $M$and a vector $\mathbf{x}$ while a secret key corresponds to a vector $\mathbf{y}$. to be encoded in the ciphertext. Decryption algorithm should be designed to return M iff $\langle \mathbf{x}, \mathbf{w} \rangle \neq 0$. To derive NIPE from LinFE, we use two instantiations of the LinFE with independent master secret keys. The public parameters and master secret key for the resulting scheme would be

$$\mathsf{pp} = \left(g, g^{\mathbf{A}}, g^{\mathbf{S}_1}, g^{\mathbf{S}_2}\right) \quad \mathsf{msk} = (\mathbf{S}_1, \mathbf{S}_2).$$

The ciphertext for $(\mathbf{x}, M)$ will result from encoding $\mathbf{x}$ and $M \cdot \mathbf{x}$ using the two individual schemes as shown below:

$$\mathsf{ct} = \left(g^{\mathbf{A}\mathbf{r}_1}, g^{\mathbf{x}+\mathbf{S}_1\mathbf{A}\mathbf{r}_1}, g^{\mathbf{A}\mathbf{r}_2}, g^{M\cdot\mathbf{x}+\mathbf{S}_2\mathbf{A}\mathbf{r}_2}\right).$$

Here $\mathbf{r}_1, \mathbf{r}_2$ are sampled uniformly at random from $\mathbb{Z}_q^k$. A secret key $\mathsf{sk}_{\mathbf{w}} = (\mathbf{w}^{\mathbf{T}}\mathbf{S}_1, \mathbf{w}^{\mathbf{T}}\mathbf{S}_2)$ helps in recovering $g^{M\langle \mathbf{x},\mathbf{w}\rangle}$ and $g^{\langle \mathbf{x},\mathbf{w}\rangle}$ with respect to g. One may recover M by simply computing the discrete logarithm of $g^{M\langle \mathbf{x},\mathbf{w}\rangle}$ by $g^{\langle \mathbf{x},\mathbf{w}\rangle}$ which is possible only when $\langle \mathbf{x}, \mathbf{w} \rangle \neq 0$. The restriction on the inner-products now shifts to the messages that is, the messages have to lie in a polynomial-sized subset of $\mathbb{Z}_q$. A similar technique has been previously used in [4] to construct public revocation and traitor-tracing from LinFE and revocation, in particular, can be seen as a special case of NIPE. However, our naive construction is not sufficient to (simultaneously) achieve attribute privacy and function privacy since the secret key already reveals too much information about $\mathbf{w}$.

To circumvent the problem, we adapt the construction to the bilinear map setting. This is because functions are associated with secret keys and a basic step to ensure privacy of the function encoded in the secret key components is to hide them in the exponents of elements coming from a discrete log hard group. Ciphertext components already live in a cyclic group. Decryption requires combining the ciphertext and key components to recover the message which can be facilitated if the two groups are equipped with a pairing/bilinear map. Furthermore, the secret key is additionally randomized with $y \in \mathbb{Z}_q$ (for the

generalized case of SNME, this would be a vector $\mathbf{y} \in \mathbb{Z}_q^m$ where $\mathbf{w}$ is replaced by a matrix $\mathbf{W} \in \mathbb{Z}_q^{m \times n}$). Randomization is essential for the function privacy proof, which exploits the hardness of a min-entropy variant of the MDDH assumption. We now discuss the construction of a NIPE scheme possessing both attribute and function privacy.

Let $\mathcal{G} = (\mathbb{G}_1, \mathbb{G}_2, \mathbb{G}_T, q, g_1, g_2, e)$ denote an asymmetric bilinear map ensemble. The public parameters and master secret key for our modified scheme Π^{NIPE} would be similar to the naive scheme we described earlier except that pp components now live in $\mathbb{G}_1$.

$$\mathsf{pp} = \left(g_1, g_1^{\mathbf{A}}, g_1^{\mathbf{S}_1}, g_1^{\mathbf{S}_2}\right) \quad \mathsf{msk} = (g_2, \mathbf{S}_1, \mathbf{S}_2)\,.$$

Similarly, the ciphertext for $(\mathbf{x}, M)$ for Π^{NIPE} is given by

$$\mathsf{ct} = \left(g_1^{\mathbf{A}\mathbf{r}_1}, g_1^{\mathbf{x}+\mathbf{S}_1\mathbf{A}\mathbf{r}_1}, g_1^{\mathbf{A}\mathbf{r}_2}, g_1^{M \cdot \mathbf{x}+\mathbf{S}_2\mathbf{A}\mathbf{r}_2}\right),$$

where $\mathbf{r}_1, \mathbf{r}_2$ are uniformly distributed in $\mathbb{Z}_q^k$. Secret key for $\mathbf{w}$ would now be defined as

$$\mathsf{sk}_{\mathbf{w}} = \left(g_2^{y \cdot \mathbf{w}}, g_2^{y \cdot \mathbf{w}^{\mathbf{T}}\mathbf{S}_1}, g_2^{y \cdot \mathbf{w}^{\mathbf{T}}\mathbf{S}_2}\right)$$

randomized by y sampled uniformly from $\mathbb{Z}_q$. During decryption, ciphertext and key components are paired to obtained $e(g_1, g_2)^{My\langle \mathbf{x}, \mathbf{w} \rangle}$ and $e(g_1, g_2)^{y\langle \mathbf{x}, \mathbf{w} \rangle}$. Message M can be recovered by computing the discrete logarithm of the former with respect to the latter, conditioned on $\langle \mathbf{x}, \mathbf{w} \rangle \neq 0$.

Unlike the SME case, we can only prove attribute privacy of our SNME scheme in the bounded collusion model. More precisely, an adversary is allowed to query at most $n - 1$ secret keys, so that the master secret key components $\mathbf{S}_0, \mathbf{S}_1 \ldots, \mathbf{S}_n$ retain sufficient entropy from the adversary's point of view. The proof then proceeds via a sequence of two hybrid experiments, in each of which the proof embeds a fresh MDDH instance in the challenge ciphertext.

We argue that when these instances are sampled from the "random" distribution instead of the "real" distribution, the challenge ciphertext perfectly hides which attribute-message pair among $(\mathbf{x}_0, M_0)$ and $(\mathbf{x}_1, M_1)$ is being encrypted. The argument for perfect hiding relies on hash proof systems [18, 19], similar to those used by Agrawal et al. in proving the security of their linear FE scheme [6]. Finally, the scheme is adaptively secure because the reduction knows the master secret key at any time, which allows it to answer all secret key queries without knowing the challenge attributes beforehand. For more details on the proof, the reader may refer to Section 5 and the full version [**?**].

To prove function privacy, we again rely on the min-entropy variant of the MDDH assumption over the group $\mathbb{G}_2$. This proof is technically very similar to the proof of function privacy for our SME scheme. The reader may refer to Section 5 for the detailed proof.

Hidden Vector Encryption (HVE). We extend our techniques to construct a hidden vector encryption wherein a secret key for a vector $\mathbf{y} \in (\Sigma \cup \{\star\})^n$ allows decryption of a ciphertext on attribute vector $\mathbf{x} \in \Sigma^n$ if for each $j \in [1, n]$,

either $y_j = x_j$ or $y_j = \star$. Although attribute-private HVE is implied by attribute-private SME, the implication does not extend to function privacy. In fact, defining function privacy for HVE itself is tricky given the presence of wildcard characters. We overcome this issue by presenting a weaker notion of function privacy for HVE that allows revealing positions of the *wildcard* ($\star$) characters in a given predicate vector, while hiding the contents of the other "non-wildcard" positions. Also presented is a construction of HVE that is provably function private in this weaker model from bilinear maps. The construction is quite similar to our SME construction, except for certain minor tweaks to account for the presence of wildcard characters. The proofs of attribute and function privacy (in the weak model) also follow analogously.

1.3 Open Problems

Several interesting questions remain unanswered. The construction of NIPE/SNME we present have a restriction – attribute privacy only holds in the bounded collusions model. It would be interesting to obtain constructions free of this restriction. Another problem is to construct efficient function-private PE for richer functionalities such as Boolean and arithmetic span programs from standard assumptions.

1.4 Organization of the Paper

In Section 2, we present the notation, a quick review of bilinear maps and related assumptions followed by definitions of PE and associated security notions. This is followed by a description of min-entropy variants of MDDH assumption required for our proofs. We formalize the relaxed computational notion of function privacy and discuss related issues in Section 3. In Section 4, we present our SME construction followed by proofs of attribute privacy and function privacy. Section 5 describes our construction of SNME. Due to lack of space, we omit details of the proofs. Interested readers are referred to the full version [**?**].The full version also describes a function private hidden vector encryption along with a sketch of its security proof.

2 Background and Preliminary Definitions

In this section, we fix notation, present background material on predicate encryption and recall certain standard computational assumptions in bilinear groups. We also introduce certain *min-entropy* variants of these assumptions useful for our proofs.

2.1 Notation

This section summarizes the notations used throughout the rest of the paper. We write $x \xleftarrow{R} \chi$ to represent that an element x is sampled uniformly at random from

a set/distribution $\mathcal{X}$. The output a of a deterministic algorithm $\mathcal{A}$ is denoted by $x = \mathcal{A}$ and the output a' of a randomized algorithm $\mathcal{A}'$ is denoted by $x' \leftarrow \mathcal{A}'$.

We refer to $\lambda \in \mathbb{N}$ as the security parameter, and denote by $\mathsf{exp}(\lambda)$, $\mathsf{poly}(\lambda)$ and $\mathsf{negl}(\lambda)$ any generic (unspecified) exponential function, polynomial function and negligible function in λ respectively. Note that a function $f : \mathbb{N} \to \mathbb{N}$ is said to be negligible in λ if for every positive polynomial p, $f(\lambda) < 1/p(\lambda)$ when λ is sufficiently large.

For $a, b \in \mathbb{Z}$ such that $a \leq b$, we denote by $[a, b]$ the set of integers lying between a and b (both inclusive). For a finite field $\mathbb{F}_q$ (q being a λ-bit prime) and $m, n \in \mathbb{N}$, we denote by $\mathbb{F}_q^{m \times n}$ the space of all $m \times n$ matrices $\mathbf{W}$ with elements from $\mathbb{F}_q$. We use the short-hand notation $\mathbb{F}_q^m$ to represent the vector space $\mathbb{F}_q^{m \times 1}$. The transpose of a matrix $\mathbf{W} \in \mathbb{F}_q^{m \times n}$ is denoted as $\mathbf{W}^{\mathbf{T}}$. The symbol $\mathbf{0}$ is used to denote an all-zero matrix of appropriate dimension.

Finally, the min-entropy of a random variable Y is denoted as $\mathbf{H}_\infty(\mathrm{Y})$ and is evaluated as $\mathbf{H}_\infty(\mathrm{Y}) = -\log(\max_y \Pr[\mathrm{Y} = y])$. A random variable Y is said to be a k-source if $\mathbf{H}_\infty(\mathrm{Y}) \geq k$.

2.2 Predicate Encryption

Definition 1. (Predicate Encryption). A predicate encryption (PE) scheme for a class of predicates $\mathcal{F}$ over an attribute space Σ and a payload-message space $\mathcal{M}$ is a quadruple of PPT algorithms $\Pi = (\mathsf{Setup}, \mathsf{KeyGen}, \mathsf{Enc}, \mathsf{Dec})$, defined as follows:

- $\mathsf{Setup}(1^\lambda)$: On input the security parameter λ, the setup algorithm generates the public parameter pp and the master secret key msk.
- $\mathsf{KeyGen}(\mathsf{pp}, \mathsf{msk}, f)$: On input the public parameter pp, the master secret key msk and a predicate $f \in \mathcal{F}$, the key-generation algorithm outputs a secret key sk_f.
- $\mathsf{Enc}(\mathsf{pp}, I, M)$: On input the public parameter pp, an attribute $I \in \Sigma$ and a payload message $M \in \mathcal{M}$, the encryption algorithm outputs a ciphertext ct.
- $\mathsf{Dec}(\mathsf{pp}, \mathsf{sk}_f, \mathsf{ct})$: On input the public parameter pp, a ciphertext ct and a secret key sk_f, the decryption algorithm outputs either a payload-message $M \in \mathcal{M}$ or the symbol $\perp$.

