

Multi-Party Functional Encryption

Shweta Agrawal
IIT Madras*

Rishab Goyal
MIT[†]

Junichi Tomida
NTT Corporation[‡]

Abstract

We initiate the study of *multi-party functional encryption* (MPFE) which unifies and abstracts out various notions of functional encryption which support distributed ciphertexts or secret keys, such as multi-input FE, multi-client FE, decentralized multi-client FE, multi-authority FE, dynamic decentralized FE, adhoc multi-input FE and such others. Using our framework, we identify several gaps in the literature and provide some constructions to fill these:

1. **Multi-Authority ABE with Inner Product Computation.** The recent work of Abdalla et al. (ASIACRYPT’20) constructed a novel “composition” of Attribute Based Encryption (ABE) and Inner Product Functional Encryption (IPFE), namely functional encryption schemes that combine the access control functionality of attribute based encryption with the possibility of performing linear operations on the encrypted data. In this work, we extend the access control component to support the much more challenging multi-authority setting, i.e. “lift” the primitive of ABE in their construction to multi-authority ABE for the same class of access control policies (LSSS structures). This yields the first construction of a nontrivial multi-authority FE beyond ABE from simple assumptions on pairings to the best of our knowledge.

Our techniques can also be used to generalize the decentralized attribute based encryption scheme of Michalevsky and Joye (ESORICS’18) to support inner product computation on the message. While this scheme only supports inner product predicates which is less general than those supported by the Lewko-Waters (EUROCRYPT’11) construction, it supports policy hiding which the latter does not. Our extension inherits these features and is secure based on the k -linear assumption, in the random oracle model.

2. **Function Hiding DDFE.** The novel primitive of *dynamic* decentralized functional encryption (DDFE) was recently introduced by Chotard et al. (CRYPTO’20), where they also provided the first construction for inner products. However, the primitive of DDFE does not support function hiding, which is a significant limitation for several applications. In this work, we provide a new construction for inner product DDFE which supports function hiding. To achieve our final result, we define and construct the first function hiding *multi-client* functional encryption (MCFE) scheme for inner products, which may be of independent interest.
3. **Distributed Ciphertext-Policy ABE.** We provide a distributed variant of the recent ciphertext-policy attribute based encryption scheme, constructed by Agrawal and Yamada (EUROCRYPT’20). Our construction supports \mathbf{NC}^1 access policies, and is secure based on “Learning With Errors” and relies on the generic bilinear group model as well as the random oracle model.

Our new MPFE abstraction predicts meaningful new variants of functional encryption as useful targets for future work.

*Email: shweta.a@cse.iitm.ac.in. Work done in part while at the Simons Institute for the Theory of Computing. Also supported by the Swarnajayanti Fellowship, an Indo-French CEFIPRA grant and the CCD Centre of Excellence.

[†]Email: goyal@utexas.edu. Research supported in part by NSF CNS Award #1718161, an IBM-MIT grant, and by the Defense Advanced Research Projects Agency (DARPA) under Contract No. HR00112020023. Any opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the United States Government or DARPA. Work done in part while at the Simons Institute for the Theory of Computing, supported by Simons-Berkeley research fellowship.

[‡]Email: junichi.tomida.vw@hco.ntt.co.jp.

Contents

1	Introduction	1
1.1	Unifying the View: Multi-Party Functional Encryption	2
1.2	Comparison to Prior Work	3
1.3	New Constructions	4
1.4	Technical Overview	5
1.5	Predicting New and Useful Primitives via MPFE	10
2	Preliminaries	10
2.1	Access Structures and Linear Secret-Sharing Schemes	11
2.2	Bilinear Map Preliminaries	11
2.3	Basic Tools	14
3	Multi-Party Functional Encryption	15
3.1	Dynamic Multi-Party Functional Encryption	17
4	Capturing Existing Primitives as MPFE	19
4.1	Feasibility of MPFE for General Circuits	21
5	Multi-Authority ABE \circ IPFE for LSSS Access Structures	22
5.1	Construction	23
5.2	Correctness and Security	24
6	Function-Hiding DDFE for Inner Products	31
6.1	Construction of Function-Hiding IP-MCFE	34
6.2	Construction of Function-Hiding IP-DDFE	37
7	Distributed Ciphertext Policy ABE	41
7.1	Definition	41
7.2	Preliminaries	42
7.2.1	Attribute Based Encryption	42
7.2.2	Lattice Preliminaries	44
7.2.3	KP-ABE Scheme by Boneh et al. [BGG ⁺ 14].	45
7.3	Construction	47
8	Decentralized ABE \circ IPFE with Policy Hiding	54
8.1	Definition	54
8.2	Construction	55
8.3	Correctness and Security	56
9	New Primitives Predicted by MPFE	60
A	Feasibility for MPFE for General Circuits	64
A.1	Definitions	64
A.1.1	Multi-Input Functional Encryption	64
A.1.2	Functional Encryption	66
A.1.3	Two-Round MPC	66
A.2	Constructions	68

1 Introduction

Functional encryption (FE) [SW05, BSW11] is a powerful generalization of public key encryption which enables a user to learn a function of the encrypted data. Concretely, in FE, a secret key SK_f is associated with a function f and the ciphertext CT_x is associated with a message x (in the domain of f). And, by combining SK_f with CT_x , the decryptor learns $f(x)$ and nothing else.

The original motivation behind the concept of functional encryption, as discussed in [BSW11], was to put forth a *new broad vision of encryption systems*. Since its introduction, the concept of FE has been massively impactful in several aspects: (i) it helped unify the existing literature on encryption systems (such as identity-based encryption [Sha84, BF01], attribute-based encryption [SW05, GPSW06], predicate encryption [KSW08] and more) and place them under a single umbrella which enabled clear comparisons, (ii) it helped in predicting new natural encryption primitives that had not been studied before, such as partially hiding predicate/functional encryption [GVW15], and (iii) it served as the right abstraction to understand the relationship of this broad concept with other notions in cryptography, such as to indistinguishability obfuscation [AJ15, BV15].

Supporting Multiple Users. Subsequently, many new primitives arose to generalize FE to the multi-user setting – multi-input functional encryption [GGG+14], multi-client functional encryption [CDSG+18a], decentralized multi-client functional encryption, adhoc multi-input functional encryption [ACF+20], multi-authority attribute based encryption [Cha07], dynamic decentralized functional encryption [CDSG+20] and such others. Similar to the many special cases of functional encryption, these notions are related yet different and it is often difficult to understand how they compare to one-another, whether they use related techniques, and what is known in terms of feasibility. Moreover, each new variant that springs up acquires a different name, leading to a plethora of acronyms which clutter the landscape, often adding to confusion rather than clarity.

In this work, we initiate the study of “Multi-Party Functional Encryption” (MPFE) which unifies and abstracts out various notions of multi-user functional encryption, such as those described above. Our starting point is the observation that all above notions of FE support some form of distributed ciphertexts or distributed keys or both. In more detail, we summarize the state of affairs as:

1. **Distributed Ciphertexts.** The primitives of multi-input functional encryption (MIFE) [GGG+14] and multi-client functional encryption (MCFE) [CDSG+18a] generalize FE to support distributed inputs. Both notions permit different parties P_1, \dots, P_n each with inputs x_1, \dots, x_n to compute joint functions on their data, namely $f(x_1, \dots, x_n)$. Each party encrypts its input x_i to obtain CT_i , a key authority holding a master secret MSK generates a functional key SK_f and these enable the decryptor to compute $f(x_1, \dots, x_n)$.

The main difference between these definitions lies in the way the inputs can be combined. In multi-*client* functional encryption (MCFE), inputs x_i are additionally associated with public “labels” lab_i and inputs can only be combined with other inputs that share the same label. On the other hand, multi-input functional encryption does not restrict the way that inputs are combined and permits all possible combinations of inputs. Both primitives are defined as *key policy* systems – namely, the access control policy or function is embedded in the secret key rather than the ciphertext.

2. **Distributed Keys.** Distribution or decentralization of keys in the context of FE has also been considered in various works, to achieve two primary objectives (not necessarily simultaneously) – a) handling the *key escrow* problem, so that there is no single entity in the system that holds a powerful master secret against which no security can hold, and b) *better fitting real world* scenarios where different authorities may be responsible for issuing keys corresponding to different attributes of a user, such as offices for passport, drivers license and such others. We summarize some relevant primitives next.

- (a) *Decentralized Attribute Based Encryption with Policy Hiding* (DABE): A decentralized policy-hiding ABE, denoted by DABE [MJ18] was proposed by Michalevsky and Joye to handle the

key escrow problem. In a DABE scheme, there are n key authorities, each of which run a local setup to generate their private and public keys. An encryptor encrypts a message m along with a general access structure C , while secret keys corresponding to (the same) attribute \mathbf{x} are issued by independent authorities. Decryption recovers m if $C(\mathbf{x}) = 1$. The access policy in the ciphertext is hidden.

- (b) *Multi-Authority Functional Encryption* (MAFE): The notion of Multi-Authority FE/ABE [Cha07, LW11, BCG⁺17] emerged to address the second objective, i.e. handling the case where different authorities are responsible for different sets of attributes. Since ABE is a special case of FE, we focus on MAFE. A MAFE scheme is defined as a ciphertext-policy scheme, namely the policy/function is embedded in the ciphertext as against the function keys. In MAFE, n key authorities may independently generate their private and public keys, without any interaction. An encryptor computes a ciphertext for a message m along with a policy f over the various authorities. Any authority i , can generate a token for a user P for attributes lab_i . A decryptor with tokens for lab_i from authority $i \in [n]$, can decrypt the ciphertext to recover $f(\text{lab}_1, \dots, \text{lab}_n, m)$.

3. **Distributed Ciphertexts and Keys.** Some primitives allow to distribute both ciphertexts and keys. Some examples below.

- (a) *Decentralized Multi-Client Functional Encryption* (D-MCFE): The notion of decentralized multi-client FE was defined by Chotard et al. [CDSG⁺18a, ABKW19, LT⁺19] in order to handle the key escrow problem in an MCFE scheme. D-MCFE is defined as a key policy primitive, and adapts MCFE as described above to ensure that there is no single master secret held by any entity – the parties participate in an interactive setup protocol to establish their individual (correlated) master secret keys. In more detail, there are n parties, each holding MSK_i for $i \in [n]$, that compute ciphertexts for their inputs $(\text{lab}_i, \mathbf{x}_i)$ as well as generate partial decryption keys $\text{SK}_{i,f}$ for a given function f . The decryptor can combine the partial secret keys and individual ciphertexts to compute $f(\mathbf{x}_1, \dots, \mathbf{x}_n)$ if and only if all the labels are equal.
- (b) *Ad Hoc MIFE* (aMIFE): Similar to D-MCFE, this notion was introduced in [ACF⁺20] to handle the key escrow problem in MIFE. This notion is key policy, and offers some additional features as compared to D-MCFE — non-interactive setup and dynamic choice of function arity as well as parties that participate in a computation. This notion does not differentiate between key authorities and users, and lets users generate their own partial decryption keys along with ciphertexts. Thus, for $i \in [n]$, party i computes a ciphertext for \mathbf{x}_i and partial key $\text{SK}_{f,i}$ which can be combined by the decryptor to obtain $f(\mathbf{x}_1, \dots, \mathbf{x}_n)$.
- (c) *Dynamic Decentralized FE* (DDFE): This primitive was introduced very recently in [CDSG⁺20] to further generalize aMIFE – it requires non-interactive, local setup and allows dynamic choice of function arity as in aMIFE, but additionally allows partial decryption keys provided by users to be combined in more general ways than in aMIFE. Also, unlike aMIFE, it supports the public key setting.

1.1 Unifying the View: Multi-Party Functional Encryption

While the above notions enable controlled manipulation of encrypted data in increasingly expressive ways, they are too related to warrant independent identities. To unify and extend the above primitives, we propose the notion of multi-party functional encryption (MPFE). All the above examples (and more) can be cast as examples of MPFE with a suitable choice of parameters: this clarifies the connections between these primitives. MPFE allows for both distributed ciphertexts and distributed keys, and specifies how these may be combined for function evaluation. To avoid bifurcating key-policy and ciphertext-policy schemes, we allow either ciphertext or key inputs to encode functions. To better capture attribute and function hiding, we allow every message or function being encoded to have a public and private part. To support schemes with

interactive, independent or centralized setup, we allow the setup algorithm of MPFE to function in any of these modes.

A bit more formally, let n_x be the number of ciphertext inputs and n_y be the number of key inputs. Let $\mathcal{X} = \mathcal{X}_{\text{pub}} \times \mathcal{X}_{\text{pri}}$ be the space of ciphertext inputs and $\mathcal{Y} = \mathcal{Y}_{\text{pub}} \times \mathcal{Y}_{\text{pri}}$ be the space of key inputs. We define two aggregation functions as $\text{Agg}_x : \mathcal{X}^{n_x} \rightarrow \mathcal{X}$, and $\text{Agg}_y : \mathcal{Y}^{n_y} \rightarrow \mathcal{Y}$, which specify how these inputs may be combined to capture a given primitive. The definitions of the algorithms that constitute an MPFE scheme are the same as in all prior work:

- a **Setup** algorithm outputs the encryption keys for n_x encryptors and master keys for n_y key authorities. This algorithm may now run in one of three modes (**Central, Local, Decentralized**), which captures centralized setup, local/independent setup or decentralized/interactive setup.
- an **Encrypt** algorithm which is run independently by n_x users, each encoding their own message $x_i = (x_{\text{pub},i}, x_{\text{pri},i})$ with their own encryption key EK_i .
- a key-generation algorithm **KeyGen** which is run independently by all n_y key authorities, each generating its own partial key for an input $y_j = (y_{\text{pub},j}, y_{\text{pri},j})$ of its choice using its own master secret key MSK_j .
- a decryption algorithm **Decrypt**, which given input the partial keys $\{\text{SK}_i\}_{i \leq n_y}$ and partial ciphertexts $\{\text{CT}_j\}_{j \leq n_x}$ can combine them to compute $\mathcal{U}(\text{Agg}_x(\{x_i\}), \text{Agg}_y(\{y_j\}))$, where \mathcal{U} is the universal circuit.