Correctness. A PE scheme is said to be functionally correct if for any security parameter $\lambda \in \mathbb{N}$, any predicate $f \in \mathcal{F}$, any attribute $I \in \Sigma$, and any payload message $M \in \mathcal{M}$, letting $(\mathsf{pp}, \mathsf{msk}) \leftarrow \mathsf{Setup}(1^\lambda)$, $\mathsf{sk}_f \leftarrow \mathsf{KeyGen}(\mathsf{pp}, \mathsf{msk}, f)$ and $\mathsf{ct} \leftarrow \mathsf{Enc}(\mathsf{pp}, I, M)$, the following hold:

1. If $f(I) = 1$, $\Pr[\mathsf{Dec}(\mathsf{pp}, \mathsf{ct}, \mathsf{sk}_f) = M] = 1$,
2. If $f(I) = 0$, $\Pr[\mathsf{Dec}(\mathsf{pp}, \mathsf{ct}, \mathsf{sk}_f) = \perp]$ with overwhelmingly large probability,

where the probabilities are computed over the randomness of the $\mathsf{Setup}, \mathsf{KeyGen}$ and Enc algorithms.

Experiment $\mathsf{Expt}^{(b)}_{\mathsf{AP},\Pi,\mathcal{A}}(\lambda)$:

1. The challenger samples $(\mathsf{pp}, \mathsf{msk}) \leftarrow \mathsf{Setup}(1^\lambda)$ and provides pp to $\mathcal{A}$.
2. The adversary $\mathcal{A}$ adaptively issues key-generation queries. For each query predicate f, the challenger responds with

$$\mathsf{sk}_f \leftarrow \mathsf{KeyGen}\left(\mathsf{pp}, \mathsf{msk}, f\right).$$

3. The adversary $\mathcal{A}$ outputs attribute-message pairs (I_0, M_0) and (I_1, M_1), such that for each predicate f queried, it holds that

$$f(I_0) = f(I_1) = 0.$$

   The challenger responds to the adversary $\mathcal{A}$ with the ciphertext

$$\mathsf{ct} \leftarrow \mathsf{Enc}\left(\mathsf{pp}, I_b, M_b\right).$$

4. The adversary $\mathcal{A}$ continues to adaptively issue key-generation queries, subject to the aforementioned restrictions. The challenger responds as above.
5. Eventually, the adversary $\mathcal{A}$ outputs a bit b'.

Fig. 1. The Attribute Privacy Experiment for Predicate Encryption

Attribute Privacy. Define the experiment $\mathsf{Expt}^{(b)}_{\mathsf{AP},\Pi,\mathcal{A}}(\lambda)$ as in Fig. 1 for a PE $\Pi = (\mathsf{Setup}, \mathsf{KeyGen}, \mathsf{Enc}, \mathsf{Dec})$, a security parameter $\lambda \in \mathbb{N}$ and a bit $b \in \{0, 1\}$. Let $\mathbf{Adv}^{\mathsf{AP}}_{\Pi,\mathcal{A}}(\lambda)$ denote the advantage of the adversary $\mathcal{A}$ in the aforementioned experiment, defined as

$$\mathbf{Adv}^{\mathsf{AP}}_{\Pi,\mathcal{A}}(\lambda) := \left| \Pr\left[\mathsf{Expt}^{(0)}_{\mathsf{AP},\Pi,\mathcal{A}}(\lambda) = 1\right] - \Pr\left[\mathsf{Expt}^{(1)}_{\mathsf{AP},\Pi,\mathcal{A}}(\lambda) = 1\right] \right| \leq \mathsf{negl}(\lambda).$$

Definition 2. (Attribute Private PE.) A PE scheme Π is said to be *attribute private* if for all security parameters $\lambda \in \mathbb{N}$ and for all PPT adversaries $\mathcal{A}$, it holds that $\mathbf{Adv}^{\mathsf{AP}}_{\Pi,\mathcal{A}}(\lambda) \leq \mathsf{negl}(\lambda)$.

2.3 Sub-Classes of Predicate Encryption

In this subsection, we recall definitions of certain sub-classes of predicate encryption that are used in the rest of the paper.

Inner Product Encryption. An inner product encryption (IPE) scheme [30, 33] is a PE over an attribute space $\Sigma = \mathbb{F}_q^n$ (q being a λ-bit prime) and a set of Boolean predicates $f_{\mathbf{y}} : \mathbb{F}_q^n \longrightarrow \{0, 1\}$ such that for each $\mathbf{y} \in \mathbb{F}_q^n$ and $\mathbf{x} \in \mathbb{F}_q^n$, we have

$$f_{\mathbf{y}}(\mathbf{x}) = \begin{cases} 1 & \text{if } \langle \mathbf{y}, \mathbf{x} \rangle = 0 \\ 0 & \text{otherwise.} \end{cases}$$

where $\langle \cdot, \cdot \rangle$ denotes the inner product (equivalently, scalar product) of two vectors over $\mathbb{Z}_q$.

Subspace Membership Encryption. Subspace membership encryption (SME) [13] is a generalization of IPE to accommodate general linear subspaces as opposed to only vector spaces. Formally, an SME scheme is is a PE over an attribute space $\Sigma = \mathbb{F}_q^n$ (q being a λ-bit prime) and a set of Boolean predicates $f_{\mathbf{W}} : \mathbb{F}_q^n \longrightarrow \{0,1\}$ such that for each $\mathbf{W} \in \mathbb{F}_q^{m \times n}$ and $\mathbf{x} \in \mathbb{F}_q^n$, we have

$$f_{\mathbf{W}}(\mathbf{x}) = \begin{cases} 1 & \text{if } \mathbf{W} \cdot \mathbf{x} = \mathbf{0} \\ 0 & \text{otherwise.} \end{cases}$$

Non-Zero IPE. Non-zero IPE (NIPE) [7, 16] is the dual of IPE in the sense that it is a PE over an attribute space $\Sigma = \mathbb{F}_q^n$ (q being a λ-bit prime) and a set of Boolean predicates $f_{\mathbf{y}} : \mathbb{F}_q^n \longrightarrow \{0,1\}$ such that for each $\mathbf{y} \in \mathbb{F}_q^n$ and $\mathbf{x} \in \mathbb{F}_q^n$, we have

$$f_{\mathbf{y}}(\mathbf{x}) = \begin{cases} 1 & \text{if } \langle \mathbf{y}, \mathbf{x} \rangle \neq 0 \\ 0 & \text{otherwise.} \end{cases}$$

Subspace Non-Membership Encryption. Subspace non-membership encryption (SNME) is a generalization of NIPE and the dual of SME in the sense that it is a PE over an attribute space $\Sigma = \mathbb{F}_q^n$ (q being a λ-bit prime) and a set of Boolean predicates $f_{\mathbf{W}} : \mathbb{F}_q^n \longrightarrow \{0,1\}$ such that for each $\mathbf{W} \in \mathbb{F}_q^{m \times n}$ and $\mathbf{x} \in \mathbb{F}_q^n$, we have

$$f_{\mathbf{W}}(\mathbf{x}) = \begin{cases} 1 & \text{if } \mathbf{W} \cdot \mathbf{x} \neq \mathbf{0} \\ 0 & \text{otherwise.} \end{cases}$$

2.4 Bilinear Maps and Matrix Diffie-Hellman Assumptions

Let $\mathsf{GroupGen}(1^\lambda)$ be a PPT algorithm that takes as input a security parameter λ, and outputs a tuple of the form $(\mathbb{G}_1, \mathbb{G}_2, \mathbb{G}_T, q, g_1, g_2, e)$, where $\mathbb{G}_1$, $\mathbb{G}_2$ and $\mathbb{G}_T$ are distinct cyclic groups of order q (q being a λ-bit prime), g_1 is a generator for $\mathbb{G}_1$, g_2 is a generator for $\mathbb{G}_2$, and $e : \mathbb{G}_1 \times \mathbb{G}_2 \longrightarrow \mathbb{G}_T$ is an efficiently computable non-degenerate asymmetric bilinear map. Also, let $\mathbf{W} \in \mathbb{Z}_q^{m \times n}$ be an $m \times n$ matrix with entries $\{w_{i,j}\}_{i \in [1,m], j \in [1,n]}$. Throughout the paper, we use the following notations:

- $g_1^{\mathbf{W}}$: set of group elements $\{g_1^{w_{i,j}}\}_{i \in [1,m], j \in [1,n]} \in \mathbb{G}_1^{m \times n}$
- $g_2^{\mathbf{W}}$: set of group elements $\{g_2^{w_{i,j}}\}_{i \in [1,m], j \in [1,n]} \in \mathbb{G}_2^{m \times n}$
- $e(g_1, g_2)^{\mathbf{W}}$: set of group elements $\{e(g_1, g_2)^{w_{i,j}}\}_{i \in [1,m], j \in [1,n]} \in \mathbb{G}_T^{m \times n}$

We now review the matrix Diffie-Hellman (MDDH) assumption over the source groups $\mathbb{G}_1$ and $\mathbb{G}_2$ of a bilinear map.

The $\mathcal{D}_{m,n}$-MDDH Assumption. Let $m, n \in \mathbb{N}$ such that $m > n$, and let $\mathcal{D}_{m,n}$ denote a matrix distribution over $Z_q^{m\times n}$ such that a matrix $\mathbf{W} \xleftarrow{R} \mathcal{D}_{m,n}$ has full rank n with overwhelmingly large probability. The $\mathcal{D}_{m,n}$-MDDH assumption [21] holds over the group $\mathbb{G}_i$ (for $i = 1, 2$) if the distribution ensembles:

$$\left\{\left(g_i^{\mathbf{W}}, g_i^{\mathbf{W}\cdot\mathbf{y}}\right)\right\}_{\mathbf{W}\xleftarrow{R}\mathcal{D}_{m,n},\ \mathbf{y}\xleftarrow{R}\mathbb{Z}_q^n} \text{ and } \left\{\left(g_i^{\mathbf{W}}, g_i^{\mathbf{u}}\right)\right\}_{\mathbf{W}\xleftarrow{R}\mathcal{D}_{m,n},\ \mathbf{u}\xleftarrow{R}\mathbb{Z}_q^m}$$

are computationally indistinguishable.

The $\mathcal{U}_{m,n}$-MDDH Assumption. The $\mathcal{U}_{m,n}$-MDDH assumption is a special instance of the $\mathcal{D}_{m,n}$-MDDH assumption where the matrix distribution $\mathcal{D}_{m,n}$ is the uniform distribution over $Z_q^{m\times n}$.

2.5 A "Min-Entropy" Variant of the MDDH Assumption

In this subsection, we introduce another special instance of the $\mathcal{D}_{k_1,k_2}$-MDDH assumption where the matrix distribution $\mathcal{D}_{k_1,k_2}$ is not uniform, but an ordered collection of $m\times n$ independent $\omega(\log\lambda)$-sources over $\mathbb{Z}_q$. We first state and prove the following lemma.

Lemma 2.1 *Let $\mathcal{W}_{k_1,k_2} = \left[\mathrm{W}_{i,j}\right]_{i\in[1,k_1],j\in[1,k_2]}$ be a matrix of independently distributed random variables such that each random variable $\mathrm{W}_{i,j}$ for $i \in [1, k_1]$ and $j \in [1, k_2]$ is an $\omega(\log\lambda)$-source over $\mathbb{Z}_q$. Then, any matrix $\mathbf{W} \xleftarrow{R} \mathcal{W}_{k_1,k_2}$ has full rank n with overwhelmingly large probability.*

Proof. Let $\mathcal{W}_{k_1,k_2} = \left[\mathrm{W}_{i,j}\right]_{i\in[1,k_1],j\in[1,k_2]}$ be a tuple of $(k_1 \times k_2)$ independently distributed random variables such that each random variable $\mathrm{W}_{i,j}$ for $i \in [1, k_1]$ and $j \in [1, k_2]$ is a t-source over $\mathbb{Z}_q$. Let $\mathbf{W} \xleftarrow{R} \mathcal{W}_{k_1,k_2}$, and let $\widetilde{\mathbf{W}}$ be any arbitrary $k_2 \times k_2$ sub-matrix of $\mathbf{W}$. Then, the probability of the event that $\widetilde{\mathbf{W}}$ has a zero determinant may be quantified as:

$$\begin{aligned}\Pr\left[\mathrm{Det}(\widetilde{\mathbf{W}}) = 0\right] &= 1 - \left(\prod_{j=1}^{k_2-1}\left(1 - 2^{-j\cdot t}\right)\right)\\ &\leq 1 - \left(1 - 2^{-t}\right)^{(k_2-1)} \leq (k_2 - 1)\cdot 2^{-t},\end{aligned}$$

which is negligible for $t = \omega(\log\lambda)$. This completes the proof of Lemma 2.1.

The Min-Entropy-MDDH Assumption. Let $k_1, k_2 \in \mathbb{N}$ with $k_1 > k_2$, and let $\mathcal{W}_{k_1,k_2} = \left[\mathrm{W}_{i,j}\right]_{i\in[1,k_1],j\in[1,k_2]}$ be a tuple of *independently distributed* random variables such that each random variable $\mathrm{W}_{i,j}$ for $i \in [1, k_1]$ and $j \in [1, k_2]$ is an $\omega(\log\lambda)$-source over $\mathbb{Z}_q$. The (k_1, k_2)-min-entropy-MDDH assumption holds over the group $\mathbb{G}_i$ (for $i = 1, 2$) if the distribution ensembles:

$$\left\{\left(g_i^{\mathbf{W}}, g_i^{\mathbf{W}\cdot\mathbf{y}}\right)\right\}_{\mathbf{W}\xleftarrow{R}\mathcal{W}_{k_1,k_2},\ \mathbf{y}\xleftarrow{R}\mathbb{Z}_q^n} \text{ and } \left\{\left(g_i^{\mathbf{W}}, g_i^{\mathbf{u}}\right)\right\}_{\mathbf{W}\xleftarrow{R}\mathcal{W}_{k_1,k_2},\ \mathbf{u}\xleftarrow{R}\mathbb{Z}_q^m}$$

are computationally indistinguishable.
All proofs of function privacy for the schemes presented in this paper are based on the $\mathcal{W}_{m,n}$-MDDH assumption over the group $\mathbb{G}_2$.

2.6 Dual Bases

We briefly recall the concept of "dual bases" [16], along with some useful lemmas that are used in the rest of the proof. Fix some integers $k_0, k_1, k_2 \geq 1$, and let $k = k_0 + k_1 + k_2$. We denote by "basis" a uniformly sampled tuple of matrices

$$(\mathbf{B}_0, \mathbf{B}_1, \mathbf{B}_2) \xleftarrow{R} \mathbb{Z}_q^{k\times k_0} \times \mathbb{Z}_q^{k\times k_1} \times \mathbb{Z}_q^{k\times k_2}.$$

The corresponding "dual basis" is the tuple of matrices

$$(\mathbf{B}_0^*, \mathbf{B}_1^*, \mathbf{B}_2^*) \in \mathbb{Z}_q^{k\times k_0} \times \mathbb{Z}_q^{k\times k_1} \times \mathbb{Z}_q^{k\times k_2},$$

such that the following "non-degeneracy" conditions hold:

$$\mathbf{B}_0^{\mathbf{T}} \cdot \mathbf{B}_0^* = \mathbf{I}_0 \mod q, \quad \mathbf{B}_1^{\mathbf{T}} \cdot \mathbf{B}_1^* = \mathbf{I}_1 \mod q, \quad \mathbf{B}_2^{\mathbf{T}} \cdot \mathbf{B}_2^* = \mathbf{I}_2 \mod q,$$

where $\mathbf{I}_0, \mathbf{I}_1$ and $\mathbf{I}_2$ are identity matrices of appropriate dimensions, and the following "orthogonality" conditions hold:

$$\mathbf{B}_i^{\mathbf{T}} \cdot \mathbf{B}_j^* = \mathbf{0} \mod q \quad \text{for } i, j \in \{0,1,2\}, i \neq j.$$

We also recall some useful lemmas related to dual bases. These lemmas have been used in many prior works, notably [23, 26, 17].

Lemma 2.2 *Let $(\mathbf{B}_0, \mathbf{B}_1, \mathbf{B}_2)$ be a uniformly sampled basis as described above with corresponding dual basis $(\mathbf{B}_0^*, \mathbf{B}_1^*, \mathbf{B}_2^*)$. Any arbitrary vector $\mathbf{u} \in \mathbb{Z}_q^k$ may be uniquely decomposed as $\mathbf{u} = \mathbf{u}_0 + \mathbf{u}_1 + \mathbf{u}_2$ such that*

$$\mathbf{u}_0 = \mathbf{B}_0^* \cdot \mathbf{s}_0, \quad \mathbf{u}_1 = \mathbf{B}_1^* \cdot s_1, \quad \mathbf{u}_2 = \mathbf{B}_2^* \cdot \mathbf{s}_2,$$

for $(\mathbf{s}_0, \mathbf{s}_1, \mathbf{s}_2) \in \mathbb{Z}_q^{k_0} \times \mathbb{Z}_q^{k_1} \times \mathbb{Z}_q^{k_2}$. Additionally, the following holds for each $i \in \{0,1,2\}$:

$$\mathbf{u}^{\mathbf{T}} \cdot \mathbf{B}_i = \mathbf{u}_i^{\mathbf{T}} \cdot \mathbf{B}_i.$$

Lemma 2.3 *Let $(\mathbf{B}_0, \mathbf{B}_1, \mathbf{B}_2)$ be a uniformly sampled basis as described above with corresponding dual basis $(\mathbf{B}_0^*, \mathbf{B}_1^*, \mathbf{B}_2^*)$. Let a uniform vector $\mathbf{u} \xleftarrow{R} \mathbb{Z}_q^k$ be decomposed as $\mathbf{u} = \mathbf{u}_0 + \mathbf{u}_1 + \mathbf{u}_2$ such that*

$$\mathbf{u}_0 = \mathbf{B}_0^* \cdot \mathbf{s}_0, \quad \mathbf{u}_1 = \mathbf{B}_1^* \cdot \mathbf{s}_1, \quad \mathbf{u}_2 = \mathbf{B}_2^* \cdot \mathbf{s}_2,$$

for $(\mathbf{s}_0, \mathbf{s}_1, \mathbf{s}_2) \in \mathbb{Z}_q^{k_0} \times \mathbb{Z}_q^{k_1} \times \mathbb{Z}_q^{k_2}$. Then, for each $i \in \{0,1,2\}$ and for uniform $\mathbf{s}_i' \xleftarrow{R} \mathbb{Z}_q^{k_i}$, it holds that the distributions of the tuples

$$(\mathbf{u}_i, \{\mathbf{u}_j\}_{j\neq i}) \quad \text{and} \quad ((\mathbf{u}_i + \mathbf{B}_i^* \cdot \mathbf{s}_i'), \{\mathbf{u}_j\}_{j\neq i})$$

are statistically indistinguishable.

To see that the aforementioned lemma holds, fix an arbitrary $i \in \{0,1,2\}$, set $\mathbf{u}' = \mathbf{u} + \mathbf{B}_i^* \cdot \mathbf{s}_i'$ for uniform $\mathbf{s}_i' \xleftarrow{R} \mathbb{Z}_q^{k_i}$, decompose $\mathbf{u}' = \mathbf{u}_0' + \mathbf{u}_1' + \mathbf{u}_2'$ and observe that:

- For each $j \in \{0,1,2\} \setminus \{i\}$, we have $\mathbf{u}_j' = \mathbf{u}_j$ by the orthogonality property.
- The distributions of $\mathbf{u}_i$ and $(\mathbf{u}_i + \mathbf{B}_i^* \cdot \mathbf{s}_i')$ are statistically indistinguishable whenever the vectors $\mathbf{u}$ and $\mathbf{s}_i'$ are uniformly random.

Lemma 2.4 *Let $(\mathbf{B}_0, \mathbf{B}_1, \mathbf{B}_2)$ be a uniformly sampled basis as described above with corresponding dual basis $(\mathbf{B}_0^*, \mathbf{B}_1^*, \mathbf{B}_2^*)$. Let (i_0, i_1, i_2) be a fixed but arbitrary permutation of the set $\{0,1,2\}$. Let $\widehat{\mathbf{B}}_{i_0,i_1}$ be a basis for the span of the matrices $\left[\mathbf{B}_{i_0}^* \mid \mathbf{B}_{i_1}^*\right]$ and let $\widehat{\mathbf{B}}_{i_2}$ be a basis for the span of the matrix $\mathbf{B}_{i_2}^*$. Let*

$$\mathbf{t}_0 = (\mathbf{B}_{i_0} \cdot \mathbf{s}_0)^{\mathbf{T}}, \quad \mathbf{t}_1 = (\mathbf{B}_{i_0} \cdot \mathbf{s}_{1,0} + \mathbf{B}_{i_1} \cdot \mathbf{s}_{1,1})^{\mathbf{T}},$$

where $\mathbf{s}_0, \mathbf{s}_{1,0}, \mathbf{s}_{1,1}$ are uniformly sampled vectors of appropriate dimensions. If the $\mathcal{U}_{(k_{i_0}+k_{i_1}),k_{i_0}}$-MDDH assumption holds over the bilinear group $\mathbb{G}_2$, then for all PPT adversaries $\mathcal{A}$, we have

$$\left|\Pr\left[\mathcal{A}\left(D, g_2^{\mathbf{t}_0}\right) = 1\right] - \Pr\left[\mathcal{A}\left(D, g_2^{\mathbf{t}_1}\right) = 1\right]\right| \leq \mathsf{negl}(\lambda),$$

where $D := \left(g_2^{\mathbf{B}_0^*}, g_2^{\mathbf{B}_1^*}, g_2^{\mathbf{B}_2^*}, \widehat{\mathbf{B}}_{i_0,i_1}, \widehat{\mathbf{B}}_{i_2}\right)$.

Note that Lemma 2.4 is essentially the prime-order analog of the well-known subgroup decision assumption over composite order groups, which has classically been used for dual system encryption [32]. The reader may refer to [17] for the proof of this lemma.

3 Function Privacy of SME and SNME

In this section, we define the indistinguishability-based framework for the function privacy of subspace membership encryption (SME) and subspace non-membership encryption (SNME). Let $\Pi = (\mathsf{Setup}, \mathsf{KeyGen}, \mathsf{Enc}, \mathsf{Dec})$ be an SME (equivalently, SNME) scheme. Define the experiment $\mathsf{Expt}^{(b)}_{\mathsf{FP},\Pi,\mathcal{A}}(\lambda)$ as in Fig. 2 for a security parameter $\lambda \in \mathbb{N}$ and a bit $b \in \{0,1\}$. Let $\mathbf{Adv}^{\mathsf{FP}}_{\Pi,\mathcal{A}}(\lambda)$ denote the advantage of the adversary $\mathcal{A}$ in the aforementioned experiment, defined as

$$\mathbf{Adv}^{\mathsf{FP}}_{\Pi,\mathcal{A}}(\lambda) := \left|\Pr\left[\mathsf{Expt}^{(0)}_{\mathsf{FP},\Pi,\mathcal{A}}(\lambda) = 1\right] - \Pr\left[\mathsf{Expt}^{(1)}_{\mathsf{FP},\Pi,\mathcal{A}}(\lambda) = 1\right]\right| \leq \mathsf{negl}(\lambda).$$

Definition 3. (Function Private SME.) An SME scheme Π is said to be *function private* if for all security parameters $\lambda \in \mathbb{N}$ and for all PPT adversaries $\mathcal{A}$, it holds that $\mathbf{Adv}^{\mathsf{FP}}_{\Pi,\mathcal{A}}(\lambda) \leq \mathsf{negl}(\lambda)$.

Definition 4. (Function Private SNME.) An SNME scheme Π is said to be *function private* if for all security parameters $\lambda \in \mathbb{N}$ and for all PPT adversaries $\mathcal{A}$, it holds that $\mathbf{Adv}^{\mathsf{FP}}_{\Pi,\mathcal{A}}(\lambda) \leq \mathsf{negl}(\lambda)$.

Experiment $\mathsf{Expt}^{(b)}_{\mathsf{FP},\Pi,\mathcal{A}}(\lambda)$:

1. The challenger samples $(\mathsf{pp}, \mathsf{msk}) \leftarrow \mathsf{Setup}(1^\lambda)$ and provides pp to $\mathcal{A}$.
2. The adversary $\mathcal{A}$ adaptively issues key-generation queries. For each queried predicate matrix $\mathbf{W}$, the challenger responds with
$$\mathsf{sk}_{\mathbf{W}} \leftarrow \mathsf{KeyGen}\left(\mathsf{pp}, \mathsf{msk}, \mathbf{W}\right).$$
3. The adversary $\mathcal{A}$ outputs circuits of the form
$$\mathcal{W}_0 = \left[\mathrm{W}^{(0)}_{i,j}\right]_{i\in[1,m],j\in[1,n]}, \quad \mathcal{W}_1 = \left[\mathrm{W}^{(1)}_{i,j}\right]_{i\in[1,m],j\in[1,n]},$$
representing joint distributions over $\mathbb{F}_q^{m\times n}$, with the following restrictions:
   (a) For each $i \in [1,m], j \in [1,n]$ and $\tilde{b} \in \{0,1\}$, $\mathrm{W}^{(\tilde{b})}_{i,j}$ represents an $\omega(\log\lambda)$-source over $\mathbb{F}_q$.
   (b) For each $i, i' \in [1,m], j, j' \in [1,n]$ and $\tilde{b} \in \{0,1\}$, $\mathrm{W}^{(\tilde{b})}_{i,j}$ and $\mathrm{W}^{(\tilde{b})}_{i',j'}$ represent mutually independent distributions.