Note that either x_i or y_j can be functions, capturing both key and ciphertext policy schemes. By suitably choosing n_x , n_y , Agg_x , Agg_y and the mode of setup, namely ($\text{mode} \in \{\text{Central, Local, Decentralized}\}$), the above abstraction lets us specify all the aforementioned primitives in a unified manner, and also allows us to instantiate these parameters in different ways to yield new, meaningful primitives. Please see Section 3 for the formal definition and Section 4 for details on how the above primitives can be expressed as instances of MPFE.

Dynamic MPFE. In the above description, we assume that the number of parties as well as the aggregation functions are input to the setup algorithm. A more powerful definition could support full dynamism, where the parties generate their own keys, join the protocol dynamically without prior agreement, and choose the functionality (in our case Agg_x and Agg_y) dynamically so that it can change for every instance of the protocol.

While dynamism is obviously desirable, it is significantly harder to instantiate since it necessitates a local setup algorithm without any co-ordination between the parties. While there do exist some constructions for dynamic FE supporting multiple users, such as adhoc MIFE [ACF⁺20] and DDFE [CDSG⁺20], most constructions in the literature are “static” and rely on centralized or interactive setup [GGG⁺14, ACGU20, CZY19, ABKW19, L†19, ABG19, CDSG⁺18a, CDSG⁺18b]. Thus, a definition which is inherently dynamic would preclude representation of most constructions in the literature.

For simplicity of notation and ease of workability, we define MPFE with and without dynamism separately. We provide the definition of the static variant in Section 3 and the dynamic variant in Appendix 3.1. We note that these two variants may be condensed to a single one using additional notation but this makes the definitions harder to work with.

Feasibility. In Appendix A, we provide a general feasibility of MPFE for circuits from the minimal assumption of MIFE for circuits.

1.2 Comparison to Prior Work

The notions of D-MCFE, aMIFE and DDFE are most closely related to our work, since they allow combining both ciphertexts and keys simultaneously. However, our notion differs from these in important ways. To begin, the setup algorithms of the above primitives have a fixed format – in D-MCFE, this is interactive,

while for **aMIFE** and **DDFE**, it is decentralized and non-interactive. Thus, **aMIFE** and **DDFE** cannot capture **D-MCFE** and vice versa. Moreover, neither of these can capture most existing constructions in the literature which have trusted, centralized setup as discussed above. In contrast, we allow setup to have either of these, as well as other formats, allowing us to capture all the above primitives and more (see Appendix 4). Next, **D-MCFE**, **aMIFE** require partial keys to represent the same function. While **DDFE** does allow partial keys to be combined in expressive ways, it does not support any function hiding. Even the support for partial input hiding in these primitives is less than complete: for instance, **aMIFE** does not support public input in the ciphertext, and while **DDFE** allows for some part of the input to be public, this is via a separate empty key ϵ . In contrast, **MPFE** captures public and private input in both the ciphertext and the function key directly, making it feasible (in the case of function inputs) and simpler (in the case of ciphertext inputs) to capture partial hiding.

The most important feature of **MPFE** is that it captures existing constructions using a uniform, simple notation, allowing to place all prior work on the same map, making these constructions easier to compare and allowing to identify gaps between these. Using our **MPFE** framework, we interpolate the space in prior work to predict several new, natural and useful primitives. Then, we provide multiple new constructions from simple, standard assumptions to address these limitations (described next), as well as identify novel new primitives (described in Section 1.5) to be constructed in future work.

1.3 New Constructions

We next describe the new constructions we provide in this work.

Multi-Authority ABE \circ IPFE. The recent work of Abdalla et al. [ACGU20] constructed a novel “composition” of **ABE** and **IPFE**, namely functional encryption schemes that combine the access control functionality of attribute based encryption with the possibility of performing linear operations on the encrypted data. In more detail, the message space contains a policy predicate ϕ and a message vector $\mathbf{v} \in \mathbb{Z}_q^\ell$, while decryption keys are jointly associated with an attribute vector $\mathbf{x} \in \{0, 1\}^n$ and a key vector $\mathbf{u} \in \mathbb{Z}_q^\ell$. The functionality provided by such a system is that a decryptor recovers the inner product value $\langle \mathbf{u}, \mathbf{v} \rangle$ if $\phi(\mathbf{x}) = 1$. Thus, it provides a fine-grained access control on top of inner product functional encryption (**IPFE**) capability. For ease of exposition, we denote this primitive, which is called “**IPFE** with fine-grained access control” in [ACGU20] by **ABE \circ IPFE** in our work¹. Abdalla et al. [ACGU20] provide a construction leveraging state of the art **ABE** from pairings to support predicates represented by Linear Secret Sharing Schemes (LSSS) in the above functional encryption scheme.

In this work, we extend the access control component to support the much more challenging multi-authority setting, i.e. “lift” the primitive of **ABE \circ IPFE** to **MA-ABE \circ IPFE**. In a multi-authority attribute-based **IPFE** system, there are n authorities, each controlling one portion of the master key material, such that every authority generates their key materials $(\text{PK}_i, \text{MSK}_i)$ locally and independently. To generate a secret key, the i -th authority, on input a single attribute bit b and a key vector \mathbf{u} , can locally sample a decryption key $\text{SK}_{i,b,\mathbf{u}}$.² Combining n such decryption keys $\{\text{SK}_{i,x_i,\mathbf{u}}\}_i$ from all n authorities yields a single-authority decryption key for attribute-key pair (\mathbf{x}, \mathbf{u}) . This yields the first construction of a nontrivial multi-authority **FE** beyond **ABE** from simple assumptions on pairings to the best of our knowledge.

Using our techniques, we also extend the decentralized attribute based encryption (**DABE**) scheme of Michalevsky and Joye [MJ18] to support inner product computations. While [MJ18] only supports inner product predicates unlike [LW11], it supports policy hiding unlike the latter – our extension inherits these features.

Function Hiding DDFE. The novel primitive of *dynamic* decentralized inner product functional encryption (**IP-DDFE**) was recently introduced by Chotard et al. [CDSG⁺20], where they also provided the first

¹We caution the reader that the notation **ABE \circ IPFE** is for readability and does not denote a formal composition.

²We are ignoring the global identifier **GID** that is necessary for tying together different authority’s keys for the purposes of this overview.

construction. In IP-DDFE, each party i can dynamically join the system by generating a public key PK_i and a master secret key MSK_i and later generate ciphertexts CT_i and secret keys SK_i by using MSK_i . For encryption, party i can encrypt a vector \mathbf{x}_i on behalf of a party set \mathcal{U}_M and a label lab_M , while in key generation, it can generate a secret key for \mathbf{y}_i on behalf of a party set \mathcal{U}_K and an agreed vector $\mathbf{y} = \{\mathbf{y}_i\}_{i \in \mathcal{U}_K}$. Decryption is possible only when a decryptor has a ciphertext set $\{\text{CT}_i\}_{i \in \mathcal{U}}$ for $(\mathbf{x}_i, \mathcal{U}, \text{lab}_M)$ and a secret-key set $\{\text{SK}_i\}_{i \in \mathcal{U}}$ for $(\mathbf{y}_i, \mathcal{U}, \mathbf{y} = \{\mathbf{y}_i\}_{i \in \mathcal{U}})$ for some $\mathcal{U}, \text{lab}_M, \mathbf{y}$ and can obtain $\sum_{i \in \mathcal{U}} \langle \mathbf{x}_i, \mathbf{y}_i \rangle$.

However, as discussed before, the primitive of DDFE does not support function hiding. We see this as a significant limitation of this notion. Function hiding is a well studied and very useful property with many applications – for instance, it allows parties to securely delegate computation to an untrusted server without the server being able to learn the functionality. In some cases, knowing the functionality and the output (which the server computes in the clear) may leak information about the underlying data. In other cases, the functionality itself may be private and protected by copyright laws. In our work, we provide a new construction for IP-DDFE which supports function hiding. In more detail, the key generator, similar to the encryptor associates a label lab_K with its vector \mathbf{y}_i and combining partial keys is only possible when their labels match. Importantly, the key vector \mathbf{y}_i may now be *hidden* analogously to the vector \mathbf{x}_i in the ciphertext. To achieve our final result, we define and construct the first *function hiding* MCFE scheme for inner products, which may be of independent interest.

Ciphertext-Policy ABE with Distributed Key Generation. We provide a distributed variant of the recent ciphertext-policy attribute based encryption scheme, constructed by Agrawal and Yamada [AY20]. In more detail, we have n key authorities, all of which run an independent, local setup algorithm, and issue partial keys for a given user GID for a given attribute \mathbf{x} . The ciphertext is associated with an NC_1 circuit C (public) and a private message m . The decryptor may obtain partial decryption keys from the n authorities and recover m if and only if $C(\mathbf{x}) = 1$. Our construction is secure based on “Learning With Errors” and relies on the generic bilinear group model as well as the random oracle model.

1.4 Technical Overview

In this section, we provide an overview of the techniques used for our constructions. We begin with our two constructions that extend multi-authority schemes [LW11, MJ18] to support inner products.

Multi-Authority ABE \circ IPFE for LSSS Access Structures. We described the functionality of MA-ABE \circ IPFE in Section 1.3. Security is defined in a multifold setting where: (1) adversary is allowed to corrupt the key authorities, (2) make key queries that do not satisfy the challenge policy predicate ϕ^* , and (3) also make key queries that satisfy the challenge policy predicate ϕ^* but decrypt to the same value for both challenge vectors (that is, $\langle \mathbf{u}, \mathbf{v}_0^* \rangle = \langle \mathbf{u}, \mathbf{v}_1^* \rangle$).

A natural first line of attack is to consider whether such a scheme can generically be built from combining these two primitives. As it turns out, any such generic construction suffers from the common problem of mix and match attacks, that is, we must prevent an authorized MA-ABE portion of the key from being used along with an IPFE portion of an unauthorized key. Another idea is to extend the ABE \circ IPFE construction of [ACGU20] to support multiple authorities. However, this work relies on the predicate encoding framework which is not suitable as-is for our application. Instead, our approach is to start with the multi-authority ABE construction by Lewko and Waters [LW11] for LSSS access structures, and show how to leverage it’s intrinsic algebraic structure to add an inner product functionality “on top” of the multi-authority ABE construction.

To begin, we provide an informal sketch of a simplified version of our construction. Recall that an access policy corresponding to a linear secret sharing scheme access structure contains a share generating matrix \mathbf{A} and a row index to party index mapping function ρ .

LSetup : The i -th authority samples a length ℓ masking vector α_i as its secret key, and publishes its encoding $[\alpha_i]_T$ in the target group as the public key.

KeyGen : To generate a secret key for key vector \mathbf{u} , the i -th authority projects α_i on the vector space defined by key vector \mathbf{u} . That is, if the attribute bit x_i is 1³, then the partial decryption key is simply $[\langle \alpha_i, \mathbf{u} \rangle]$.

Enc : For encrypting a message vector \mathbf{v} under an access policy (\mathbf{A}, ρ) , the encryptor first secret shares the message vector \mathbf{v} using the access policy \mathbf{A} into a share matrix \mathbf{S}_v . That is, \mathbf{S}_v is a random matrix with the property that for each accepting attribute \mathbf{x} there exists a reconstruction vector \mathbf{z}_x such that $\mathbf{z}_x^\top \cdot \mathbf{S}_v = \mathbf{v}^\top$. It next arranges the authority public keys $[\alpha_i]_T$ row-wise in a matrix $\mathbf{\Delta}$ as per the function ρ , that is i -th row on $\mathbf{\Delta}$ is $\rho(i)$ -th public key $[\alpha_{\rho(i)}]_T$. Finally, it output the ciphertext as the following matrix

$$\text{CT}_0 = [\mathbf{S}_v + \mathbf{\Delta} \odot (\mathbf{r} \otimes \mathbf{1}^\top)]_T, \quad \text{CT}_1 = [\mathbf{r}],$$

where \mathbf{r} is a random vector of appropriate dimension and \odot denotes the component-wise multiplications between two matrices of same dimensions.

Dec : A decryptor then simply left-multiplies CT_0 with the reconstruction vector \mathbf{z}_x and right-multiplies with the key vector \mathbf{u} to compute the following:

$$\begin{aligned} \mathbf{z}_x^\top \cdot \text{CT}_0 \cdot \mathbf{u} &= \mathbf{z}_x^\top \cdot [\mathbf{S}_v + \mathbf{\Delta} \odot (\mathbf{r} \otimes \mathbf{1}^\top)]_T \cdot \mathbf{u} \\ &= [\mathbf{v}^\top \cdot \mathbf{u} + \mathbf{z}_x^\top \cdot (\mathbf{\Delta} \odot (\mathbf{r} \otimes \mathbf{1}^\top)) \cdot \mathbf{u}]_T \\ &= [\mathbf{v}^\top \cdot \mathbf{u} + \mathbf{z}_x^\top \cdot (\mathbf{\Delta} \odot (\mathbf{r} \otimes \mathbf{u}^\top))]_T \end{aligned}$$

It next arranges the partial decryption keys $[\langle \alpha_i, \mathbf{u} \rangle]$ row-wise in a vector \mathbf{K} as per the function ρ , that is i -th element of \mathbf{K} is $\rho(i)$ -th decryption key $[\langle \alpha_{\rho(i)}, \mathbf{u} \rangle]$. It performs the component pairing between \mathbf{K} and CT_1 , and then takes the linear combination as specified by \mathbf{z}_x which can be simplified as follows:

$$\mathbf{z}_x^\top \cdot e(\mathbf{K}, \text{CT}_1) = [\mathbf{z}_x^\top \cdot (\mathbf{\Delta} \odot (\mathbf{r} \otimes \mathbf{u}^\top))]_T$$

Finally, it can recover $[\mathbf{v}^\top \cdot \mathbf{u}]_T$ from the above two terms, and learn the exponent value by brute force search.

Now in the above sketch we ignored the global identifier GID that is necessary for tying together the partial decryption keys provided by each authority, and we also ignore the modifications necessary for proving security under standard bilinear assumptions. At a very high level, for proving security we rely on ideas from the dual system paradigm [Wat09] as in the multi-authority ABE scheme of [LW11]. However, we must deal with several new challenges to adapt this paradigm to our setting, as we describe next.

In the dual system paradigm, the intuition is that the reduction algorithm first switches all the secret keys to semi-functional keys, and thereafter it also switches the challenge ciphertext to a semi-functional ciphertext, and after both these changes security follows directly from the property that semi-functional secret keys and ciphertexts are not compatible for decryption. In IPFE, we cannot hope to execute the same strategy directly since now we cannot switch all secret keys to semi-functional keys since some secret keys might allow decrypting the challenge ciphertext (but they still would not help in distinguishing by admissibility constraints on the attacker). At this point, we define the concept of *partial* semi-functional ciphertexts such that (at a high level) we first switch all the rejecting secret keys to semi-functional while leaving the accepting keys as is, and thereafter we switch the challenge ciphertext to be a “partial” semi-functional ciphertext such that this hides the non-trivial information about the encrypted message vectors.

Although this intuition seems to work at a high level, it is still insufficient since it is unclear how to switch the entire ciphertext to semi-functional in the standard model. To that end, our idea is to switch all the accepting secret keys (including the ones for satisfying predicates) to their semi-functional counterparts as well, but now ensure that the challenge ciphertext components that the accepting keys interact with are only *nominally* semi-functional. Here the difference between a regular ciphertext, a nominally semi-functional, and a standard semi-functional ciphertext is that – regular ciphertexts lie in a special subgroup with no

³As in prior ABE schemes based on bilinear maps, the key is empty when $x_i = 0$.

special blinding terms; while nominally semi-functional ciphertexts have structured blinding factors outside the special subgroup but it does not affect decryption irrespective of the type of secret key being used; and a standard semi-functional ciphertext has unstructured blinding factors outside the special subgroup such that it affects decryption when using semi-functional keys. Now switching portions of the challenge ciphertext as nominally semi-functional is necessary because of two reasons: first, making the entire challenge ciphertext semi-functional will affect decryption w.r.t. accepting keys which will be distinguishable for the adversary; second, it is unclear how to sample the challenge ciphertext in which only one component is semi-functional while other are regular sub-encryptions due to the fact that these different ciphertext components are significantly correlated. Thus, we get around this barrier by ensuring that the challenge ciphertext is sampled as what we call a partial semi-functional ciphertext (which has nominally semi-functional components along with a standard semi-functional component).

Please see Section 5 for the formal construction and proof. Our construction relies on standard assumptions over composite-order bilinear groups, but could be also be easily adapted to prime-order groups with a security proof in the generic group model as in [LW11].

DABE \circ IPFE for Inner Product Predicates, with Policy Hiding. Next, we extend the construction of decentralized attribute based encryption by Michalevsky and Joye [MJ18] to incorporate inner product functional encryption. Observe that [MJ18] supports only inner product predicates but allows for hiding the policy in the ciphertext. While our extension to the scheme of [MJ18] also yields a multi-authority ABE extended to support inner products as above, the details of the transformation are quite different. We observe that the algebraic structure of [MJ18] makes it amenable to incorporating the IPFE functionality using ideas developed in the literature for constructing IPFE generically from public-key encryption which have special structural and homomorphic properties [ABDCP15, ALS16, BJK15]. We proceed to describe this transformation next.

In an overly simplified version of the Michalevsky-Joye construction, one could interpret the i -th key authority as simply sampling a pair of secret exponents $\delta_i, w_i \leftarrow \mathbb{Z}_p$, where δ_i is regarded as the partial message masking term, while w_i is considered the i -th attribute bit binding term. Now each authority's public key is simply set as the group encodings $[\delta_i]$ and $[w_i]$. Implicitly, the scheme uses the linear combination of partial message masking terms $\delta = \sum_i \delta_i$ to derive the main message masking term (used for deriving the secret key encapsulating the message, or the KEM key in short).

To encrypt a message m under attribute \mathbf{x} , the user chooses randomness $r \leftarrow \mathbb{Z}_p$ and computes $[r\delta]_T$ to be used as the KEM key, and binds each attribute bit to a ciphertext component as $[(x_i\alpha + w_i)r]$ (where $[\alpha]$ is taken from the CRS). It sets the ciphertext to be $C_0 = [r]$, $C_m = m \cdot [r\delta]_T$, and $C_i = [(x_i\alpha + w_i)r]$ for $i \in [n]$. While a partial secret key for policy vector \mathbf{y} for user GID is simply generated as $K_{i,\mathbf{y}} = [\delta_i - y_i w_i h]$ where $[h]$ is computed as $H(\text{GID})$ so as to bind the different authorities' secret keys. The decryption can be simply performed given the bilinear operation as:

$$\begin{aligned} \text{Dec}(\{K_{i,\mathbf{y}}\}_i, \text{CT}) &= \frac{C_m}{\prod_i e(C_i, H(\text{GID})^{y_i}) \cdot \prod_i e(K_{i,\mathbf{y}}, C_0)} \\ &= \frac{m \cdot [r\delta]_T}{[\langle \mathbf{x}, \mathbf{y} \rangle \alpha r h + \sum_i w_i h y_i r]_T \cdot [\delta r - \sum_i y_i w_i h r]_T} = \frac{m}{[\langle \mathbf{x}, \mathbf{y} \rangle \alpha r h]_T} \end{aligned}$$

As discussed above, we upgrade the [MJ18] construction using ideas from [ABDCP15, ALS16, BJK15] as follows. During key generation, each authority now samples a vector of partial masking terms instead of a single element, i.e. $\boldsymbol{\delta}_i \leftarrow \mathbb{Z}_p^\ell$, and appropriately sets the public key too. Implicitly, the main message masking term is now set as $\boldsymbol{\delta} = \sum_i \boldsymbol{\delta}_i$. To encrypt a message vector \mathbf{u} under attribute vector \mathbf{x} , the user chooses randomness $r \leftarrow \mathbb{Z}_p$ and computes $[r\boldsymbol{\delta}]_T$ to be used as the KEM key for encrypting \mathbf{u} index-by-index, and binds the attribute bit as before. Thus, only the message binding ciphertext component changes to $C_m = [r\boldsymbol{\delta} + \mathbf{u}]_T$. Looking ahead, it will be decryptor's job to first homomorphically take an inner product between the C_m vector and the inner product key vector \mathbf{v} . Next, a partial secret key for policy vector \mathbf{y} and inner-product vector \mathbf{v} for user GID is generated as $K_{i,\mathbf{y},\mathbf{v}} = [\sum_j \delta_{i,j} v_j - y_i w_i h]$. In words, the idea here is

that the partial secret key now uses a linear combination of its partial (un)masking term $\sum_j \delta_{i,j} v_j$ depending on the underlying inner-product vector \mathbf{v} . The decryption can be naturally extended by performing inner products via the bilinear operations.

As in the case of our first construction, the proof techniques in [MJ18] do not apply directly as they were specially designed for ABE which is an all-or-nothing encryption primitive, and do not translate directly to IPFE. Again, we solve this issue by a careful analysis in the dual system paradigm [Wat09]. We refer the reader to Section 8 for more details.

Function-Hiding DDFE for Inner Products. In this section, we describe the main ideas in the construction of our function hiding DDFE for inner products. The functionality of IP-DDFE was discussed in Section 1.3. Informally, the security of DDFE requires that the adversary cannot distinguish two sets $\{\text{CT}_i^0\}$ and $\{\text{CT}_i^1\}$ of ciphertexts even given a set $\{\text{SK}_i\}$ of secret keys and a set $\{\text{MSK}_i\}$ of master secret keys of corrupted parties as long as two sets of values are the same that are legitimately obtained from $\{\text{CT}_i^0\}$ and $\{\text{CT}_i^1\}$ using $\{\text{SK}_i\}$ and $\{\text{MSK}_i\}$. Let us recall dynamic decentralized inner product functional encryption (IP-DDFE) by Chotard et al. [CDSG⁺20].

The starting point of the IP-DDFE scheme of [CDSG⁺20] is the multi-client inner product functional encryption (IP-MCFE) scheme in [CDSG⁺18a], where participants $\{1, \dots, n\}$ in the system are a priori fixed, and there is an authority who generates encryption keys mcEK_i for each party and a master secret key mcMSK , which is used to generate secret keys mcSK . Here, $\text{mcMSK} = \{\text{mcMSK}_i\}_{i \in [n]}$ and $\text{mcEK}_i = \text{mcMSK}_i$ (and we denote an encryption key for i by mcMSK_i in what follows). We also recall that in MCFE, only a set of ciphertexts with the same label can be decrypted. Chotard et al. [CDSG⁺20] lifted the IP-MCFE scheme to an IP-DDFE scheme via following steps. First, each party joins the system dynamically by generating a key K_i of a pseudorandom function (PRF) as a master secret key MSK_i . In encryption and key generation for party set \mathcal{U} , party $i \in \mathcal{U}$ generates $\text{mcMSK}_{i,\mathcal{U}}$ on the fly, which corresponds to mcMSK_i of the IP-MCFE scheme for participants \mathcal{U} . Then, it can generate $\text{mcCT}_{i,\mathcal{U}}$ and $\text{mcSK}_{i,\mathcal{U}}$ with $\text{mcMSK}_{i,\mathcal{U}}$, which corresponds to CT_i and SK_i of the IP-DDFE scheme. Second, they introduce a class of DDFE called DSum, which allows a decryptor to securely obtain $\text{mcSK}_{\mathcal{U}} = \sum_{i \in \mathcal{U}} \text{mcSK}_{i,\mathcal{U}}$ from encryption of partial secret keys $\{\text{mcSK}_{i,\mathcal{U}}\}_{i \in \mathcal{U}}$. Then, the decryptor can compute $\text{mcDec}(\text{mcSK}_{\mathcal{U}}, \{\text{mcCT}_{i,\mathcal{U}}\}_{i \in \mathcal{U}})$. DSum also plays a role in preventing combination of partial secret keys for which the agreed vectors are inconsistent.

Our Function-Hiding IP-DDFE. Our approach is to lift function-hiding IP-MCFE to function-hiding IP-DDFE following their blueprint. Unfortunately, there are no function-hiding IP-MCFE schemes, and we need to start with constructing this. Our first idea is to leverage the recent conversion by Abdalla et al. [ABG19] from IPFE to IP-MCFE. However, this idea does not work since all parties share the same encryption key of an IPFE scheme in their converted schemes, and once a single party is corrupted, the adversary can completely decode the entire secret keys. Thus, their conversion is not applicable to the function-hiding setting.

To address this challenge, we devise a new technique to convert function-hiding IPFE to function-hiding IP-MCFE, which is inspired by the function-hiding multi-input IPFE scheme by Datta et al. [DOT18]. In their scheme, each party i has a master secret key $i\text{MSK}_i$ of a function-hiding IPFE scheme, the ciphertext miCT_i of \mathbf{x}_i is $i\text{CT}_i[(\mathbf{x}_i, 1)]$, and the secret key miSK of $\{\mathbf{y}_i\}_{i \in [n]}$ is $\{i\text{SK}_i[(\mathbf{y}_i, r_i)]\}_{i \in [n]}$ where r_i are randomly chosen so that $\sum_{i \in [n]} r_i = 0$. $i\text{CT}_i[\mathbf{x}]$ and $i\text{SK}_i[\mathbf{y}]$ denotes the ciphertext of \mathbf{x} and the secret key of \mathbf{y} in the function-hiding IPFE scheme, respectively. To lift their MIFE to MCFE, we need to add the label checking mechanism and security against corruption of parties. Fortunately, we can achieve the latter almost for free since each party uses independent master secret key and corruption of a party does not affect other parties' ciphertexts and secret keys. We can achieve the former by changing miCT_i to $i\text{CT}_i[(\mathbf{x}_i, t_i)]$ where $t_i = H(\text{lab})$ is a hash of a label. Then, a decryptor can learn $\sum (\langle \mathbf{x}_i, \mathbf{y}_i \rangle + t_i r_i)$, which reveals $\sum \langle \mathbf{x}_i, \mathbf{y}_i \rangle$ only when $t_1 = \dots = t_n$. We can prove the masking term $t_i r_i$ hides $\langle \mathbf{x}_i, \mathbf{y}_i \rangle$ under the SXDH assumption in the random oracle model.