   The challenger samples $\mathbf{W} \xleftarrow{R} \mathcal{W}_b$ and responds to the adversary $\mathcal{A}$ with the secret-key
$$\mathsf{sk}_{\mathbf{W}} = \mathsf{KeyGen}(\mathsf{msk}, \mathbf{W}).$$
4. The adversary $\mathcal{A}$ continues to adaptively issue key-generation queries. The challenger responds as above.
5. Eventually, the adversary $\mathcal{A}$ outputs a bit b'.

Fig. 2. The Function Privacy Experiment for SME and SNME

The Mutual Independence Condition. Observe that the function privacy experiment requires the adversarially chosen distributions $\mathcal{W}_0$ and $\mathcal{W}_1$ to be constructed such that the individual component distributions are both "mutually independent" and "sufficiently unpredictable". A stronger notion of function privacy could allow these components to be "arbitrarily correlated", so long as they are "individually" sufficiently unpredictable. As shown in [13], such a notion is impossible to satisfy. In other words, if arbitrary correlations were allowed, the adversary $\mathcal{A}$ in the function privacy experiment can always create challenge distributions that satisfy the unpredictability requirement, but secret keys for matrices from these distributions are easily distinguishable. We present a brief illustration here for the sake of completeness.

Consider an IPE scheme (equivalently, an SME scheme of dimension $m = 1$) and an adversary $\mathcal{A}$ in the function privacy experiment that chooses the challenge distributions as:

$$\mathcal{W}_0 = \left[\mathrm{W}^{(0)}_1, 2\mathrm{W}^{(0)}_1, \mathrm{W}^{(0)}_2, \ldots, \mathrm{W}^{(0)}_{n-1}\right], \quad \mathcal{W}_1 = \left[\mathrm{W}^{(1)}_1, \mathrm{W}^{(1)}_2, \ldots, \mathrm{W}^{(1)}_{n-1}, 2\mathrm{W}^{(1)}_{n-1}\right],$$

where for each $j \in [1, n-1]$ and $\tilde{b} \in \{0,1\}$, $\mathrm{W}_j^{(\tilde{b})}$ represents a *uniform* source over $\mathbb{F}_q$. Clearly, each individual distribution has min-entropy $\log q = \omega(\log \lambda)$; yet, secret keys for vectors sampled from $\mathcal{W}_0$ can be distinguished from secret keys for vectors sampled from $\mathcal{W}_1$ with non-negligible advantage as follows: encrypt a message M under two attribute vectors $\mathbf{x}_0$ and $\mathbf{x}_1$ defined as:

$$\mathbf{x}_0 = (2, -1, 0, \ldots, 0)\,, \quad \mathbf{x}_1 = (0, \ldots, 0, 2, -1)\,,$$

and see which of the two ciphertexts decrypts correctly under the challenge secret key. This justifies the mutual independence criteria imposed in the function privacy experiment.

Multi-Challenge vs. Single-Challenge. Observe that the aforementioned function privacy definition for SME/SNME is "single-challenge" in the sense that the function privacy experiment allows the adversary a single challenge query. In fact, as the adversary is also given access to the key-generation oracle, the "single-challenge" definition is polynomially equivalent to a "multi-challenge" variant where the adversary is allowed polynomially many challenge queries. This equivalence may be proved by a hybrid argument (originally proposed in [13]), 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.

4 Function Private SME

In this section, we present the construction of an SME scheme that achieves computational function privacy whenever the predicate matrices are sampled from distributions with min-entropy $\omega(\log \lambda)$. In contrast, the SME scheme of Boneh *et al.* [13] is statistically function private, albeit for predicate matrices sampled from distributions with min-entropy slightly larger than λ.

Attribute and function privacy guarantees of our scheme follow from variants of the general $\mathcal{D}$-MDDH assumption in the standard model. More specifically, attribute privacy can be based on the $\mathcal{U}_{k+1,k}$-MDDH assumption in $\mathbb{G}_1$ and $\mathcal{U}_{2k,k}$-MDDH assumption in $\mathbb{G}_2$, while function privacy follows from the $\mathcal{W}_{m,n}$-MDDH assumption described in Section 2.4. The scheme is described below, while the proofs of attribute and function privacy are presented subsequently.

4.1 The Construction

Let $\mathsf{GroupGen}(1^\lambda)$ be a PPT algorithm that takes as input a security parameter $\lambda \in \mathbb{N}$, and outputs the tuple $(\mathbb{G}_1, \mathbb{G}_2, \mathbb{G}_T, q, g_1, g_2, e)$, where $\mathbb{G}_1$, $\mathbb{G}_2$ and $\mathbb{G}_T$ are cyclic groups of prime order q (q being a λ-bit prime), g_1 is a generator for $\mathbb{G}_1$, g_2 is a generator for $\mathbb{G}_2$, and $e : \mathbb{G}_1 \times \mathbb{G}_2 \longrightarrow \mathbb{G}_T$ is an efficiently computable non-degenerate asymmetric bilinear map. Our scheme Π^{SME} is parameterized by $m, n = \mathsf{poly}(\lambda)$ in the sense that it supports predicate matrices of the form

$\mathbf{W} \in \mathbb{Z}_q^{m \times n}$, and attribute vectors of the form $\mathbf{x} \in \mathbb{Z}_q^n$. Finally, the payload message space $\mathcal{M}$ is assumed to a "super-polynomially smaller" subset of $\mathbb{G}_T$, namely $|\mathcal{M}| < |\mathbb{G}_T|^{1/2}$. Our scheme works as follows. [2]

- $\mathsf{Setup}(1^\lambda)$: Uniformly sample $(\mathbb{G}_1, \mathbb{G}_2, \mathbb{G}_T, q, g_1, g_2, e) \leftarrow \mathsf{GroupGen}(1^\lambda)$. Also, uniformly sample

$$\mathbf{A} \xleftarrow{R} \mathbb{Z}_q^{(k+1)\times k}, \quad \mathbf{S}_0, \mathbf{S}_1, \ldots, \mathbf{S}_n \xleftarrow{R} \mathbb{Z}_q^{(2k+1)\times(k+1)},$$

$$\mathbf{K} \xleftarrow{R} \mathbb{Z}_q^{1\times(k+1)}, \quad \mathbf{B}_0 \xleftarrow{R} \mathbb{Z}_q^{(2k+1)\times k}$$

  for some constant $k > 0$. Output

$$\mathsf{pp} = \left(g_1, g_1^{\mathbf{A}}, g_1^{\mathbf{S}_0 \cdot \mathbf{A}}, g_1^{\mathbf{S}_1 \cdot \mathbf{A}}, \ldots, g_1^{\mathbf{S}_n \cdot \mathbf{A}}, e(g_1, g_2)^{\mathbf{K} \cdot \mathbf{A}} \right),$$
$$\mathsf{msk} = (g_2, \mathbf{S}_0, \mathbf{S}_1, \ldots, \mathbf{S}_n, \mathbf{K}, \mathbf{B}_0)\,.$$

- $\mathsf{KeyGen}(\mathsf{pp}, \mathsf{msk}, \mathbf{W})$: Parse the predicate matrix $\mathbf{W} \in \mathbb{Z}_q^{m \times n}$ as

$$\mathbf{W} = \left[w_{i,j} \right]_{i \in [1,m], j \in [1,n]}.$$

  Uniformly sample $\mathbf{s} \xleftarrow{R} \mathbb{Z}_q^k$ and set $\mathbf{t} = (\mathbf{B}_0 \cdot \mathbf{s})^{\mathbf{T}}$. Finally, pick uniform $y_1, \ldots, y_m \xleftarrow{R} \mathbb{Z}_q$ and output $\mathsf{sk}_{\mathbf{W}} = \left(\{h_j\}_{j \in [0,n]} \right)$ where

$$h_0 = g_2^{\left(\mathbf{K} + \sum_{i=1}^m y_i \cdot \left(\sum_{j=1}^n w_{i,j} \cdot \mathbf{t} \cdot \mathbf{S}_j \right) \right)^{\mathbf{T}}},$$
$$h_j = g_2^{\left(\sum_{i=1}^m y_i \cdot w_{i,j} \cdot \mathbf{t} \right)^{\mathbf{T}}} \quad \text{for } j \in [1, n].$$

- $\mathsf{Enc}(\mathsf{pp}, \mathbf{x}, M)$: Given an attribute vector $\mathbf{x} = \left[x_1 \ \ldots \ x_n \right]^{\mathbf{T}} \in \mathbb{Z}_q^n$ and a message $M \in \mathcal{M} \subset \mathbb{G}_T$, uniformly sample $\mathbf{r} \xleftarrow{R} \mathbb{Z}_q^k$ and output $\mathsf{ct} = \left(\{c_j\}_{j \in [0,n+1]} \right)$ where

$$c_0 = g_1^{(\mathbf{A} \cdot \mathbf{r})^{\mathbf{T}}}$$
$$c_j = g_1^{((x_j \cdot \mathbf{S}_0 + \mathbf{S}_j) \cdot \mathbf{A} \cdot \mathbf{r})^{\mathbf{T}}} \quad \text{for } j \in [1, n]$$
$$c_{n+1} = M \cdot e(g_1, g_2)^{(\mathbf{K} \cdot \mathbf{A} \cdot \mathbf{r})^{\mathbf{T}}}$$

- $\mathsf{Dec}(\mathsf{pp}, \mathsf{sk}_{\mathbf{W}}, \mathsf{ct})$: Parse the ciphertext as $\mathsf{ct} = \left(\{c_j\}_{j \in [0,n+1]} \right)$ and the secret key as $\mathsf{sk}_{\mathbf{W}} = \left(\{h_j\}_{j \in [0,n]} \right)$. Compute

$$M = \left(c_{n+1} \cdot \prod_{j=1}^{n} e\left(c_j, h_j\right) \right) \Big/ e\left(c_0, h_0\right).$$

  If $M \in \mathcal{M}$, output M. Otherwise, output $\perp$.

[2] The restriction on the size of the message space $\mathcal{M}$ is necessary for correctness as explained subsequently. Note that this restriction does not prevent $\mathcal{M}$ from being exponentially large.

Correctness. To see that the aforementioned scheme is functionally correct, observe the following.

$$\begin{aligned}
\prod_{j=1}^{n} e\left(c_j, h_j\right) &= \prod_{j=1}^{n} e(g_1, g_2)^{\left(\sum_{i=1}^{m} y_i \cdot w_{i,j} \cdot \mathbf{t} \cdot (x_j \cdot \mathbf{S}_0 + \mathbf{S}_j) \cdot \mathbf{A} \cdot \mathbf{r}\right)^{\mathbf{T}}} \\
&= e(g_1, g_2)^{\left(\left(\sum_{j=1}^{n} \sum_{i=1}^{m} y_i \cdot w_{i,j} \cdot x_j \cdot \mathbf{t} \cdot \mathbf{S}_0 + \sum_{j=1}^{n} \sum_{i=1}^{m} y_i \cdot w_{i,j} \cdot \mathbf{t} \cdot \mathbf{S}_j\right) \cdot \mathbf{A} \cdot \mathbf{r}\right)^{\mathbf{T}}} \\
&= e(g_1, g_2)^{\left(\sum_{i=1}^{m} y_i \cdot \sum_{j=1}^{n} w_{i,j} \cdot x_j \cdot \mathbf{t} \cdot \mathbf{S}_0 \cdot \mathbf{A} \cdot \mathbf{r}\right)^{\mathbf{T}}} \cdot e(g_1, g_2)^{\left(\sum_{i=1}^{m} y_i \cdot \left(\sum_{j=1}^{n} w_{i,j} \cdot \mathbf{t} \cdot \mathbf{S}_j\right) \cdot \mathbf{A} \cdot \mathbf{r}\right)^{\mathbf{T}}} \\
&= e(g_1, g_2)^{\left(\sum_{i=1}^{m} y_i \cdot \sum_{j=1}^{n} w_{i,j} \cdot x_j \cdot \mathbf{t} \cdot \mathbf{S}_0 \cdot \mathbf{A} \cdot \mathbf{r}\right)^{\mathbf{T}}} \cdot e\left(g_1^{(\mathbf{A} \cdot \mathbf{r})^{\mathbf{T}}}, g_2^{\left(\sum_{i=1}^{m} y_i \cdot \left(\sum_{j=1}^{n} w_{i,j} \cdot \mathbf{t} \cdot \mathbf{S}_j\right)\right)^{\mathbf{T}}}\right) \\
&= M \cdot \left(c_{n+1}\right)^{-1} \cdot e\left(c_0, h_0\right) \cdot e(g_1, g_2)^{((\mathbf{y} \cdot \mathbf{W} \cdot \mathbf{x}) \cdot \mathbf{t} \cdot \mathbf{S}_0 \cdot \mathbf{A} \cdot \mathbf{r})^{\mathbf{T}}}
\end{aligned}$$

where $\mathbf{y} = \begin{bmatrix} y_1 \dots y_m \end{bmatrix}$. Hence, when $\mathbf{W} \cdot \mathbf{x} = \mathbf{0} \mod q$, the decryption algorithm recovers M correctly. On the other hand, when $\mathbf{W} \cdot \mathbf{x} \neq \mathbf{0} \mod q$ the distribution of M such that M satisfies the decryption equation is uniformly random over $\mathbb{G}_T$, and hence, with overwhelmingly large probability over the randomness of KeyGen and Enc, the decryption algorithm returns $\perp$. [3]

4.2 Attribute Privacy

We state and prove the following theorem.