Our next step is to lift function-hiding IP-MCFE to function-hiding IP-DDFE. To do so, we must address additional technical challenges as described next. In the original definition of IP-DDFE, recall that each

secret key is associated with $(\mathbf{y}_i, \mathcal{U}, \mathbf{y} = \{\mathbf{y}_{i'}\}_{i' \in \mathcal{U}})$ where the first element \mathbf{y}_i is a vector for a linear function while the third element \mathbf{y} is an agreed vector that controls combination of partial secret keys. More precisely, a decryptor can combine partial secret keys to obtain a full secret key for \mathbf{y} only when it has $\{\text{SK}_i\}_{i \in \mathcal{U}}$ associated with \mathbf{y} . However, \mathbf{y} cannot be hidden in the blueprint by Chotard et al. To tackle this, we observe that the role of the agreed vector is analogous to a label in the ciphertext, controlling combination of partial secret keys. Thus, we alternatively use an independent label lab_K to create a natural symmetry between inputs for encryption and key generation. Now, since the vector \mathbf{y}_i for a linear function can be hidden by function-hiding IP-MCFE, we obtain function-hiding IP-DDFE.

Another deviation from their blueprint arises in the part that securely generates $\text{mcSK}_{\mathcal{U}}$ from $\text{mcSK}_{i, \mathcal{U}}$. In our IP-MCFE construction, $\text{mcSK}_{i, \mathcal{U}} = \text{iSK}_i[(\mathbf{y}_i, r_i)]$ and $\text{mcSK}_{\mathcal{U}} = \{\text{iSK}_i[(\mathbf{y}_i, r_i)]\}_{i \in \mathcal{U}}$. Thus, we do not need to sum up $\text{mcSK}_{i, \mathcal{U}}$ to obtain $\text{mcSK}_{\mathcal{U}}$, instead, each party has to somehow generate a secret share r_i without interaction such that $\sum_{i \in [\mathcal{U}]} r_i = 0$ only when all $\text{mcSK}_{i, \mathcal{U}}$ are generated on behalf of the same label. To handle this issue, we employ a technique by Chase and Chow [CC09] to generate such shares via pseudorandom function. Please see Section 6 for more details.

Distributed Ciphertext-Policy ABE The recent construction of Agrawal-Yamada [AY20] proposed a succinct ciphertext-policy ABE for log-depth circuits provably secure under LWE in the bilinear generic group model. In our work, we extend the setup and key generation in [AY20] among a polynomial number of authorities that are working completely *non-interactively* and asynchronously. We start by describing the syntax of a distributed CP-ABE scheme. In a fully distributed setting, the authorities run their local setup algorithms individually to generate a fresh master public-secret key pair (PK, MSK) per authority such that given a sequence of, say N , master public keys $\{\text{PK}_i\}_{i \in [N]}$, an encryptor could encrypt a message μ for a predicate circuit F of its choice. Such ciphertexts can be decrypted after obtaining a partial predicate key from all N authorities for a consistent identifier GID , and attribute vector \mathbf{x} such that $F(\mathbf{x}) = 1$. Note that here the key generation algorithm is run locally (and independently) by each authority, which on input its master key MSK_i along with GID and attribute \mathbf{x} , computes a partial key $\text{SK}_{i, \text{GID}, \mathbf{x}}$. While correctness is natural, security must be defined carefully.

In this work, we consider the strongest form of corruption, where we allow the adversary to pick the key parameters for all corrupt authorities, and also allow it to query honest authorities on identifier-attribute pairs (GID, \mathbf{x}) such that $F^*(\mathbf{x}) = 1$ (where F^* is the challenge predicate circuit) as long as there is at least one honest authority to which the adversary did not query the pair (GID, \mathbf{x}) . All other queries are unconstrained since if $F^*(\mathbf{x}) = 0$, then such keys should not be useful for decryption to begin with. The intuition behind allowing the queries to honest authorities such that $F^*(\mathbf{x}) = 1$ is that we want to prevent partial secret keys for two distinct accepting attributes provided by two distinct authorities to be usable for decryption.

To describe our construction, we recall the high level structure of the Agrawal-Yamada scheme [AY20], which in turn uses the BGG^+ ABE construction [BGG⁺14]. Roughly, a BGG^+ ciphertext is sampled in two steps — first, it samples a sequence of 2ℓ encodings $\{\psi_{i,b}\}_{i,b}$; second, depending upon the attribute \mathbf{x} the final ciphertext consists of ℓ encodings $\{\psi_{i,x_i}\}_i$. (Note that BGG^+ is a key-policy scheme, whereas we are building a ciphertext-policy system.) The main idea behind the ciphertext-policy ABE of [AY20] is as follows:

Setup : Sample 2ℓ random exponents $w_{i,b}$, store it as master secret key, and give its encoding $\{[w_{i,b}]_1\}_{i,b}$ as the public key.

KeyGen : To generate a secret key for attribute $\mathbf{x} \in \{0, 1\}^\ell$, first sample a random exponent δ and then give out $[\delta/w_{i,x_i}]_2$ for $i \in [\ell]$ as the secret key.

Enc : To encrypt under predicate F , the encryptor samples all 2ℓ BGG^+ encodings $\{\psi_{i,b}\}_{i,b}$, and also samples a random exponent γ . It then gives out the ciphertext as a BGG^+ secret key for predicate C along with encodings $[\gamma w_{i,b} \psi_{i,b}]_1$ for $i \in [\ell], b \in \{0, 1\}$.

Dec : A decryptor pairs the encodings $[\gamma w_{i,x_i} \psi_{i,x_i}]_1$ with $[\delta/w_{i,x_i}]_2$ to learn $[\gamma \delta \psi_{i,x_i}]_T$, and then it performs the BGG^+ decryption in the exponent to learn the plaintext.

For the multi-authority extension, each authority samples its own sequence of 2ℓ random exponents $w_{i,b}^{(j)}$ for $j \in [N]$. Then during encryption, the encryptor N -out-of- N (additively) secret shares the BGG^+ encodings $\{\psi_{i,b}\}_{i,b}$ into $\{\phi_{i,b}^{(j)}\}_{i,b}$ for $j \in [N]$. Now it encodes each sequence of $\{\phi_{i,b}^{(j)}\}_{i,b}$ terms under the corresponding authority’s master public key as above. During decryption, a decryptor will first recover $\{\phi_{i,x_i}^{(j)}\}_i$ for all j in the exponent, then add them to reconstruct the actual BGG^+ ciphertext $\{\psi_{i,x_i}\}_i$ which it can decrypt as before. In order to let multiple independent authorities sample the same δ , we rely on a hash function which we model as a ROM, and set $[\delta]_2 = H(\text{GID})$.

Although our multi-authority transformation is natural, the proof does not follow trivially from [AY20]. This is primarily due to the fact that in the distributed setting, the adversary could potentially make key queries on accepting attributes as long as there is at least one honest party that does not receive the same query. Such queries did not exist in the single-authority setting. However, we can extend the single-authority proof to the multi-authority setting by a careful analysis of the additional “bad” zero-test queries that an adversary can make. Please see Section 7 for more details.

1.5 Predicting New and Useful Primitives via MPFE

One of the most exciting benefits of MPFE is that it provides the right framework to pose new, compelling questions that have not been studied before. For example, a very interesting question is what new kinds of dynamic key accumulation are possible, namely how to combine keys of different users chosen dynamically. So far, most existing literature on FE systems that enable aggregation of multiple decryption keys still consider very restricted scenarios: (i) each partial decryption key corresponds to a portion of a much larger decryption key of a single user (e.g., distributed/decentralized/multi-authority FE etc), (ii) each partial decryption key corresponds to a function and many such keys may be combined if they each encode the *same* function (e.g. adhoc MIFE, D-MCFE).

However, the ability to combine keys in much more creative ways can enable several cool new applications. As an example, consider the following notion of “reputation point based encryption” – in this setting, each user key is associated with a subject tag T (say math, history etc) and a reputation value v (that is, a point score denoted as an integer). Now an encryptor specifies a tag T' along with a threshold reputation value w , and hides its message m under it. That is, $\text{CT}(T', w, m)$ denotes such a ciphertext, and we require the functionality that such a ciphertext should be decryptable by any sequence of user keys $\text{SK}(T_1, v_1), \dots, \text{SK}(T_\ell, v_\ell)$ where all the subject tags match ($T' = T_1 \dots = T_\ell$) and the combined reputation value of the group $\sum_{i \leq \ell} v_i$ is greater than threshold w . For example, an encryption of a message under subject ‘math’ and minimum reputation value of 1000 points can be decrypted by not only a single user with 1000 reputation points in ‘math’ but also by say a group of three users with 400, 250, 350 reputation points (respectively) in ‘math’, but not by a group of users who satisfy either the subject check or the reputation point check *but not both*. To the best of our knowledge, such an encryption framework has not been studied before, but our MPFE framework enables expressing and introducing such an encryption functionality.

To supplement this new abstraction, we sketch a generic construction of this new multi-key FE primitive from a multi-authority ABE scheme in Section 9. We also identify other useful new primitives that have not been studied before as targets for future work.

Acknowledgements. We thank Fabrice Mouhartem for helpful discussions and collaboration during early stages of this work. We are grateful to Edouard Dufour Sans, David Pointcheval and Romain Gay for insightful discussions about DDFE which led to refinements in the definition of dynamic MPFE.

2 Preliminaries

Notation. We begin by defining the notation that we will use throughout the paper. We use bold letters to denote vectors and the notation $[a, b]$ to denote the set of integers $\{k \in \mathbb{N} \mid a \leq k \leq b\}$. We use $[n]$ to

denote the set $[1, n]$. Concatenation is denoted by the symbol \parallel .

Throughout the paper, we use λ to denote the security parameter. We say a function $f(\lambda)$ is *negligible* if it is $O(\lambda^{-c})$ for all $c > 0$, and we use $\text{negl}(\lambda)$ to denote a negligible function of λ . We say $f(\lambda)$ is *polynomial* if it is $O(\lambda^c)$ for some constant $c > 0$, and we use $\text{poly}(\lambda)$ to denote a polynomial function of λ . We use the abbreviation PPT for probabilistic polynomial-time. The function $\log x$ is the base 2 logarithm of x .

2.1 Access Structures and Linear Secret-Sharing Schemes

We recall the concepts of access structures and linear secret-sharing schemes (LSSS). We follow the notation from prior works [GPSW06, LW11].

Definition 2.1 (Access Structures). Let $\{P_i\}_{i \in [n]}$ be a set of parties. A collection $\mathbb{A} \subseteq 2^{\{P_1, \dots, P_n\}}$ is monotone if $\forall B, C : \text{if } B \in \mathbb{A} \text{ and } B \subseteq C, \text{ then } C \in \mathbb{A}$. An access structure (respectively, monotone access structure) is a collection (respectively, monotone collection) \mathbb{A} of non-empty subsets of $\{P_i\}_{i \in [n]}$. The sets in \mathbb{A} are called the authorized sets, and the sets not in \mathbb{A} are called the unauthorized sets.

As in prior work in the bilinear setting, attributes will play the role of parties and we will only consider monotone access structures. We observe that more general access structures can be (inefficiently) realized with our techniques by letting the negation of an attribute be a separate attribute (this doubles the total number of attributes).

Definition 2.2 (Linear Secret-Sharing Schemes (LSSS)). A secret sharing scheme Π over a set of parties \mathcal{P} is called linear (over \mathbb{Z}_p) if:

1. The shares for each party form a vector over \mathbb{Z}_p .
2. There exists a $\ell \times m$ matrix \mathbf{A} called the share-generating matrix for Π such that the x -th row of \mathbf{A} is labeled by a party $\rho(x)$ (ρ is a function from row index to party index). When we consider the column vector $\mathbf{v} = (s, r_2, \dots, r_n)^\top$, where $s \in \mathbb{Z}_p$ is the secret to be shared and $r_2, \dots, r_n \in \mathbb{Z}_p$ are randomly chosen, then $\mathbf{A} \cdot \mathbf{v}$ is the vector of ℓ shares of the secret s according to Π . The share $(\mathbf{A} \cdot \mathbf{v})_x$ belongs to party $\rho(x)$.

Linear reconstruction property: Let S denote an authorized set, and define $I \subseteq [\ell]$ as $I = \{x : \rho(x) \in S\}$. Then the first basis vector of canonical basis of \mathbb{Z}_p^m , that is $(1, 0, \dots, 0)^\top$ is in the span of rows of \mathbf{A} indexed by I , and there exist constants $\{w_x \in \mathbb{Z}_p\}_{x \in I}$ such that, for any valid shares $\{\text{sh}_x\}_x$ of a secret s according to Π , we have: $\sum_{x \in I} w_x \text{sh}_x = s$. And, these constants can be found in polynomial time with respect to the size of the share-generating matrix \mathbf{A} .

2.2 Bilinear Map Preliminaries

Here, we introduce our notation for bilinear maps and the bilinear generic group model following Baltico et.al [BCFG17], who specializes the framework by Barthe [BFF+14] for defining generic k -linear groups to the bilinear group settings. The definition closely follows that of Maurer [Mau05], which is equivalent to the alternative formulation by Shoup [Sho97].

Notation on Bilinear Maps. In this work, we use bilinear groups for our various MPFE constructions. Some of our constructions rely on composite order bilinear groups, while other rely on prime order asymmetric bilinear groups. Below we define our notations for both.

Asymmetric prime order bilinear groups. An asymmetric bilinear group generator takes as input 1^λ and outputs a group description

$\mathbb{G} = (q, \mathbb{G}_1, \mathbb{G}_2, \mathbb{G}_T, e, g_1, g_2)$, where q is a prime of $\Theta(\lambda)$ bits, $\mathbb{G}_1, \mathbb{G}_2$, and \mathbb{G}_T are cyclic groups of order q , $e : \mathbb{G}_1 \times \mathbb{G}_2 \rightarrow \mathbb{G}_T$ is a non-degenerate bilinear map, and g_1 and g_2 are generators of \mathbb{G}_1 and \mathbb{G}_2 , respectively. In more detail, we have:

- Bilinearity: $\forall g_1 \in \mathbb{G}_1, g_2 \in \mathbb{G}_2, a, b \in \mathbb{Z}_p, e(g_1^a, g_2^b) = e(g_1, g_2)^{ab}$,
- Non-Degeneracy: $e(g_1, g_2) \neq 1_{\mathbb{G}_T}$ for $g_1 \neq 1_{\mathbb{G}_1}$ and $g_2 \neq 1_{\mathbb{G}_2}$, where $1_{\mathbb{G}_1}, 1_{\mathbb{G}_2}$ and $1_{\mathbb{G}_T}$ are the identity elements of groups $\mathbb{G}_1, \mathbb{G}_2$ and \mathbb{G}_T respectively.

We require that the group operations in $\mathbb{G}_1, \mathbb{G}_2$, and \mathbb{G}_T as well as the bilinear map e can be efficiently computed. We employ the implicit representation of group elements: for a matrix \mathbf{A} over \mathbb{Z}_q , we define $[\mathbf{A}]_1 := g_1^{\mathbf{A}}, [\mathbf{A}]_2 := g_2^{\mathbf{A}}, [\mathbf{A}]_T := g_T^{\mathbf{A}}$, where exponentiation is carried out component-wise.

We also use the following less standard notations. For vectors $\mathbf{w} = (w_1, \dots, w_\ell)^\top \in \mathbb{Z}_q^\ell$ and $\mathbf{v} = (v_1, \dots, v_\ell)^\top \in \mathbb{Z}_q^\ell$ of the same length, $\mathbf{w} \odot \mathbf{v}$ denotes the vector that is obtained by component-wise multiplications. Namely, $\mathbf{v} \odot \mathbf{w} = (v_1 w_1, \dots, v_\ell w_\ell)^\top$. When $\mathbf{w} \in (\mathbb{Z}_q^*)^\ell$, $\mathbf{v} \oslash \mathbf{w}$ denotes the vector $\mathbf{v} \oslash \mathbf{w} = (v_1/w_1, \dots, v_\ell/w_\ell)^\top$. It is easy to verify that for vectors $\mathbf{c}, \mathbf{d} \in \mathbb{Z}_q^\ell$ and $\mathbf{w} \in (\mathbb{Z}_q^*)^\ell$, we have $(\mathbf{c} \odot \mathbf{w}) \odot (\mathbf{d} \oslash \mathbf{w}) = \mathbf{c} \odot \mathbf{d}$. For group elements $[\mathbf{v}]_1 \in \mathbb{G}_1^\ell$ and $[\mathbf{w}]_1 \in \mathbb{G}_2^\ell$, $[\mathbf{v}]_1 \odot [\mathbf{w}]_2$ denotes $([v_1 w_1]_T, \dots, [v_\ell w_\ell]_T)^\top$, which is efficiently computable from $[\mathbf{v}]_1$ and $[\mathbf{w}]_2$ using the bilinear map e .

We now recall the k -linear assumption on prime order bilinear groups over source group \mathbb{G}_1 . It can be analogously define for group \mathbb{G}_2 as well.

Assumption 2.3 (k -linear in \mathbb{G}_1). For every PPT adversary \mathcal{A} , there exists a negligible function $\text{negl}(\cdot)$ such that for all $\lambda \in \mathbb{N}$,

$$\Pr \left[\begin{array}{l} \mathcal{A}(\Pi, g_1, g_2, g_1^{\mathbf{a}}, g_1^{\mathbf{z}^\beta}) = b : \\ \Pi = (q, \mathbb{G}_1, \mathbb{G}_2, \mathbb{G}_T, e(\cdot, \cdot)) \leftarrow \text{Gen}(1^\lambda) \\ g_1 \leftarrow \mathbb{G}_1, g_2 \leftarrow \mathbb{G}_2, \beta \leftarrow \{0, 1\} \\ \mathbf{a} \leftarrow \mathbb{Z}_q^k, \mathbf{b} \leftarrow \mathbb{Z}_q^k, \mathbf{z}_1 \leftarrow \mathbb{Z}_q^{k+1} \\ \mathbf{z}_0 = (a_1 b_1, \dots, a_k b_k, \sum_i b_i) \end{array} \right] \leq 1/2 + \text{negl}(\lambda).$$

Generic Bilinear Group Model. Let $\mathbb{G} = (q, \mathbb{G}_1, \mathbb{G}_2, \mathbb{G}_T, e, g_1, g_2)$ be a bilinear group setting, L_1, L_2 , and L_T be lists of group elements in $\mathbb{G}_1, \mathbb{G}_2$, and \mathbb{G}_T respectively, and let \mathcal{D} be a distribution over L_1, L_2 , and L_T . The generic group model for a bilinear group setting \mathbb{G} and a distribution \mathcal{D} is described in Fig. 2.1. In this model, the challenger first initializes the lists L_1, L_2 , and L_T by sampling the group elements according to \mathcal{D} , and the adversary receives handles for the elements in the lists. For $s \in \{1, 2, T\}$, $L_s[h]$ denotes the h -th element in the list L_s .

The handle to this element is simply the pair (s, h) . An adversary running in the generic bilinear group model can apply group operations and bilinear maps to the elements in the lists. To do this, the adversary has to call the appropriate oracle specifying handles for the input elements. The challenger computes the result of a query, stores it in the corresponding list, and returns to the adversary its (newly created) handle. Handles are not unique (i.e., the same group element may appear more than once in a list under different handles).

We remark that we slightly simplify the generic group model of Baltico et. al [BCFG17]. Whereas they allow the adversary to access the equality test oracle, which is given two handles (s, h_1) and (s, h_2) and returns 1 if $L_s[h_1] = L_s[h_2]$ and 0 otherwise for all $s \in \{1, 2, T\}$, we replace this oracle with the zero-test oracle, which is given a handle (s, h) and returns 1 if $L_s[h] = 0$ and 0 otherwise only for the case of $s = T$. We claim that even with this modification, the model is equivalent to the original one. This is because we can perform the equality test for (s, h_1) and (s, h_2) using our restricted oracles as follows. Let us first consider the case of $s = T$. In this case, we can get the handle (T, h') corresponding to $L_T[h_1] - L_T[h_2]$ by calling neg_T and add_T . We then make a zero-test query for (T, h') . Clearly, we get 1 if $L_s[h_1] = L_s[h_2]$ and 0 otherwise. We next consider the case of $s \in \{1, 2\}$. This case can be reduced to the case of $s = T$ by lifting the group elements corresponding to h_1 and h_2 to the group elements in \mathbb{G}_T by taking bilinear maps with an arbitrary non-unit group element in \mathbb{G}_{3-s} , which is possible by calling map_e .

Symbolic Group Model. The symbolic group model for a bilinear group setting \mathbb{G} and a distribution \mathcal{D}_P gives to the adversary the same interface as the corresponding generic group model, except that internally the

State: Lists L_1, L_2, L_T over $\mathbf{G}_1, \mathbf{G}_2, \mathbf{G}_T$ respectively.

Initializations: Lists L_1, L_2, L_T sampled according to distribution \mathcal{D} .

Oracles: The oracles provide black-box access to the group operations, the bilinear map, and equalities.

- For all $s \in \{1, 2, T\}$: $\text{add}_s(h_1, h_2)$ appends $L_s[h_1] + L_s[h_2]$ to L_s and returns its handle $(s, |L_s|)$.
- For all $s \in \{1, 2, T\}$: $\text{neg}_s(h_1, h_2)$ appends $L_s[h_1] - L_s[h_2]$ to L_s and returns its handle $(s, |L_s|)$.
- $\text{map}_e(h_1, h_2)$ appends $e(L_1[h_1], L_2[h_2])$ to L_T and returns its handle $(T, |L_T|)$.
- $\text{zt}_T(h)$ returns 1 if $L_T[h] = 0$ and 0 otherwise.

All oracles return \perp when given invalid indices.

Figure 2.1: Generic group model for bilinear group setting $\mathbf{G} = (q, \mathbf{G}_1, \mathbf{G}_2, \mathbf{G}_T, e, g_1, g_2)$ and distribution \mathcal{D} .

challenger stores lists of element in the field $\mathbf{Z}_p(X_1, \dots, X_n)$ instead of lists of group elements. The oracles add_s , neg_s , map , and zt computes addition, negation, multiplication, and equality in the field. In our work, we will use the subring $\mathbf{Z}_p[X_1, \dots, X_n, 1/X_1, \dots, 1/X_n]$ of the entire field $\mathbf{Z}_p(X_1, \dots, X_n)$. Note that any element f in $\mathbf{Z}_p[X_1, \dots, X_n, 1/X_1, \dots, 1/X_n]$ can be represented as

$$f(X_1, \dots, X_n) = \sum_{(c_1, \dots, c_n) \in \mathbf{Z}^n} a_{c_1, \dots, c_n} X_1^{c_1} \dots X_n^{c_n}$$

using $\{a_{c_1, \dots, c_n} \in \mathbf{Z}_p\}_{(c_1, \dots, c_n) \in \mathbf{Z}^n}$, where we have $a_{c_1, \dots, c_n} = 0$ for all but finite $(c_1, \dots, c_n) \in \mathbf{Z}^n$. Note that this expression is unique.

Composite order bilinear groups. A composite order bilinear group generator takes as input 1^λ and number of prime-order subgroups k , and outputs a group description $((p_1, p_2, \dots, p_k), (N, \mathbb{G}, \mathbb{G}_T, e(\cdot, \cdot)))$, where p_i 's are primes such that $p_i \in \{2^{\lambda-1}, \dots, 2^\lambda - 1\}$ for all $i \in [k]$, $N = \prod_i p_i$, \mathbb{G} and \mathbb{G}_T are groups of order N and e is an efficiently computable function computing a non-degenerate bilinear map taking two group elements of \mathbb{G} to a group element in \mathbb{G}_T . For any set $S \subseteq [k]$, let $\mathbb{G}_S \subseteq \mathbb{G}$ denote the (unique) subgroup of order $p_S = \prod_{i \in S} p_i$. Observe that for any $S \subseteq [k]$, given $\{p_i\}_{i \in S}$, one can sample a uniformly random element from \mathbb{G}_S .

We now state the assumptions we make on composite order bilinear groups. All our assumptions belong in the family of subgroup decision assumptions introduced by Bellare, Waters, and Yilek [BWY11]. These same assumptions appear in the work of Lewko and Waters [LW11].

Assumption 2.4. We say that the assumption holds with respect to Gen for three primes if for any stateful PPT adversary \mathcal{A} , there exists a negligible function $\text{negl}(\cdot)$ such that for all $\lambda \in \mathbb{N}$,

$$\Pr \left[\mathcal{A}(\Pi, g_1, T_b) = b : \begin{array}{l} ((p_1, p_2, p_3), \Pi) \leftarrow \text{Gen}(1^\lambda, 3); \Pi = (N, \mathbb{G}, \mathbb{G}_T, e) \\ g_1 \leftarrow \mathbb{G}_1; T_0 \leftarrow \mathbb{G}; T_1 \leftarrow \mathbb{G}_1; b \leftarrow \{0, 1\} \end{array} \right] \leq 1/2 + \text{negl}(\lambda).$$

Assumption 2.5. We say that the assumption holds with respect to Gen for three primes if for any stateful PPT adversary \mathcal{A} , there exists a negligible function $\text{negl}(\cdot)$ such that for all $\lambda \in \mathbb{N}$,

$$\Pr \left[\mathcal{A}(\Pi, g_1, g_3, h_1 h_2, T_b) = b : \begin{array}{l} ((p_1, p_2, p_3), \Pi) \leftarrow \text{Gen}(1^\lambda, 3); \Pi = (N, \mathbb{G}, \mathbb{G}_T, e) \\ g_1, h_1 \leftarrow \mathbb{G}_1; h_2 \leftarrow \mathbb{G}_2; g_3 \leftarrow \mathbb{G}_3 \\ T_0 \leftarrow \mathbb{G}_1; T_1 \leftarrow \mathbb{G}_{\{1,2\}}; b \leftarrow \{0, 1\} \end{array} \right] \leq 1/2 + \text{negl}(\lambda).$$