Theorem 4.1 *If the* $\mathcal{U}_{k+1,k}$-MDDH *assumption holds over the group* $\mathbb{G}_1$ *and the* $\mathcal{U}_{2k,k}$-MDDH *assumption holds over the group* $\mathbb{G}_2$*, then for all PPT adversaries* $\mathcal{A}$*, we have* $\mathbf{Adv}^{\mathsf{AP}}_{\Pi^{\mathrm{SME}},\mathcal{A}}(\lambda) \leq \mathsf{negl}(\lambda)$.

Proof. The proof proceeds through a sequence of experiments, beginning with the "real" attribute privacy experiment and ending with an experiment where the adversary has no advantage. We consider a variant of the "real" attribute privacy experiment where the challenge messages M_0 and M_1 are chosen to be equal by the adversary. One can reduce the case for $M_0 \neq M_1$ to this case by arguing that an encryption of M_b for $b \in \{0, 1\}$ is indistinguishable from an encryption of M_0 [16, 37]. Hence, it is sufficient to assume that $M_0 = M_1$ in the hybrid experiments presented next. [4]

Expt-0. This is the "real" experiment. In this experiment, the adversary $\mathcal{A}$ is given the public parameter pp. The adversary chooses two (distinct) vector-message pairs $(\mathbf{x}_0, M_0), (\mathbf{x}_1, M_1) \in \mathbb{Z}_q^n \times \mathcal{M}$, such that

$$\mathbf{x}_b = \begin{bmatrix} x_{1,b} & x_{2,b} & \dots & x_{n,b} \end{bmatrix}^{\mathbf{T}} \quad \text{for each } b \in \{0, 1\}.$$

[3] The argument follows from the fact that both $\mathbf{y}$ and $\mathbf{r}$ are uniformly random vectors in $\mathbb{Z}_q^m$ and $\mathbb{Z}_q^k$, respectively, and $|\mathcal{M}| < |\mathbb{G}_T|^{1/2}$.

[4] Due to paucity of space, we only provide brief proof sketches in several cases. We refer the reader to the full version of the paper [**?**] for the detailed proofs.

and $M_0 = M_1$. In addition, the adversary (adaptively) issues a maximum of Q key generation queries (for some fixed polynomial $Q = Q(\lambda)$) corresponding to predicate matrices the form $\mathbf{W}_1, \ldots, \mathbf{W}_Q \in \mathbb{Z}_q^{m \times n}$, subject to the restriction that

$$(\mathbf{W}_\ell \cdot \mathbf{x}_0 \neq 0 \mod q) \wedge (\mathbf{W}_\ell \cdot \mathbf{x}_1 \neq \mathbf{0} \mod q) \text{ for each } \ell \in [1, Q].$$

It receives in response $\left(\mathsf{ct}^*, \mathsf{sk}_{\mathbf{W}_1}, \ldots, \mathsf{sk}_{\mathbf{W}_Q}\right)$, where

$$\begin{aligned} \mathsf{ct}^* &\leftarrow \mathsf{Enc}(\mathsf{pp}, \mathbf{x}_b, M_0) \quad \text{for some random } b \xleftarrow{R} \{0,1\}, \\ \mathsf{sk}_{\mathbf{W}_\ell} &\leftarrow \mathsf{KeyGen}(\mathsf{pp}, \mathsf{msk}, \mathbf{W}_\ell) \quad \text{for each } \ell \in [1, Q]. \end{aligned}$$

Finally, it outputs a bit b'. Let $P_{\mathcal{A},0}$ denote the probability that $b = b'$.

Expt-1. This experiment is identical to Expt-0 except for the manner in which the challenge ciphertext ct^* is generated. Namely, the challenger $\mathcal{B}$ uniformly samples $\mathbf{r} \xleftarrow{R} \mathbb{Z}_q^k$ and uses the master secret key components $\mathbf{S}_0, \mathbf{S}_1, \ldots, \mathbf{S}_n, \mathbf{K}$ to generate the ciphertext $\mathsf{ct}^* = \left(\{c_j\}_{j \in [0,n+1]}\right)$ as

$$c_0 = g_1^{(\mathbf{A} \cdot \mathbf{r})^\mathbf{T}}, \quad \left\{ c_j = (c_0)^{(x_{j,b} \cdot \mathbf{S}_0 + \mathbf{S}_j)^\mathbf{T}} \right\}_{j \in [1,n]}, \quad c_{n+1} = M_0 \cdot e\left(c_0, g_2\right)^{\mathbf{K}^\mathbf{T}}.$$

Note that for each $j \in [1, n]$, we essentially have $c_j = g_1^{\left(\mathbf{v}_j^{(1)}\right)^\mathbf{T}}$, where

$$\mathbf{v}_j^{(1)} = (x_{j,b} \cdot \mathbf{S}_0 + \mathbf{S}_j) \cdot \mathbf{A} \cdot \mathbf{r}.$$

Let $P_{\mathcal{A},1}$ denote the probability that $b = b'$, where b' is the bit output by the adversary $\mathcal{A}$ at the end of Expt-1. Observe that the challenge ciphertext ct^* in Expt-1 has the same distribution as in Expt-0. Hence, we have $P_{\mathcal{A},1} = P_{\mathcal{A},0}$.

Expt-2. This experiment is identical to Expt-1 except for the manner in which the challenge ciphertext ct^* is generated. Namely, the challenger $\mathcal{B}$ uniformly samples $\mathbf{u} \xleftarrow{R} \mathbb{Z}_q^{k+1}$, and generates the ciphertext $\mathsf{ct}^* = \left(\{c_j\}_{j \in [0,n+1]}\right)$ as

$$c_0 = g_1^{\mathbf{u}^\mathbf{T}}, \quad \left\{ c_j = (c_0)^{(x_{j,b} \cdot \mathbf{S}_0 + \mathbf{S}_j)^\mathbf{T}} \right\}_{j \in [1,n]}, \quad c_{n+1} = M_0 \cdot e\left(c_0, g_2\right)^{\mathbf{K}^\mathbf{T}}.$$

Note that for each $j \in [1, n]$, we essentially have $c_j = g_1^{\left(\mathbf{v}_j^{(2)}\right)^\mathbf{T}}$, where

$$\mathbf{v}_j^{(2)} = \boxed{(x_{j,b} \cdot \mathbf{S}_0 + \mathbf{S}_j) \cdot \mathbf{u}}.$$

Let $P_{\mathcal{A},2}$ denote the probability that $b = b'$, where b' is the bit output by the adversary $\mathcal{A}$ at the end of Expt-2. We state the following lemma.

Lemma 4.1 *For all PPT adversaries $\mathcal{A}$, $|P_{\mathcal{A},2} - P_{\mathcal{A},1}| \leq \mathsf{negl}(\lambda)$.*

The proof of this lemma follows directly from the $\mathcal{U}_{k+1,k}$-MDDH assumption over the group $\mathbb{G}_1$. More specifically, given a PPT adversary $\mathcal{A}$ that can distinguish between between its views in Expt-1 and Expt-2 with non-negligible probability, one can construct a PPT algorithm that can distinguish between the ensembles

$$\left\{\left(g_1^{\mathbf{A}}, g_1^{\mathbf{A}\cdot\mathbf{r}}\right)\right\}_{\mathbf{A}\xleftarrow{R}\mathbb{Z}_q^{(k+1)\times k}, \mathbf{r}\xleftarrow{R}\mathbb{Z}_q^k} \text{ and } \left\{\left(g_1^{\mathbf{A}}, g_1^{\mathbf{u}}\right)\right\}_{\mathbf{A}\xleftarrow{R}\mathbb{Z}_q^{(k+1)\times k},\ \mathbf{u}\xleftarrow{R}\mathbb{Z}_q^{k+1}}$$

with non-negligible probability. Quite evidently, the existence of such a PPT algorithm violates the $\mathcal{U}_{k+1,k}$-MDDH assumption over the group $\mathbb{G}_1$.

Expt-3. This experiment is identical to Expt-2 except for the manner in which the challenge ciphertext ct^* is generated. Namely, the challenger $\mathcal{B}$ uniformly samples a basis

$$(\mathbf{B}_0, \mathbf{B}_1, \mathbf{B}_2) \in \mathbb{Z}_q^{(2k+1)\times k} \times \mathbb{Z}_q^{(2k+1)\times 1} \times \mathbb{Z}_q^{(2k+1)\times k},$$

with corresponding dual basis $(\mathbf{B}_0^*, \mathbf{B}_1^*, \mathbf{B}_2^*)$, and uses $\mathbf{B}_0$ as part of the master secret key msk. It samples $\mathbf{u} \xleftarrow{R} \mathbb{Z}_q^{k+1}$ and decomposes $\mathbf{S}_0 \cdot \mathbf{u} \in \mathbb{Z}_q^{2k+1}$ as

$$\mathbf{S}_0 \cdot \mathbf{u} = \mathbf{u}_0 + \mathbf{u}_1 + \mathbf{u}_2,$$

such that

$$\mathbf{u}_0 = \mathbf{B}_0^* \cdot \mathbf{s}_0, \quad \mathbf{u}_1 = \mathbf{B}_1^* \cdot s_1, \quad \mathbf{u}_2 = \mathbf{B}_2^* \cdot \mathbf{s}_2 \quad \text{for some } \mathbf{s}_0, \mathbf{s}_2 \in \mathbb{Z}_q^k, s_1 \in \mathbb{Z}_q.$$

Note that such a decomposition always exists by Lemma 2.2. The challenger $\mathcal{B}$ then generates the ciphertext $\mathsf{ct}^* = \left(\{c_j\}_{j\in[0,n+1]}\right)$ as

$$c_0 = g_1^{\mathbf{u}^\mathbf{T}}, \quad \left\{c_j = g_1^{\left(\mathbf{v}_j^{(3)}\right)^\mathbf{T}}\right\}_{j\in[1,n]}, \quad c_{n+1} = M_0 \cdot e\left(c_0, g_2\right)^{\mathbf{K}^\mathbf{T}},$$

where for each $j \in [1, n]$, we have

$$\mathbf{v}_j^{(3)} = x_{j,b} \cdot \mathbf{u}_0 + \boxed{x_{j,1-b} \cdot \mathbf{u}_1} + x_{j,b} \cdot \mathbf{u}_2 + \mathbf{S}_j \cdot \mathbf{u}.$$

Let $P_{\mathcal{A},3}$ denote the probability that $b = b'$, where b' is the bit output by the adversary $\mathcal{A}$ at the end of Expt-3. We state the following lemma.

Lemma 4.2 *For all unbounded adversaries $\mathcal{A}$, $|P_{\mathcal{A},3} - P_{\mathcal{A},2}| \leq \mathsf{negl}(\lambda)$.*

Proof Sketch. To prove Lemma 4.2, it is sufficient to prove that for each $j \in [1, n]$ and for all $\mathbf{x}_0, \mathbf{x}_1 \in \mathbb{Z}_q^n$, the distributions of $\mathbf{v}_j^{(2)}$ and $\mathbf{v}_j^{(3)}$ are statistically close. Informally, the proof is based on the following observations and a simple application of Lemma 2.3.

1. If one were to decompose $\mathbf{S}_j \cdot \mathbf{u}$ for $j \in [1, n]$ as

$$\mathbf{S}_j \cdot \mathbf{u} = \mathbf{u}_{j,0} + \mathbf{u}_{j,1} + \mathbf{u}_{j,2},$$

such that

$$\mathbf{u}_{j,0} = \mathbf{B}_0^* \cdot \mathbf{s}_{j,0}, \quad \mathbf{u}_{j,1} = \mathbf{B}_1^* \cdot s_{j,1}, \quad \mathbf{u}_{j,2} = \mathbf{B}_2^* \cdot \mathbf{s}_{j,2},$$

for some $\mathbf{s}_{j,0}, \mathbf{s}_{j,2} \in \mathbb{Z}_q^k, s_{j,1} \in \mathbb{Z}_q$, then the public parameter pp and the secret keys $\mathsf{sk}_{\mathbf{W}_1}, \ldots, \mathsf{sk}_{\mathbf{W}_Q}$ statistically hide $\mathbf{u}_{j,1}$ for $j \in [1, n]$. In other words, in the view of an unbounded adversary, the distribution of $\mathbf{u}_{j,1}$ is statistically indistinguishable from that of a uniformly random vector in the span of $\mathbf{B}_1^*$. The reasoning behind this observation is detailed in the full version [**?**].

2. For each $j \in [1, n]$, for all $\mathbf{x}_0, \mathbf{x}_1 \in \mathbb{Z}_q^n$ and for all $\mathbf{u}_1$ in the span of $\mathbf{B}_1^*$, the distributions of

$$(x_{j,b} \cdot \mathbf{u}_1 + \mathbf{u}_{j,1}) \quad \text{and} \quad (x_{j,1-b} \cdot \mathbf{u}_1 + \mathbf{u}_{j,1})$$

are statistically indistinguishable whenever $\mathbf{u}_{j,1}$ is uniform in the span of $\mathbf{B}_1^*$.

Expt-4-ℓ. For each $\ell \in [0, Q]$, the experiment Expt-4-ℓ is identical to Expt-3 except for the manner in which the first ℓ secret key queries are answered by the challenger $\mathcal{B}$. More specifically, $\mathcal{B}$ uniformly samples a basis $(\mathbf{B}_0, \mathbf{B}_1, \mathbf{B}_2)$ with corresponding dual basis $(\mathbf{B}_0^*, \mathbf{B}_1^*, \mathbf{B}_2^*)$, and uses $\mathbf{B}_0$ as part of the master secret key msk. For each $\ell' \in [1, \ell]$, $\mathcal{B}$ uniformly samples $\mathbf{s}_{\ell',0} \xleftarrow{R} \mathbb{Z}_q^k$ and $s_{\ell',1} \xleftarrow{R} \mathbb{Z}_q$, and sets

$$\mathbf{t}_{\ell'} = \left(\mathbf{B}_0 \cdot \mathbf{s}_{\ell',0} + \mathbf{B}_1 \cdot s_{\ell',1}\right)^{\mathbf{T}}.$$

In other words, the vector $(\mathbf{t}_{\ell'})^{\mathbf{T}}$ now lies in the span of $\left[\mathbf{B}_0 \mid \mathbf{B}_1\right]$ and not in the span of $\mathbf{B}_0$, as in the real experiment. The challenger $\mathcal{B}$ then generates the secret key corresponding to the predicate matrix $\mathbf{W}_{\ell'}$ as $\mathsf{sk}_{\mathbf{W}_{\ell'}} = \left(\{h_{j,\ell'}\}_{j \in [0,n]}\right)$ where

$$h_{0,\ell'} = g_2^{\left(\mathbf{K} + \sum_{i=1}^m y_{\ell',i} \cdot \left(\sum_{j=1}^n w_{i,j} \cdot \mathbf{t}_{\ell'} \cdot \mathbf{S}_j\right)\right)^{\mathbf{T}}},$$
$$h_{j,\ell'} = g_2^{\left(\sum_{i=1}^m y_{\ell',i} \cdot w_{i,j} \cdot \mathbf{t}_{\ell'}\right)^{\mathbf{T}}} \quad \text{for } j \in [1, n].$$

where $y_{\ell',1}, \ldots, y_{\ell',m} \xleftarrow{R} \mathbb{Z}_q$.

Let $P_{\mathcal{A},4,\ell}$ denote the probability that $b = b'$, where b' is the bit output by the adversary $\mathcal{A}$ at the end of Expt-4-ℓ. We state the following lemma.

Lemma 4.3 *For all PPT adversaries $\mathcal{A}$, $\left|P_{\mathcal{A},4,\ell} - P_{\mathcal{A},4,(\ell-1)}\right| \leq \mathsf{negl}(\lambda)$ for each $\ell \in [1, Q]$.*

Proof. The proof proceeds through another sequence of hybrid experiments, beginning with an experiment identical to Expt-4-$(\ell-1)$ and ending with an experiment identical to Expt-4-ℓ. Each experiment in this sequence differs from its predecessor in one of two ways: either the ℓ^{th} secret key sk_ℓ is generated in a different manner, or the challenge ciphertext ct^* is generated in a different manner. The corresponding indistinguishability arguments between pairs of successive experiments rely heavily on Lemmas 2.2, 2.3 and 2.4.

Expt-5. This experiment is identical to Expt-4-Q except for the manner in which the challenge ciphertext ct^* is generated. More specifically, the challenger $\mathcal{B}$ samples $\mathbf{u}, \mathbf{u}', \mathbf{u}'' \xleftarrow{R} \mathbb{Z}_q^{k+1}$ and uses the dual basis to decompose these as

$$\begin{aligned}\mathbf{S}_0 \cdot \mathbf{u} &= (\mathbf{u}_0 + \mathbf{u}_1 + \mathbf{u}_2)\\ \mathbf{S}_0 \cdot \mathbf{u}' &= (\mathbf{u}'_0 + \mathbf{u}'_1 + \mathbf{u}'_2)\\ \mathbf{S}_0 \cdot \mathbf{u}'' &= (\mathbf{u}''_0 + \mathbf{u}''_1 + \mathbf{u}''_2)\end{aligned}$$

It then generates the ciphertext $\mathsf{ct}^* = \left(\{c_j\}_{j\in[0,n+1]}\right)$ as

$$c_0 = g_1^{\mathbf{u}^\mathbf{T}}, \quad \left\{c_j = g_1^{\left(\mathbf{v}_j^{(5)}\right)^\mathbf{T}}\right\}_{j\in[1,n]}, \quad c_{n+1} = M_0 \cdot e\left(c_0, g_2\right)^{\mathbf{K}^\mathbf{T}},$$

where for each $j \in [1, n]$, we have

$$\mathbf{v}_j^{(5)} = \boxed{x_{j,0} \cdot (\mathbf{u}'_0 + \mathbf{u}'_1)} + \boxed{x_{j,1} \cdot (\mathbf{u}''_0 + \mathbf{u}''_1)} + x_{j,b} \cdot \mathbf{u}_2 + \mathbf{S}_j \cdot \mathbf{u}.$$

Let $P_{\mathcal{A},5}$ denote the probability that $b = b'$, where b' is the bit output by the adversary $\mathcal{A}$ at the end of Expt-5. We state and prove the following lemma.

Lemma 4.4 *For all unbounded adversaries $\mathcal{A}$, $|P_{\mathcal{A},5} - P_{\mathcal{A},4-Q}| \leq \mathsf{negl}(\lambda)$.*

Proof Sketch. To prove this lemma, we employ the standard "change of basis" technique used in dual pairing vector spaces [31, 33, 34]. More specifically, we argue that the distributions of

$$(\mathbf{u}_1, \mathbf{u}_2) \quad \text{and} \quad ((\mathbf{u}'_0 + \mathbf{u}'_1), (\mathbf{u}''_0 + \mathbf{u}''_1))$$

are statistically indistinguishable whenever the vectors $\mathbf{u}, \mathbf{u}', \mathbf{u}''$ and the basis matrices $\mathbf{B}_0, \mathbf{B}_1$ are uniformly random. Informally, the argument follows from the following observations:

- The randomness $\mathbf{t}_i$ in each secret key $\mathsf{sk}_{\mathbf{W}_i}$ for $i \in [1, Q]$ statistically hides the span of $\left[\mathbf{B}_0 \mid \mathbf{B}_1\right]$. This allows for an alternative simulation of Expt-4, where the basis matrices $\mathbf{B}_0, \mathbf{B}_1$ are "changed", i.e., replaced by two other specially constructed basis matrices, such that the replacement matrices are also distributed uniformly.
- The alternative simulation of Expt-4 is statistically indistinguishable from the original simulation of Expt-4.
- The alternative simulation of Expt-4 with respect to the changed basis matrices is statistically indistinguishable from the simulation of Expt-5 with respect to the original basis matrices.

Expt-6. This experiment is identical to Expt-5 except for the manner in which the challenge ciphertext ct^* is generated. Namely, the challenger $\mathcal{B}$ uniformly samples $\mathbf{u}, \mathbf{u}', \mathbf{u}'' \xleftarrow{R} \mathbb{Z}_q^{k+1}$ and generates the ciphertext $\mathsf{ct}^* = \left(\{c_j\}_{j\in[0,n+1]}\right)$ as

$$c_0 = g_1^{\mathbf{u}^\mathbf{T}}, \quad \left\{c_j = g_1^{\left(\mathbf{v}_j^{(6)}\right)^\mathbf{T}}\right\}_{j\in[1,n]}, \quad c_{n+1} = M_0 \cdot e\left(c_0, g_2\right)^{\mathbf{K}^\mathbf{T}},$$

where for each $j \in [1, n]$, we have

$$\mathbf{v}_j^{(6)} = \boxed{x_{j,0} \cdot \mathbf{S}_0 \cdot \mathbf{u}' + x_{j,1} \cdot \mathbf{S}_0 \cdot \mathbf{u}''} + \mathbf{S}_j \cdot \mathbf{u}.$$

Let $P_{\mathcal{A},6}$ denote the probability that $b = b'$, where b' is the bit output by the adversary $\mathcal{A}$ at the end of Expt-6. We state and prove the following lemma.

Lemma 4.5 *For all unbounded adversaries $\mathcal{A}$, $|P_{\mathcal{A},6} - P_{\mathcal{A},5}| \leq \mathsf{negl}(\lambda)$.*

Proof. The proof is similar to the proof of indistinguishability of Expt 2 and Expt 3.

Finally, observe that in Expt-6, the adversary $\mathcal{A}$ has no advantage in guessing b, since the ciphertext ct^* is entirely independent of b. In other words, for all PPT adversaries $\mathcal{A}$, we must have $P_{\mathcal{A},6} = 1/2$. This completes the proof of Theorem 4.1. □

4.3 Function Privacy

We state and prove the following theorem.

Theorem 4.2 *If the (n, m)-*min-entropy-MDDH *assumption holds over the group $\mathbb{G}_2$, then for all PPT adversaries $\mathcal{A}$, we have $\mathbf{Adv}^{\mathsf{FP}}_{\Pi^{\mathrm{SME}},\mathcal{A}}(\lambda) \leq \mathsf{negl}(\lambda)$.*

Proof. The proof proceeds through a sequence of experiments, beginning with the "real" function privacy experiment and ending with an experiment where the adversary has no advantage.

Expt-0. This is the "real" function privacy experiment. In this experiment, the adversary $\mathcal{A}$ is given the public parameter pp. The adversary chooses two circuits corresponding to matrix distributions of the form

$$\mathcal{W}_0 = \left[\mathrm{W}_{i,j}^{(0)}\right]_{i\in[1,m],j\in[1,n]}, \quad \mathcal{W}_1 = \left[\mathrm{W}_{i,j}^{(1)}\right]_{i\in[1,m],j\in[1,n]},$$

representing joint distributions over $\mathbb{Z}_q^{m\times n}$, subject to the following restrictions:

1. For each $i \in [1, m], j \in [1, n]$ and $\tilde{b} \in \{0, 1\}$, $\mathrm{W}_{i,j}^{(\tilde{b})}$ represents an $\omega(\log \lambda)$-source over $\mathbb{F}_q$.
2. For each $i, i' \in [1, m], j, j' \in [1, n]$ and $\tilde{b} \in \{0, 1\}$, $\mathrm{W}_{i,j}^{(\tilde{b})}$ and $\mathrm{W}_{i',j'}^{(\tilde{b})}$ represent mutually independent distributions.