Assumption 2.6. We say that the assumption holds with respect to Gen for three primes if for any stateful PPT adversary \mathcal{A} , there exists a negligible function $\text{negl}(\cdot)$ such that for all $\lambda \in \mathbb{N}$,

$$\Pr \left[\begin{array}{l} \mathcal{A}(\Pi, g_1, h_1 g_3, g_2 h_3, T_b) = b : \\ \begin{array}{l} ((p_1, p_2, p_3), \Pi) \leftarrow \text{Gen}(1^\lambda, 3); \Pi = (N, \mathbb{G}, \mathbb{G}_T, e) \\ g_1, h_1 \leftarrow \mathbb{G}_1; g_2 \leftarrow \mathbb{G}_2; g_3, h_3 \leftarrow \mathbb{G}_3 \\ T_0 \leftarrow \mathbb{G}_{\{1,2\}}; T_1 \leftarrow \mathbb{G}_{\{1,3\}}; b \leftarrow \{0, 1\} \end{array} \end{array} \right] \leq 1/2 + \text{negl}(\lambda).$$

Assumption 2.7. We say that the assumption holds with respect to Gen for three primes if for any stateful PPT adversary \mathcal{A} , there exists a negligible function $\text{negl}(\cdot)$ such that for all $\lambda \in \mathbb{N}$,

$$\Pr \left[\begin{array}{l} \mathcal{A} \left(\begin{array}{l} \Pi, \{g_i\}_{i \in [3]}, T_b, \\ g_1^a, g_1^b g_3^b, g_1^c, g_1^{ac} g_3^d \end{array} \right) = b : \\ \begin{array}{l} ((p_1, p_2, p_3), \Pi) \leftarrow \text{Gen}(1^\lambda, 3); \Pi = (N, \mathbb{G}, \mathbb{G}_T, e) \\ g_1 \leftarrow \mathbb{G}_1; g_2 \leftarrow \mathbb{G}_2; g_3 \leftarrow \mathbb{G}_3 \\ a, b, c, d \leftarrow \mathbb{Z}_N, T_0 \leftarrow e(g_1, g_1)^{abc}; T_1 \leftarrow \mathbb{G}_T; b \leftarrow \{0, 1\} \end{array} \end{array} \right] \leq 1/2 + \text{negl}(\lambda).$$

2.3 Basic Tools

Pseudorandom Functions. A pseudorandom function (PRF) family $\mathcal{F} = \{\text{PRF}^K\}_{K \in \mathcal{K}}$ with a key space \mathcal{K} , a domain \mathcal{X} , and a range \mathcal{Y} is a function family that consists of functions $\text{PRF}^K : \mathcal{X} \rightarrow \mathcal{Y}$.

Definition 2.8. Let \mathcal{R} be a set of functions consisting of all functions whose domain and range are \mathcal{X} and \mathcal{Y} respectively. A PRF family \mathcal{F} is said to be secure if for any PPT adversary \mathcal{A} , the following condition holds,

$$|\Pr[\mathcal{A}^{\text{PRF}^K(\cdot)}(1^\lambda) = 1] - \Pr[\mathcal{A}^{\mathcal{R}(\cdot)}(1^\lambda) = 1]| \leq \text{negl}(\lambda),$$

where $K \leftarrow \mathcal{K}$ and $R \leftarrow \mathcal{R}$.

Non-interactive key exchange (NIKE). A NIKE scheme for shared key space \mathcal{K} consists of the three algorithms.

$\text{Setup}(1^\lambda) \rightarrow \text{PP}$. It takes a security parameter 1^λ and outputs a public parameter PP .

$\text{KeyGen}(\text{PP}) \rightarrow (\text{PK}, \text{SK})$. It takes PP and outputs a public key PK and the corresponding secret key SK .

$\text{SharedKey}(\text{PK}, \text{SK}) \rightarrow \text{K}$. It takes PK and SK and deterministically outputs a shared key $\text{K} \in \mathcal{K}$.

Correctness. A NIKE scheme is correct if, for all $\lambda \in \mathbb{N}$, we have

$$\Pr \left[\begin{array}{l} \text{PP} \leftarrow \text{Setup}(1^\lambda) \\ (\text{PK}_i, \text{SK}_i), (\text{PK}_j, \text{SK}_j) \leftarrow \text{KeyGen}(\text{PP}) \\ \text{K}_{i,j} = \text{SharedKey}(\text{PK}_i, \text{SK}_j) \\ \text{K}_{j,i} = \text{SharedKey}(\text{PK}_j, \text{SK}_i) \end{array} \right] = 1.$$

Security. We say a NIKE scheme is IND-secure if, for all stateful PPT adversaries \mathcal{A} , we have

$$\Pr \left[\begin{array}{l} \beta \leftarrow \{0, 1\}, \text{PP} \leftarrow \text{Setup}(1^\lambda) \\ \mathcal{S} \leftarrow \mathcal{A}(\text{PP}) \\ (\text{PK}_i, \text{SK}_i) \leftarrow \text{KeyGen}(\text{PP}) \\ \mathcal{CS}, (i', j') \leftarrow \mathcal{A}(\{\text{PK}_i\}_{i \in \mathcal{S}}) \text{ where } i', j' \in \mathcal{S} \setminus \mathcal{CS} \text{ and } i' \neq j' \\ \text{K}_{i',j'}^0 = \text{SharedKey}(\text{PK}_{i'}, \text{SK}_{j'}), \text{K}_{i',j'}^1 \leftarrow \mathcal{K} \\ \beta' \leftarrow \mathcal{A}(\{\text{SK}_i\}_{i \in \mathcal{CS}}, \text{K}_{i',j'}^\beta) \end{array} \right] \leq 1/2 + \text{negl}(\lambda).$$

3 Multi-Party Functional Encryption

In this section, we define our notion of multi-party functional encryption (MPFE). Let n_x be the number of ciphertext inputs and n_y be the number of key inputs. Let $\mathcal{X} = \mathcal{X}_{\text{pub}} \times \mathcal{X}_{\text{pri}}$ be the space of ciphertext inputs and $\mathcal{Y} = \mathcal{Y}_{\text{pub}} \times \mathcal{Y}_{\text{pri}}$ be the space of key inputs. We define two aggregation functions as $\text{Agg}_x : \mathcal{X}^{n_x} \rightarrow \mathcal{X}^*$, and $\text{Agg}_y : \mathcal{Y}^{n_y} \rightarrow \mathcal{Y}^*$.

An MPFE scheme is defined as a tuple of 4 algorithms/protocols $\text{MPFE} = (\text{Setup}, \text{KeyGen}, \text{Encrypt}, \text{Decrypt})$. To suitably capture existing primitives, we define our **Setup** algorithm to run in three modes, described next.

Setup modes. The **Setup** algorithm can be run in different modes: central, local, or interactive. For $\text{mode} \in \{\text{Central}, \text{Local}, \text{Interactive}\}$, consider the following.

Central : Here the **Setup** algorithm is run by one trusted third party which outputs the master secret keys and encryption keys for all users in the system.

Local : Here it is run independently by different parties without any interaction, and each party outputs its own encryption key and/or master secret key.

Interactive : Here it is an interactive protocol run by a set of users, at the end of which, each user has its encryption key and/or master secret key. We note that these keys may be correlated across multiple users.

A multi-party functional encryption (MPFE) consists of the following:

Setup ($1^\lambda, n_x, n_y, \text{Agg}_x, \text{Agg}_y$): This algorithm can be executed in any one of the three modes described above.⁴ Given input the security parameter, number of ciphertext inputs n_x , number of key inputs n_y and two aggregation functions $\text{Agg}_x, \text{Agg}_y$ as defined above, this algorithm outputs a set of encryption keys $\{\text{EK}_i\}_{i \leq n_x}$, master secret keys $\{\text{MSK}_i\}_{i \leq n_y}$ and public key PK.

KeyGen (PK, MSK, $i, y = (y_{\text{pub}}, y_{\text{pri}})$): Given input the public key PK, a master secret key MSK, user index $i \in [n_y]$ and a function input $y = (y_{\text{pub}}, y_{\text{pri}})$, this algorithm outputs a secret key SK_y .

Encrypt (PK, EK, $i, x = (x_{\text{pub}}, x_{\text{pri}})$): Given input the public key PK, an encryption key EK, user index $i \in [n_x]$, an input $x = (x_{\text{pub}}, x_{\text{pri}})$, this algorithm outputs a ciphertext CT_x .

Decrypt (PK, $\{\text{SK}_i\}_{i \leq n_y}, \{\text{CT}_j\}_{j \leq n_x}$): Given input the public key PK, a set of secret keys $\{\text{SK}_i\}_{i \leq n_y}$ and a set of ciphertexts $\{\text{CT}_j\}_{j \leq n_x}$, this algorithm outputs a value z or \perp .

We remark that in the *local* setup mode, it will be helpful to separate the setup algorithm into a global setup, denoted by **GSetup** along with a local setup, denoted by **LSetup**, where the former is used only to generate common parameters of the system, such as group descriptions and such.

Correctness. We say that an MPFE scheme is *correct* if, $\forall (n_x, n_y) \in \mathbb{N}^2$, ciphertext inputs $x_i \in \mathcal{X}$ for $i \in [n_x]$, key inputs $y_j \in \mathcal{Y}$ for $j \in [n_y]$, message and function aggregation circuits Agg_x and Agg_y , it holds that:

$$\Pr \left[\begin{array}{l} (\text{PK}, \{\text{EK}_i\}, \{\text{MSK}_j\}) \leftarrow \text{Setup}(1^\lambda, n_x, n_y, \text{Agg}_x, \text{Agg}_y) \\ \text{CT}_i \leftarrow \text{Encrypt}(\text{PK}, \text{EK}_i, i, x_i) \quad \forall i \in [n_x] \\ \text{SK}_j \leftarrow \text{KeyGen}(\text{PK}, \text{MSK}_j, j, y_j) \quad \forall j \in [n_y] \\ z \leftarrow \text{Decrypt}(\text{PK}, \{\text{SK}_j\}_{j \leq n_y}, \{\text{CT}_i\}_{i \leq n_x}) \\ z' = \mathcal{U}(\text{Agg}_x(\{x_i\}), \text{Agg}_y(\{y_j\})) \end{array} \right] = 1.$$

Recall that \mathcal{U} is the universal circuit with appropriate input and output size.

⁴We omit specifying the mode in the syntax for notational brevity.

Indistinguishability based security. Next, we define security of MPFE. The security definition is modelled in a similar fashion to MIFE security [GGG⁺14, §2.2] while taking into account corruption queries.

For any choice of parameters λ, n_x, n_y , aggregation functions $\text{Agg}_x, \text{Agg}_y$, and master keys $\mathbf{K} = (\text{PK}, \{\text{EK}_i\}_{i \in [n_x]}, \{\text{MSK}_j\}_{j \in [n_y]}) \leftarrow \text{Setup}(1^\lambda, n_x, n_y, \text{Agg}_x, \text{Agg}_y)$, we define the following list of oracles:

$\text{Corrupt}^{\mathbf{K}}(\cdot)$, upon a call to this oracle for any $i \in [n_x]$ or $j \in [n_y]$, the adversary gets the corresponding encryption key EK_i or master secret key MSK_j . (Let $\mathcal{S}_x \subseteq [n_x]$ and $\mathcal{S}_y \subseteq [n_y]$ denote the set of user indices for which the corresponding encryption and master keys have been corrupted.)⁵

$\text{Key}^{\mathbf{K}, \beta}(\cdot, \cdot)$, upon a call to this oracle for an honest user index $j \in [n_y]$, function inputs $(y_j^{k,0}, y_j^{k,1})$ (where $y_j^{k,b} = (y_{j,\text{pub}}^{k,b}, y_{j,\text{pri}}^{k,b})$ for $b \in \{0,1\}$), the challenger first checks whether the user j was already corrupted or not. That is, if $j \in \mathcal{S}_y$, then it sends nothing, otherwise it samples a decryption key for function input $y_j^{k,\beta}$ using key MSK_j and sends it to the adversary. (Here β is the challenge bit chosen at the start of the experiment.)

$\text{Enc}^{\mathbf{K}, \beta}(\cdot, \cdot)$, upon a call to this oracle for an honest user index $i \in [n_x]$, message inputs $(x_i^{\ell,0}, x_i^{\ell,1})$ (where $x_i^{\ell,b} = (x_{i,\text{pub}}^{\ell,b}, x_{i,\text{pri}}^{\ell,b})$ for $b \in \{0,1\}$), the challenger first checks whether the user i was already corrupted or not. That is, if $i \in \mathcal{S}_x$, then it sends nothing, otherwise it samples a ciphertext for input $x_i^{\ell,\beta}$ using key EK_i and sends it to the adversary.

We let Q_x and Q_y be the number of encryption and key generation queries (respectively) that had non-empty responses. Let $\mathcal{Q}_x = \{(i, (x_i^{\ell,0}, x_i^{\ell,1}))\}_{\ell \in [Q_x]}$ be the set of ciphertext queries and $\mathcal{Q}_y = \{(j, (y_j^{k,0}, y_j^{k,1}))\}_{k \in [Q_y]}$ be the set of key queries.