The adversary $\mathcal{A}$ also (adaptively) issues key generation queries corresponding to predicate matrices the form $\mathbf{W}_1, \ldots, \mathbf{W}_Q \in \mathbb{Z}_q^{m \times n}$ for some $Q = \mathsf{poly}(\lambda)$. The challenger samples $\mathbf{W}^* \xleftarrow{R} \mathcal{W}_b$ for some random $b \xleftarrow{R} \{0,1\}$, where

$$\mathbf{W}^* = \left[w^*_{i,j}\right]_{i \in [1,m], j \in [1,n]},$$

and uses the master secret key $\mathsf{msk} = (g_2, \mathbf{S}_0, \mathbf{S}_1, \ldots, \mathbf{S}_n, \mathbf{K}, \mathbf{B}_0)$ to set the challenge secret key $\mathsf{sk}_{\mathbf{W}^*} = \left(\{h_j\}_{j \in [0,n]}\right)$ where

$$h_0 = g_2^{\left(\mathbf{K} + \sum_{i=1}^m y_i \cdot \left(\sum_{j=1}^n w^*_{i,j} \cdot \mathbf{t} \cdot \mathbf{S}_j\right)\right)^{\mathbf{T}}},$$
$$h_j = g_2^{\left(\sum_{i=1}^m y_i \cdot w^*_{i,j} \cdot \mathbf{t}\right)^{\mathbf{T}}} \quad \text{for } j \in [1,n],$$

where $y_1, \ldots, y_m \xleftarrow{R} \mathbb{Z}_q$ and $\mathbf{t} = (\mathbf{B} \cdot \mathbf{s})^{\mathbf{T}}$ for some $\mathbf{s} \xleftarrow{R} \mathbb{Z}_q^k$. The adversary $\mathcal{A}$ receives $\left(\mathsf{sk}_{\mathbf{W}^*}, \mathsf{sk}_{\mathbf{W}_1}, \ldots, \mathsf{sk}_{\mathbf{W}_Q}\right)$, where

$$\mathsf{sk}_{\mathbf{W}_\ell} \leftarrow \mathsf{KeyGen}(\mathsf{pp}, \mathsf{msk}, \mathbf{W}_\ell) \quad \text{for each } \ell \in [1, Q].$$

Finally, it outputs a bit b'. Let $P_{\mathcal{A},0}$ denote the probability that $b = b'$.

Expt-1. This experiment is identical to Expt-0 except for the manner in which the challenge secret key $\mathsf{sk}_{\mathbf{W}^*}$ is generated. Namely, the challenger $\mathcal{B}$ uniformly samples $u_1, \ldots, u_n \xleftarrow{R} \mathbb{Z}_q$ and sets the challenge secret key $\mathsf{sk}_{\mathbf{W}^*} = \left(\{h_j\}_{j \in [0,n]}\right)$ as follows:

$$h_0 = g_2^{\left(\mathbf{K} + \sum_{j=1}^n u_j \cdot \mathbf{t} \cdot \mathbf{S}_j\right)^{\mathbf{T}}},$$
$$h_j = g_2^{(u_j \mathbf{t})^{\mathbf{T}}} \quad \text{for } j \in [1,n],$$

where $\mathbf{t} = (\mathbf{B} \cdot \mathbf{s})^{\mathbf{T}}$ for some $\mathbf{s} \xleftarrow{R} \mathbb{Z}_q^k$. Let $P_{\mathcal{A},1}$ denote the probability that $b = b'$, where b' is the bit output by the adversary $\mathcal{A}$ at the end of Expt-1. By the (n,m)-min-entropy-MDDH assumption, we must have $|P_{\mathcal{A},2} - P_{\mathcal{A},1}| \leq \mathsf{negl}(\lambda)$.

Finally, observe that the challenge secret key $\mathsf{sk}_{\mathbf{W}^*}$ in Expt-1 is independent of the bit b chosen by the challenger. Hence, for all PPT adversaries $\mathcal{A}$, we must have $P_{\mathcal{A},1} = 1/2$. This completes the proof of Theorem 4.2. □

5 Function Private SNME

In this section, we present an SNME scheme that is (computationally) function private whenever the predicate matrices are sampled from distributions with super-logarithmic min-entropy. Similar to the SME scheme, attribute privacy of our SNME scheme can be based on the $\mathcal{U}_{k+1,k}$-MDDH assumption, albeit in the bounded collusion setting, while function privacy follows from the min-entropy-MDDH assumption described in Section 2.4.

5.1 The Construction

Let $\mathsf{GroupGen}(1^\lambda)$ be a PPT algorithm that takes as input a security parameter $\lambda \in \mathbb{N}$, and outputs the tuple $(\mathbb{G}_1, \mathbb{G}_2, \mathbb{G}_T, q, g_1, g_2, e)$, where $\mathbb{G}_1$, $\mathbb{G}_2$ and $\mathbb{G}_T$ are cyclic groups of prime order q (q being a λ-bit prime), g_1 is a generator for $\mathbb{G}_1$, g_2 is a generator for $\mathbb{G}_2$, and $e : \mathbb{G}_1 \times \mathbb{G}_2 \longrightarrow \mathbb{G}_T$ is an efficiently computable non-degenerate asymmetric bilinear map. Our scheme Π^{SNME} is parameterized by $m, n = \mathsf{poly}(\lambda)$ in the sense that it supports predicate matrices of the form $\mathbf{W} \in \mathbb{Z}_q^{m \times n}$, and attribute vectors of the form $\mathbf{x} \in \mathbb{Z}_q^n$. The payload message space $\mathcal{M}$ for this scheme is assumed to be a "small" subset of $\mathbb{Z}_q$ such that $|\mathcal{M}| \leq \mathsf{poly}(\lambda)$.

- $\mathsf{Setup}(1^\lambda)$: Uniformly sample $(\mathbb{G}_1, \mathbb{G}_2, \mathbb{G}_T, q, g_1, g_2, e) \leftarrow \mathsf{GroupGen}(1^\lambda)$. Also, uniformly sample $\mathbf{A} \xleftarrow{R} \mathbb{Z}_q^{(k+1)\times k}$ and $\mathbf{S}_1, \mathbf{S}_2 \xleftarrow{R} \mathbb{Z}_q^{n\times(k+1)}$ for some constant $k > 0$. Output

$$\mathsf{pp} = \left(g_1, g_1^{\mathbf{A}}, g_1^{\mathbf{S}_1 \cdot \mathbf{A}}, g_1^{\mathbf{S}_2 \cdot \mathbf{A}}\right), \quad \mathsf{msk} = \left(g_2, \mathbf{S}_1, \mathbf{S}_2\right).$$

- $\mathsf{KeyGen}(\mathsf{pp}, \mathsf{msk}, \mathbf{W})$: Given a predicate matrix $\mathbf{W} \in \mathbb{Z}_q^{m\times n}$, sample $\mathbf{y} \xleftarrow{R} \mathbb{Z}_q^m$ and output $\mathsf{sk}_{\mathbf{W}} = (h_0, h_1, h_2)$, where

$$h_0 = g_2^{\mathbf{W}^\mathbf{T} \cdot \mathbf{y}}, \quad h_1 = g_2^{(\mathbf{W} \cdot \mathbf{S}_1)^\mathbf{T} \cdot \mathbf{y}}, \quad h_2 = g_2^{(\mathbf{W} \cdot \mathbf{S}_2)^\mathbf{T} \cdot \mathbf{y}}.$$

- $\mathsf{Enc}(\mathsf{pp}, \mathbf{x}, M)$: Given an attribute vector $\mathbf{x} \in \mathbb{Z}_q^n$ and a message $M \in \mathcal{M} \subset \mathbb{Z}_q$, uniformly sample $\mathbf{r}_1, \mathbf{r}_2 \xleftarrow{R} \mathbb{Z}_q^k$ and output $\mathsf{ct} = (c_{1,0}, c_{1,1}, c_{2,0}, c_{2,1})$ where

$$\begin{aligned} c_{1,0} &= g_1^{(\mathbf{A}\cdot\mathbf{r}_1)^\mathbf{T}}, \quad c_{1,1} = g_1^{(\mathbf{x}+\mathbf{S}_1\cdot\mathbf{A}\cdot\mathbf{r}_1)^\mathbf{T}}, \\ c_{2,0} &= g_1^{(\mathbf{A}\cdot\mathbf{r}_2)^\mathbf{T}}, \quad c_{2,1} = g_1^{(M\cdot\mathbf{x}+\mathbf{S}_2\cdot\mathbf{A}\cdot\mathbf{r}_2)^\mathbf{T}}. \end{aligned}$$

- $\mathsf{Dec}(\mathsf{pp}, \mathsf{sk}_{\mathbf{W}}, \mathsf{ct})$: Parse the ciphertext as $\mathsf{ct} = (c_{1,0}, c_{1,1}, c_{2,0}, c_{2,1})$ and the secret key as $\mathsf{sk}_{\mathbf{W}} = (h_0, h_1, h_2)$. Check if there exists a *unique* $M \in \mathcal{M}$ such that

$$e\left(c_{2,1}, h_0\right) \cdot e\left(c_{2,0}, h_2\right)^{-1} = \left(e\left(c_{1,1}, h_0\right) \cdot e\left(c_{1,0}, h_1\right)^{-1}\right)^M.$$

  If yes, return M. Else return $\perp$.

Correctness. To see that the aforementioned scheme is functionally correct, observe the following.

$$\begin{aligned} e\left(c_{1,1}, h_0\right) \cdot e\left(c_{1,0}, h_1\right)^{-1} &= e\left(g_1, g_2\right)^{\left(\mathbf{y}^\mathbf{T}\cdot\mathbf{W}\cdot(\mathbf{x}+\mathbf{S}_1\cdot\mathbf{A}\cdot\mathbf{r}_1) - \mathbf{y}^\mathbf{T}\cdot\mathbf{W}\cdot\mathbf{S}_1\cdot\mathbf{A}\cdot\mathbf{r}_1\right)^\mathbf{T}} \\ &= e\left(g_1, g_2\right)^{\left(\mathbf{y}^\mathbf{T}\cdot\mathbf{W}\cdot\mathbf{x}\right)^\mathbf{T}} \\ e\left(c_{2,1}, h_0\right) \cdot e\left(c_{2,0}, h_2\right)^{-1} &= e\left(g_1, g_2\right)^{\left(\mathbf{y}^\mathbf{T}\cdot\mathbf{W}\cdot(M\cdot\mathbf{x}+\mathbf{S}_2\cdot\mathbf{A}\cdot\mathbf{r}_2) - \mathbf{y}^\mathbf{T}\cdot\mathbf{W}\cdot\mathbf{S}_2\cdot\mathbf{A}\cdot\mathbf{r}_2\right)^\mathbf{T}} \\ &= e\left(g_1, g_2\right)^{M\cdot\left(\mathbf{y}^\mathbf{T}\cdot\mathbf{W}\cdot\mathbf{x}\right)^\mathbf{T}} \end{aligned}$$

When $\mathbf{W} \cdot \mathbf{x} \neq \mathbf{0} \mod q$, we have $\mathbf{y}^{\mathbf{T}} \cdot \mathbf{W} \cdot \mathbf{x} \neq 0 \mod q$ with overwhelmingly large probability over the randomness of KeyGen, and the decryption algorithm correctly recovers the message M. But when $\mathbf{W} \cdot \mathbf{x} = \mathbf{0} \mod q$, the message M cannot be uniquely recovered and the decryption algorithm returns $\perp$.

5.2 Attribute Privacy

We state the following theorem.

Theorem 5.1 *If the $\mathcal{U}_{k+1,k}$-MDDH assumption holds over the group $\mathbb{G}_1$, then for all PPT adversaries $\mathcal{A}$ that issue as most $(n-1)$ secret key queries during the attribute privacy experiment, we have $\mathbf{Adv}^{\mathsf{AP}}_{\Pi^{\mathrm{SNME}},\mathcal{A}}(\lambda) \leq \mathsf{negl}(\lambda)$.*

Proof Sketch. Due to lack of space, we only provide a brief proof sketch. We refer the reader to the full version of the paper [**?**] for the detailed proof.

The proof essentially relies on hash proof systems [18, 19], and uses arguments similar to those used by Agrawal et al. in proving the security of their linear FE scheme [6]. The analysis exploits the following fact: given the public parameter pp and no more than $(n-1)$ secret keys, the master secret key components $\mathbf{S}_0, \mathbf{S}_1 \ldots, \mathbf{S}_n$ retain sufficient entropy from an (unbounded) adversary's point of view. This in turn ensures that at some stage, if the challenge ciphertext is generated using the master-secret-key instead of the public parameter, it will perfectly hide which attribute-message pair among $(\mathbf{x}_0, M_0)$ and $(\mathbf{x}_1, M_1)$ is encrypted.

Finally, the scheme is adaptively secure because the reduction knows the master secret key at any time, which allows it to answer all secret key queries without knowing the challenge attributes beforehand. This feature is common to nearly all security proofs relying on hash proof systems [18, 19].

5.3 Function Privacy

We state and prove the following theorem.

Theorem 5.2 *If the (n, m)-min-entropy-MDDH assumption holds over the group $\mathbb{G}_2$, then for all PPT adversaries $\mathcal{A}$, we have $\mathbf{Adv}^{\mathsf{FP}}_{\Pi^{\mathrm{SNME}},\mathcal{A}}(\lambda) \leq \mathsf{negl}(\lambda)$.*

Proof. The proof proceeds through a sequence of experiments, beginning with the "real" function privacy experiment and ending with an experiment where the adversary has no advantage.

Expt-0. This is the "real" function privacy experiment. In this experiment, the adversary $\mathcal{A}$ is given the public parameter pp. The adversary chooses two circuits corresponding to matrix distributions of the form

$$\mathcal{W}_0 = \left[\mathrm{W}^{(0)}_{i,j}\right]_{i\in[1,m],j\in[1,n]}, \quad \mathcal{W}_1 = \left[\mathrm{W}^{(1)}_{i,j}\right]_{i\in[1,m],j\in[1,n]},$$

representing joint distributions over $\mathbb{Z}_q^{m\times n}$, subject to the following restrictions:

1. For each $i \in [1,m], j \in [1,n]$ and $\tilde{b} \in \{0,1\}$, $\mathrm{W}_{i,j}^{(\tilde{b})}$ represents an $\omega(\log \lambda)$-source over $\mathbb{F}_q$.
2. For each $i, i' \in [1,m], j, j' \in [1,n]$ and $\tilde{b} \in \{0,1\}$, $\mathrm{W}_{i,j}^{(\tilde{b})}$ and $\mathrm{W}_{i',j'}^{(\tilde{b})}$ represent mutually independent distributions.

The adversary $\mathcal{A}$ also (adaptively) issues key generation queries corresponding to predicate matrices the form $\mathbf{W}_1, \ldots, \mathbf{W}_Q \in \mathbb{Z}_q^{m \times n}$ for some $Q = \mathsf{poly}(\lambda)$. The challenger samples $\mathbf{W}^* \xleftarrow{R} \mathcal{W}_b$ for some random $b \xleftarrow{R} \{0,1\}$, and and uses the master secret key $\mathsf{msk} = \left(\mathbf{S}_1, \mathbf{S}_2\right)$ to set $\mathsf{sk}_{\mathbf{W}^*} = (h_0, h_1, h_2)$, where

$$h_0 = g_2^{(\mathbf{W}^*)^{\mathbf{T}} \cdot \mathbf{y}}, \quad h_1 = g_2^{(\mathbf{W} \cdot \mathbf{S}_1^*)^{\mathbf{T}} \cdot \mathbf{y}}, \quad h_2 = g_2^{(\mathbf{W} \cdot \mathbf{S}_2^*)^{\mathbf{T}} \cdot \mathbf{y}},$$

where $\mathbf{y} \xleftarrow{R} \mathbb{Z}_q^m$. The adversary $\mathcal{A}$ receives $\left(\mathsf{sk}_{\mathbf{W}^*}, \mathsf{sk}_{\mathbf{W}_1}, \ldots, \mathsf{sk}_{\mathbf{W}_Q}\right)$, where

$$\mathsf{sk}_{\mathbf{W}_\ell} \leftarrow \mathsf{KeyGen}(\mathsf{pp}, \mathsf{msk}, \mathbf{W}_\ell) \quad \text{for each } \ell \in [1, Q].$$

Finally, it outputs a bit b'. Let $P_{\mathcal{A},0}$ denote the probability that $b = b'$.

Expt-1. This experiment is identical to Expt-0 except for the manner in which the challenge secret key $\mathsf{sk}_{\mathbf{W}^*}$ is generated. Namely, the challenger $\mathcal{B}$ uniformly samples $\mathbf{u} \xleftarrow{R} \mathbb{Z}_q^n$ and uses the master secret key $\mathsf{msk} = (g_2, \mathbf{S}_1, \mathbf{S}_2)$ to output $\mathsf{sk}_{\mathbf{W}^*} = (h_0, h_1, h_2)$ where

$$h_0 = g_2^{\mathbf{u}}, \quad h_1 = g_2^{(\mathbf{S}_1)^{\mathbf{T}} \cdot \mathbf{u}}, \quad h_2 = g_2^{(\mathbf{S}_2)^{\mathbf{T}} \cdot \mathbf{u}}.$$

Let $P_{\mathcal{A},1}$ denote the probability that $b = b'$, where b' is the bit output by the adversary $\mathcal{A}$ at the end of Expt-1. By the (n,m)-min-entropy-MDDH assumption, we must have $|P_{\mathcal{A},2} - P_{\mathcal{A},1}| \leq \mathsf{negl}(\lambda)$.

Finally, observe that the challenge secret key $\mathsf{sk}_{\mathbf{W}^*}$ in Expt-1 is independent of the bit b chosen by the challenger. Hence, for all PPT adversaries $\mathcal{A}$, we must have $P_{\mathcal{A},1} = 1/2$. This completes the proof of Theorem 5.2. □

6 Acknowledgments

We thank the anonymous reviewers of PKC 2019 for useful comments. Patranabis and Mukhopadhyay are patially supported by Qualcomm India Innovation Fellowship grant. Mukhopadhyay is partially supported by a DST India Swarnajayanti Fellowship. Ramanna is partially supported by DST India Inspire Faculty award. We stress that the opinions, findings and conclusions expressed in this material are those of the authors and do not necessarily reflect the views of the funding organizations.

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. *J. Cryptology*, pages 350–391, 2008.
2. Michel Abdalla, Florian Bourse, Angelo De Caro, and David Pointcheval. Simple functional encryption schemes for inner products. In *PKC 2015*, pages 733–751, 2015.
3. Shashank Agrawal, Shweta Agrawal, Saikrishna Badrinarayanan, Abishek Kumarasubramanian, Manoj Prabhakaran, and Amit Sahai. On the practical security of inner product functional encryption. In *PKC 2015*, pages 777–798, 2015.
4. Shweta Agrawal, Sanjay Bhattacherjee, Duong Hieu Phan, Damien Stehlé, and Shota Yamada. Efficient public trace and revoke from standard assumptions: Extended abstract. In *CCS 2017*, pages 2277–2293, 2017.
5. Shweta Agrawal, David Mandell Freeman, and Vinod Vaikuntanathan. Functional Encryption for Inner Product Predicates from Learning with Errors. In *ASIACRYPT 2011*, pages 21–40, 2011.
6. Shweta Agrawal, Benoît Libert, and Damien Stehlé. Fully secure functional encryption for inner products, from standard assumptions. In *CRYPTO 2016*, pages 333–362, 2016.
7. Nuttapong Attrapadung and Benoît Libert. Functional encryption for inner product: Achieving constant-size ciphertexts with adaptive security or support for negation. In *PKC 2010*, pages 384–402, 2010.
8. Allison Bishop, Abhishek Jain, and Lucas Kowalczyk. Function-hiding inner product encryption. In *ASIACRYPT 2015*, pages 470–491, 2015.
9. Dan Boneh, Giovanni Di Crescenzo, Rafail Ostrovsky, and Giuseppe Persiano. Public Key Encryption with Keyword Search. In *EUROCRYPT 2004*, pages 506–522, 2004.
10. Dan Boneh and Matthew K. Franklin. Identity-based encryption from the weil pairing. In *CRYPTO 2001*, pages 213–229, 2001.
11. Dan Boneh, Craig Gentry, Sergey Gorbunov, Shai Halevi, Valeria Nikolaenko, Gil Segev, Vinod Vaikuntanathan, and Dhinakaran Vinayagamurthy. Fully key-homomorphic encryption, arithmetic circuit ABE and compact garbled circuits. In *EUROCRYPT 2014*, pages 533–556, 2014.
12. Dan Boneh, Ananth Raghunathan, and Gil Segev. Function-Private Identity-Based Encryption: Hiding the Function in Functional Encryption. In *CRYPTO 2013*, pages 461–478, 2013.
13. Dan Boneh, Ananth Raghunathan, and Gil Segev. Function-Private Subspace-Membership Encryption and Its Applications. In *ASIACRYPT 2013*, pages 255–275, 2013.
14. Dan Boneh and Brent Waters. Conjunctive, Subset, and Range Queries on Encrypted Data. In *TCC 2007*, pages 535–554, 2007.
15. Zvika Brakerski and Gil Segev. Function-Private Functional Encryption in the Private-Key Setting. In *TCC 2015*, pages 306–324, 2015.
16. Jie Chen, Romain Gay, and Hoeteck Wee. Improved dual system ABE in prime-order groups via predicate encodings. In *EUROCRYPT 2015*, pages 595–624, 2015.
17. Jie Chen, Junqing Gong, Lucas Kowalczyk, and Hoeteck Wee. Unbounded ABE via bilinear entropy expansion, revisited. In *EUROCRYPT 2018*, pages 503–534, 2018.

18. Ronald Cramer and Victor Shoup. A practical public key cryptosystem provably secure against adaptive chosen ciphertext attack. In *CRYPTO '98*, pages 13–25, 1998.
19. Ronald Cramer and Victor Shoup. Universal hash proofs and a paradigm for adaptive chosen ciphertext secure public-key encryption. In *EUROCRYPT 2002*, pages 45–64, 2002.
20. Pratish Datta, Ratna Dutta, and Sourav Mukhopadhyay. Functional encryption for inner product with full function privacy. In *PKC 2016*, pages 164–195, 2016.
21. Alex Escala, Gottfried Herold, Eike Kiltz, Carla Ràfols, and Jorge L. Villar. An algebraic framework for diffie-hellman assumptions. In *CRYPTO 2013*, pages 129–147, 2013.
22. Sanjam Garg, Craig Gentry, Shai Halevi, Amit Sahai, and Brent Waters. Attribute-based encryption for circuits from multilinear maps. In *CRYPTO 2013*, pages 479–499, 2013.
23. Romain Gay, Dennis Hofheinz, Eike Kiltz, and Hoeteck Wee. Tightly cca-secure encryption without pairings. In *EUROCRYPT 2016*, pages 1–27, 2016.
24. Craig Gentry, Chris Peikert, and Vinod Vaikuntanathan. Trapdoors for hard lattices and new cryptographic constructions. In *ACM STOC 2008*, pages 197–206, 2008.
25. Shafi Goldwasser, Yael Tauman Kalai, Raluca A. Popa, Vinod Vaikuntanathan, and Nickolai Zeldovich. How to run turing machines on encrypted data. In *CRYPTO 2013*, pages 536–553, 2013.
26. Junqing Gong, Xiaolei Dong, Jie Chen, and Zhenfu Cao. Efficient IBE with tight reduction to standard assumption in the multi-challenge setting. In *ASIACRYPT 2016*, pages 624–654, 2016.
27. Sergey Gorbunov, Vinod Vaikuntanathan, and Hoeteck Wee. Attribute-based encryption for circuits. *J. ACM*, 62(6):45:1–45:33, 2015.
28. Sergey Gorbunov, Vinod Vaikuntanathan, and Hoeteck Wee. Predicate encryption for circuits from LWE. In *CRYPTO 2015*, pages 503–523, 2015.
29. Vipul Goyal, Omkant Pandey, Amit Sahai, and Brent Waters. Attribute-based encryption for fine-grained access control of encrypted data. In *ACM CCS 2006*, pages 89–98, 2006.
30. Jonathan Katz, Amit Sahai, and Brent Waters. Predicate Encryption Supporting Disjunctions, Polynomial Equations, and Inner Products. *J. Cryptology*, 26(2):191–224, 2013.
31. Allison B. Lewko. Tools for simulating features of composite order bilinear groups in the prime order setting. In *EUROCRYPT 2012*, pages 318–335, 2012.
32. Allison B. Lewko and Brent Waters. New techniques for dual system encryption and fully secure HIBE with short ciphertexts. In *TCC 2010*, pages 455–479, 2010.
33. Tatsuaki Okamoto and Katsuyuki Takashima. Adaptively attribute-hiding (hierarchical) inner product encryption. In *EUROCRYPT 2012*, pages 591–608, 2012.
34. Tatsuaki Okamoto and Katsuyuki Takashima. Fully secure unbounded inner-product and attribute-based encryption. In *ASIACRYPT 2012*, pages 349–366, 2012.
35. Sikhar Patranabis, Debdeep Mukhopadhyay, and Somindu C. Ramanna. Function private predicate encryption for low min-entropy predicates. *IACR Cryptology ePrint Archive*, page 1250, 2018.
36. Brent Waters. Functional encryption for regular languages. In *CRYPTO 2012*, pages 218–235, 2012.
37. Hoeteck Wee. Attribute-hiding predicate encryption in bilinear groups, revisited. In *TCC 2017*, pages 206–233, 2017.