We say that an adversary \mathcal{A} is *admissible* if:

1. For each of the encryption and key challenges, the public components of the two challenges are equal, namely $x_{\text{pub}}^{\ell,0} = x_{\text{pub}}^{\ell,1}$ for all $\ell \in [Q_x]$, and $y_{\text{pub}}^{k,0} = y_{\text{pub}}^{k,1}$ for all $k \in [Q_y]$.
2. For each of the encryption and key challenges, the *private* components of the two challenges are also equal, namely $x_{\text{pri}}^{\ell,0} = x_{\text{pri}}^{\ell,1}$ for all $\ell \in [Q_x]$ whenever $(i, (x_i^{\ell,0}, x_i^{\ell,1})) \in \mathcal{Q}_x$ and $i \in \mathcal{S}_x$, and $y_{\text{pri}}^{k,0} = y_{\text{pri}}^{k,1}$ for all $k \in [Q_y]$ whenever $(j, (y_j^{k,0}, y_j^{k,1})) \in \mathcal{Q}_y$ and $j \in \mathcal{S}_y$. That is, the private components must be the same as well if the user index i or j , that the query was made for, was corrupted during the execution.
3. There do not exist two sequences (\vec{x}^0, \vec{y}^0) and (\vec{x}^1, \vec{y}^1) such that:

$$\mathcal{U}(\text{Agg}_x(\{x_i^0\}), \text{Agg}_y(\{y_j^0\})) \neq \mathcal{U}(\text{Agg}_x(\{x_i^1\}), \text{Agg}_y(\{y_j^1\}))$$

and i) for every $i \in [n_x]$, either x_i^b was queried or EK_i was corrupted, and ii) for every $j \in [n_y]$, either y_j^b was queried or MSK_j was corrupted, and iii) at least one of inputs $\{x_i^b\}, \{y_j^b\}$ were queried and indices i, j were not corrupted. (Note that if $i \in [n_x]$ or $j \in [n_y]$ were queried to the Corrupt oracle, the adversary can generate partial keys or ciphertexts for any value of its choice.)

An MPFE scheme ($\text{Setup}, \text{KeyGen}, \text{Encrypt}, \text{Decrypt}$) is said to be IND secure if for any *admissible* PPT adversary \mathcal{A} , all length parameters $n_x, n_y \in \mathbb{N}$, and aggregation functions $\text{Agg}_x, \text{Agg}_y$, there exists a negligible function $\text{negl}(\cdot)$ such that for all $\lambda \in \mathbb{N}$, the following holds

$$\Pr \left[\mathcal{A}^{\text{Corrupt}^{\mathbf{K}}(\cdot), \text{Key}^{\mathbf{K}, \beta}(\cdot), \text{Enc}^{\mathbf{K}, \beta}(\cdot)}(1^\lambda, \text{PK}) = \beta : \begin{array}{l} \mathbf{K} \leftarrow \text{Setup}(1^\lambda, n_x, n_y, \text{Agg}_x, \text{Agg}_y), \\ \mathbf{K} = (\text{PK}, \{\text{EK}_i\}_i, \{\text{MSK}_j\}_j), \\ \beta \leftarrow \{0,1\} \end{array} \right] \leq \frac{1}{2} + \text{negl}(\lambda).$$

⁵Note that in case EK_i is completely contained in some MSK_j then make a master secret corruption query for j will also add the corresponding index i to \mathcal{S}_x , and vice versa. At a very high level, although having separate aggregation functions for partial secret key and ciphertexts as part of the framework allows us to capture a highly expressive class of encryption scheme; defining the most general notion of security for MPFE that captures all different types of setup and key distribution settings is very cumbersome. To that end, we provide a clean security game above which captures the existing encryption primitives, but we point out that a better mechanism for capturing security would be specialize for each setup mode and corruption model.

Remark 3.1 (Weaker notions of security). We say the scheme is selective IND secure if the adversary outputs the challenge message and function pairs at the very beginning of the game, before it makes any queries or receives the PK. One may also consider the semi-honest setting, where the **Corrupt** oracle is not provided, or the case of static corruptions where the adversary provides all its corruptions once and for all at the start of the game.

3.1 Dynamic Multi-Party Functional Encryption

In this section, we define the dynamic notion of multi-party functional encryption (MPFE). We consider the fully dynamic setting where the number of key/ciphertext inputs is unspecified during setup time, and the aggregation functions are also specified only during key generation and encryption times. In the dynamic setting, an interactive or centralized setup is not meaningful since the number of parties is itself not known during setup time, hence we restrict ourselves to the local setup mode for simplicity.

Let $\mathcal{X} = \mathcal{X}_{\text{pub}} \times \mathcal{X}_{\text{pri}}$ be the space of ciphertext inputs and $\mathcal{Y} = \mathcal{Y}_{\text{pub}} \times \mathcal{Y}_{\text{pri}}$ be the space of key inputs. A dynamic multi-party functional encryption scheme (MPFE) with local setup is defined as a tuple of 5 algorithms/protocols $\text{MPFE} = (\text{GSetup}, \text{LSetup}, \text{KeyGen}, \text{Encrypt}, \text{Decrypt})$ with the following syntax:

GSetup(1^λ): On input the security parameter, the global setup algorithm samples a globally shared set of public parameters PP .

LSetup(PP): Given input the public parameters, the local setup algorithm outputs a tuple consisting of local public key PK , an encryption key EK , and a master secret key MSK . (Here the local public key is just regarded as a public identifier for the user, and not given as explicit input to other algorithms since it could always be added to the encryption and/or master secret key.)

KeyGen ($\text{MSK}, j, y = (y_{\text{pub}}, y_{\text{pri}}), \text{Agg}_y$): Given input a master secret key MSK , user index $j \in [n_y]$ ⁶, a function input $y = (y_{\text{pub}}, y_{\text{pri}})$, and an aggregation function $\text{Agg}_y : \mathcal{Y}^{n_y} \rightarrow \mathcal{Y}^*$ (for some $n_y \in \mathbb{N}$), this algorithm outputs a secret key SK_j .

Encrypt ($\text{EK}, i, x = (x_{\text{pub}}, x_{\text{pri}}), \text{Agg}_x$): Given input an encryption key EK , user index $i \in [n_x]$, an input $x = (x_{\text{pub}}, x_{\text{pri}})$, and an aggregation function $\text{Agg}_x : \mathcal{X}^{n_x} \rightarrow \mathcal{X}^*$ (for some $n_x \in \mathbb{N}$), this algorithm outputs a ciphertext CT_i .

Decrypt ($\{\text{SK}_i\}_i, \{\text{CT}_j\}_j$): Given input a sequence of secret keys $\{\text{SK}_i\}_i$ and a sequence of ciphertexts $\{\text{CT}_j\}_j$, this algorithm outputs a value z or \perp .

Correctness. We say that an MPFE scheme is *correct* if, $\forall (N, n_x, n_y) \in \mathbb{N}^3$, ciphertext inputs $x_i \in \mathcal{X}$ for $i \in [n_x]$, key inputs $y_j \in \mathcal{Y}$ for $j \in [n_y]$, message and function aggregation circuits Agg_x and Agg_y , and indexing functions $\text{index}_x : [n_x] \rightarrow [N]$, $\text{index}_y : [n_y] \rightarrow [N]$ it holds that:

$$\Pr \left[\begin{array}{l} \text{PP} \leftarrow \text{GSetup}(1^\lambda) \\ (\text{PK}_\ell, \text{EK}_\ell, \text{MSK}_\ell) \leftarrow \text{LSetup}(\text{PP}) \quad \forall \ell \in [N] \\ \text{CT}_{\text{index}_x(i)} \leftarrow \text{Encrypt}(\text{EK}_{\text{index}_x(i)}, i, x_i, \text{Agg}_x) \quad \forall i \in [n_x] \\ \text{SK}_{\text{index}_y(j)} \leftarrow \text{KeyGen}(\text{MSK}_{\text{index}_y(j)}, j, y_j, \text{Agg}_y) \quad \forall j \in [n_y] \\ z \leftarrow \text{Decrypt}(\{\text{SK}_j\}_{j \leq n_y}, \{\text{CT}_i\}_{i \leq n_x}) \\ z' = \mathcal{U}(\text{Agg}_x(\{x_i\}), \text{Agg}_y(\{y_j\})) \end{array} \right] = 1.$$

Recall that \mathcal{U} is the universal circuit with appropriate input and output size.

⁶Note that both the key generator and encryptor take as input the user index as well. This is not an additional requirement as in our current definition we are modeling the aggregation functions as a circuit, and thus both the key generator and encryptor could explicitly provide at which slot in the aggregation circuit do they want their inputs to be read. As we remark later in Remark 3.2, if we model the aggregation function in the uniform computation model, then we could avoid providing the user index as an input to the key generation and encryption algorithms.

Indistinguishability based security. Here we extend the security experiment for multi-party functional encryption that we provided in Section 3 to the dynamic user setting in the local setup mode. Since we are working in the dynamic setting, we need to define the following oracles⁷:

HonestGen(), upon a call to this oracle, the challenger samples a fresh tuple of local public key, encryption key, and master key (PK, EK, MSK), and stores them in a list L_{setup} . It sends PK to the adversary. (Note that if the scheme is a public key scheme, then the challenger sends the encryption key EK to the adversary.)

Corrupt(\cdot, \cdot), upon a call to this oracle for an honest user local public key PK and key type $\text{type} \in \{\text{enc}, \text{master}\}$, the challenger first checks whether the list L_{setup} contains a key pair associated with PK. If there is such a key pair (PK, EK, MSK), then it sends either the EK or MSK depending on the type queried. Otherwise, it sends nothing.⁸

Key $^\beta(\cdot, \cdot, \cdot, \cdot)$, upon a call to this oracle for an honest user local public key PK, function inputs $(y_j^{k,0}, y_j^{k,1})$ (where $y_j^{k,b} = (y_{j,\text{pub}}^{k,b}, y_{j,\text{pri}}^{k,b})$ for $b \in \{0, 1\}$), index j , aggregation function $\text{Agg}_{y,j}^k$, the challenger first checks whether the list L_{setup} contains a key pair associated with PK. If there is such a key pair (PK, EK, MSK), then it samples a decryption key for function input $y_j^{k,\beta}$ using key MSK and sends it to the adversary. Otherwise, it sends nothing. (Here β is the challenge bit chosen at the start of the experiment.)

Enc $^\beta(\cdot, \cdot, \cdot, \cdot)$, upon a call to this oracle for an honest user local public key PK, inputs $(x_j^{\ell,0}, x_j^{\ell,1})$ (where $x_j^{\ell,b} = (x_{j,\text{pub}}^{\ell,b}, x_{j,\text{pri}}^{\ell,b})$ for $b \in \{0, 1\}$), index j , aggregation function $\text{Agg}_{x,j}^\ell$, the challenger first checks whether the list L_{setup} contains a key pair associated with PK. If there is such a key pair (PK, EK, MSK), then it samples a ciphertext for input $x_j^{\ell,\beta}$ using key EK and sends it to the adversary. Otherwise, it sends nothing. (Here β is the challenge bit chosen at the start of the experiment.)

We let Q_x and Q_y be the number of encryption and key generation queries (respectively) that had non-empty responses. Let $\mathcal{Q}_x = \{(\text{PK}^\ell, (x_j^{\ell,0}, x_j^{\ell,1}), j, \text{Agg}_{x,j}^\ell)\}_{\ell \in [Q_x]}$ be the set of ciphertext challenge queries and $\mathcal{Q}_y = \{(\text{PK}^k, (y_j^{k,0}, y_j^{k,1}), j, \text{Agg}_{y,j}^k)\}_{k \in [Q_y]}$ be the set of key challenge queries.

We say that an adversary \mathcal{A} is *admissible* if:

1. For each of the encryption and key challenges, the public components of the two challenges are equal, namely $x_{j,\text{pub}}^{\ell,0} = x_{j,\text{pub}}^{\ell,1}$ for all $\ell \in [Q_x]$, and $y_{j,\text{pub}}^{k,0} = y_{j,\text{pub}}^{k,1}$ for all $k \in [Q_y]$.
2. For each of the encryption and key challenges, the *private* components of the two challenges are also equal, namely $x_{j,\text{pri}}^{\ell,0} = x_{j,\text{pri}}^{\ell,1}$ for all $\ell \in [Q_x]$, and $y_{j,\text{pri}}^{k,0} = y_{j,\text{pri}}^{k,1}$ for all $k \in [Q_y]$ if the encryption key EK^ℓ or the master secret key MSK^k , that the query was made for, was corrupted during the execution (respectively).
3. There do not exist two sequences $(\vec{x}^0, \vec{y}^0) \neq (\vec{x}^1, \vec{y}^1)$ and aggregation functions $\text{Agg}_x, \text{Agg}_y$ such that:

$$\mathcal{U}(\text{Agg}_x(\{x_j^0\}), \text{Agg}_y(\{y_j^0\})) \neq \mathcal{U}(\text{Agg}_x(\{x_j^1\}), \text{Agg}_y(\{y_j^1\}))$$

and i) x_j^b was queried for aggregation function Agg_x and index j , and ii) y_j^b was queried for aggregation function Agg_y and index j , and iii) at least one of inputs $\{x_j^b\}, \{y_j^b\}$ were queried and index j was not corrupted. Note that if some x_j^b or y_j^b was not queried by the adversary, then it can generate partial keys or ciphertexts for any value of its choice by performing a fresh key generation since this is a fully dynamic system.

⁷Note that all these oracles have access to all other oracle's state.

⁸As we point out in the static setting, in case EK_i is completely contained in some MSK_i (or vice versa), then making a master secret corruption query for i will also imply that encryption key for i has been corrupted too (and vice versa).

An MPFE scheme (GSetup, LSetup, KeyGen, Encrypt, Decrypt) is said to be IND secure if for any *admissible* PPT adversary \mathcal{A} , there exists a negligible function $\text{negl}(\cdot)$ such that for all $\lambda \in \mathbb{N}$, the following holds

$$\Pr \left[\mathcal{A}^{\text{HonestGen}(), \text{Corrupt}(), \text{Key}^\beta(), \text{Enc}^\beta()}(1^\lambda) = \beta : \beta \leftarrow \{0, 1\} \right] \leq \frac{1}{2} + \text{negl}(\lambda).$$

Remark 3.2 (Potential variations). The above multi-party function encryption system that we define allows the users to dynamically join the system in the permissionless model, where each incoming user only needs to know the public parameters and not interact with any authority. A slightly weaker setting could be a permissioned model in which users can still dynamically join the system but they need to contact the global authority (which sampled the public parameters) either for some identification tokens or its encryption and master secret key pair in order to prevent totally unrestricted computation which happens in the permissionless model.

Also, we want to point out that in our current framework we let the users select the aggregation functions during individual functional key and ciphertext generation to allow for more flexibility. This could be relaxed even further by letting the aggregation functions be either be described in a uniform computation model, or using an ensemble of non-uniform functions. Also, one could instead restrict the flexibility in aggregation by asking each user to choose their aggregation functions at setup time. Such flexibilities will be important in capturing the notion of DDFE described in Section 4.

4 Capturing Existing Primitives as MPFE

We now show that our model is general enough to capture several existing notions in the literature as special cases. To simplify notation, we assume that every input is associated with an index which indicates its position in the set of arguments to the universal circuit \mathcal{U} . Here, we assume that the input to \mathcal{U} which represents a circuit that must be emulated, is always associated with index 1 in either the first argument or the second argument.

Multi-Input Functional Encryption (MIFE): A Multi-Input Functional Encryption scheme, denoted by MIFE [GGG+14], can be viewed as a special case of MPFE with a *centralized* setup. In an MIFE scheme, the setup algorithm outputs n encryption keys $\text{EK}_1, \dots, \text{EK}_n$ and a single MSK. Each encryptor $i \in [n]$, chooses its message z_i and computes a ciphertext of z_i using EK_i . Here, the entire z_i is private. The key generator chooses a function f and computes a corresponding secret key SK_f using its MSK. Here, f is public.

In our notation, MIFE is captured by setting $(n_x, n_y) = (n, 1)$, $y = (f, \perp)$, $x_i = (x_{\text{pub},i}, x_{\text{pri},i}) = (\perp, z_i)$. We also set $\text{Agg}_x(x_1, \dots, x_n) = (x_{\text{pri},1}, \dots, x_{\text{pri},n_x}) = (z_1, \dots, z_n)$. Finally, Agg_y outputs $y_{\text{pub}} = f$ and upon decryption this yields $\mathcal{U}((z_1, \dots, z_n), f) = f(z_1, \dots, z_n)$ as desired.

Multi-Client Functional Encryption (MCFE): A Multi-Client Functional Encryption scheme, denoted by MCFE [GKL+13, GGG+14, CDSG+18b] can be viewed as a special case of MPFE with a *centralized* setup. In an MCFE scheme, the setup algorithm outputs n encryption keys $\text{EK}_1, \dots, \text{EK}_n$ and a single MSK. Each encryptor $i \in [n]$, chooses its message z_i and a *public label* lab_i and computes a ciphertext for (lab_i, z_i) using EK_i . Here, lab_i is public but z_i is private. The key generator chooses a function f and computes a corresponding secret key SK_f using its MSK. Here, f is public.

In our notation, set $(n_x, n_y) = (n, 1)$, $y = (f, \perp)$, $x_i = (x_{\text{pub},i}, x_{\text{pri},i}) = (\text{lab}_i, z_i)$. The function $\text{Agg}_x(x_1, \dots, x_n)$ checks that $x_{\text{pub},1} = \dots = x_{\text{pub},n}$. If this verification passes, it outputs $(x_{\text{pri},1}, \dots, x_{\text{pri},n_x}) = (z_1, \dots, z_n)$. Finally, Agg_y outputs $y_{\text{pub}} = f$ and upon decryption this yields $\mathcal{U}((z_1, \dots, z_n), f) = f(z_1, \dots, z_n)$ as desired.

Distributed Multi-Client Functional Encryption (D-MCFE): The decentralized version of MCFE, denoted by D-MCFE [CDSG+18a], can be viewed as a special case of MPFE with an *interactive* setup. This primitive removes the need for a centralized authority by shifting the task of generating functional secret

keys to the clients themselves. In D-MCFE, during the setup phase, the users generate public parameters and their individual master secret keys by running an interactive protocol. No further interaction is needed among clients for subsequent generation of functional keys. When an evaluator wishes to obtain a functional key for some function f , it requests each user for a partial decryption key corresponding to f . Ciphertexts for inputs (lab_i, z_i) are generated by each user independently as in MCFE. The decryptor, upon receiving ciphertexts and partial decryption keys from all the parties can compute $f(z_1, \dots, z_n)$ as in MCFE, so long as all the labels match.

In our notation, set $(n_x, n_y) = (n, n)$ and $y_i = (f, \perp)$, $x_i = (x_{\text{pub},i}, x_{\text{pri},i}) = (\text{lab}_i, z_i), \forall i \in [n]$. As in MCFE, $\text{Agg}_x(x_1, \dots, x_n)$ checks that $x_{\text{pub},1} = \dots = x_{\text{pub},n}$. If this verification passes, it outputs $(x_{\text{pri},1}, \dots, x_{\text{pri},n_x}) = (z_1, \dots, z_n)$. Next, Agg_y first checks that $y_{\text{pub},i} = f, \forall i \in [n]$, and if so, outputs the corresponding value $y_{\text{pub},1} = f$. Upon decryption, this yields $f(z_1, \dots, z_n)$ as in MCFE. In this primitive, $\text{MSK}_i = \text{EK}_i$ for all $i \in [n]$.

Multi-Authority Functional or Attribute-Based Encryption (MAFE or MA-ABE): A Multi-Authority Functional Encryption, denoted by MAFE (or its special case Multi-Authority Attribute Based Encryption, MA-ABE) [Cha07, LW11, BCG⁺17] can be viewed as an instance of MPFE with a *local* setup. In MAFE, there are n key authorities who may independently generate their private and public keys, without even being aware of the existence of other authorities. MAFE is a ciphertext-policy scheme, so an encryptor encrypts a message z along with a policy f over the various authorities. Here, z is private but f is public. Any authority i (say), should be able to generate a token for a user with a global identifier GID for attributes lab_i . A user with tokens for attribute-identifier pair $(\text{lab}_i, \text{GID}_i)$ from authority $i \in [n]$, should be able to decrypt the ciphertext to recover $f(z, \text{lab}_1, \dots, \text{lab}_n)$ as long as all the n identifiers are the same.

In our notation, we set $(n_x, n_y) = (1, n)$, $x = (f, z)$, $y_i = (y_{\text{pub},i}, y_{\text{pri},i}) = ((\text{lab}_i, \text{GID}_i), \perp)$. The function Agg_x simply outputs the pair (f, z) . The aggregation function Agg_y checks that $\text{GID}_1 = \dots = \text{GID}_n$. If this verification passes, it outputs $(\text{lab}_1, \dots, \text{lab}_n)$. Decryption computes $\mathcal{U}(f, z, (\text{lab}_1, \dots, \text{lab}_n)) = f(z, \text{lab}_1, \dots, \text{lab}_n)$ as desired.

Decentralized Policy-Hiding Attribute Based Encryption (DABE): A decentralized policy-hiding ABE, denoted by DABE [MJ18] proposed by Michalevsky and Joye can also be captured by an MPFE scheme with local setup. In a DABE scheme, there are n key authorities, each of which run a local setup to generate their private and public keys. They do so without communicating with each other, although they are required to obtain some global parameters of the system (such as group generators and such), which are generated by a one-time trusted setup. This is a ciphertext-policy scheme, in which an encryptor can compute a ciphertext under a general access structure, while the corresponding secret keys are issued by independent authorities as in MAFE. However, in this system, the access policy embedded in the ciphertext is hidden.

In our notation, we set $(n_x, n_y) = (1, n)$, $x = (\perp, (f, z))$, $y_i = (\text{lab}_i, \perp)$. The aggregation function Agg_y outputs $(\text{lab}_1, \dots, \text{lab}_n)$. The aggregation function Agg_x returns (f', z) where $f'(z, \cdot) = (f, z)$ iff $f(\text{lab}_1, \dots, \text{lab}_n) = 1$ and \perp otherwise. The decryptor computes $\mathcal{U}(f', z, \text{lab}_1, \dots, \text{lab}_n) = f'(z, \text{lab}_1, \dots, \text{lab}_n)$ which checks if $f(\text{lab}_1, \dots, \text{lab}_n) = 1$ and outputs (f, z) if and only if this is the case, as desired.

Partially Hiding Functional/Predicate Encryption (PHFE and PHPE): Partially Hiding Functional/Predicate Encryption (PHFE or PHPE) [GVW12, GVW15] can be seen as a special case of MPFE with a centralized setup algorithm in the public-key setting. This is a single user scheme, in which the setup outputs a PK and MSK. The encryptor, given the public key PK chooses a public label lab and a private input z and computes a ciphertext for the pair (lab, z) . The key generator, given the master secret key, computes a secret key for some function f , which is public. The decryptor computes $f(\text{lab}, z)$.

This scheme can be expressed as an MPFE scheme with $(n_x, n_y) = (1, 1)$, $x = (\text{lab}, z)$, $y = (f, \perp)$. The Agg_x and Agg_y functions are trivial in this setting, since it is a single user scheme, and simply output their inputs. The decryptor computes $\mathcal{U}(f, \text{lab}, z) = f(\text{lab}, z)$ as desired.

Ad Hoc Multi-Input Functional Encryption (aMIFE): The primitive of adhoc multi-input functional encryption (aMIFE) recently proposed by Agrawal et al. [ACF+20] can be instantiated as an MPFE scheme with local setup. Adhoc MIFE enables multiple parties to generate their own private master keys as well as corresponding public keys. To encrypt data, each party (say i) uses its private master key to compute a ciphertext for any message of its choice (say z_i), and to issue function specific keys, the party uses the same master key to compute a partial function key corresponding to some function (say f). These ciphertexts and partial keys are sent to the decryptor who can aggregate them and compute $f(z_1, \dots, z_n)$ as long as all the partial decryption keys correspond to the same function f .

In our notation, we set $(n_x, n_y) = (n, n)$, $x_i = (\perp, z_i)$, $y_i = (f_i, \perp)$. The function Agg_x outputs (z_1, \dots, z_n) , the function Agg_y checks that all $y_{\text{pub},1} = f_1 = \dots = y_{\text{pub},n} = f_n$ are equal to the same value f , which it outputs. The decryptor computes $\mathcal{U}(f, z_1, \dots, z_n) = f(z_1, \dots, z_n)$ as desired.

Dynamic Decentralized Functional Encryption (DDFE): A dynamic decentralized FE scheme, denoted by DDFE [CDSG+20], can be viewed as a special case of MPFE with a local setup in the dynamic setting. In a DDFE scheme for functionality F and empty key ϵ , the setup outputs the public parameters PP, and each party samples its own public and secret key pair PK, SK. Each encryptor on input a public key PK (and the secret key SK if the DDFE scheme is secret key) and a message m , outputs a ciphertext CT_{PK} . The key generator chooses a key space object k and computes a corresponding decryption key $\text{DK}_{\text{PK},k}$ using a user master secret key SK. Here, key space object k is public. The decryption algorithm takes as input a list of decryption keys $(\text{DK}_{\text{PK}_j, k_j})_j$ and ciphertexts $(\text{CT}_{\text{PK}'_i})_i$, and outputs $F((\text{PK}_j, k_j)_j, (\text{PK}'_i, m_i)_i)$. Also, the system enables learning the empty function on a possibly non-singleton list of ciphertexts as: $F(\epsilon, (\text{PK}'_i, m_i)_i)$.

To formally capture DDFE, we require the generalized dynamic MPFE, which is discussed in Section 3.1, which allows to specify (n_x, n_y) as well as $(\text{Agg}_x, \text{Agg}_y)$ dynamically. To set the remaining parameters in our notation, we let $y_j = ((F, \text{PK}_j, k_j), \perp)$, $x_i = (\perp, (\text{PK}'_i, m_i))$. The function Agg_x outputs $((\text{PK}'_1, m_1), \dots, (\text{PK}'_{n_x}, m_{n_x}))$ for some $n_x \in \mathbb{N}$, and Agg_y checks that $F_1 = \dots = F_{n_y} = F$ is the same in all y_j and outputs $(F, (\text{PK}_1, k_1), \dots, (\text{PK}_{n_y}, k_{n_y}))$ if so (\perp otherwise) for some $n_y \in \mathbb{N}$. Upon decryption this yields $\mathcal{U}(((\text{PK}'_1, m_1), \dots, (\text{PK}'_{n_x}, m_{n_x})), (F, (\text{PK}_1, k_1), \dots, (\text{PK}_{n_y}, k_{n_y}))) = F((\text{PK}_j, k_j)_j, (\text{PK}'_i, m_i)_i)$ as desired. And, note that if we consider the aggregation functions to be described as TMs or an ensemble of non-uniform functions Remark 3.2, then the key generator and does not need to specify the user indices during the key generation and encryption phase as desired in DDFE. Lastly, to enable learning the empty function ϵ , each user during their local setup also samples a secret key for input $y_j^\epsilon = ((F, \epsilon), \perp)$ and keeps it part of the public key and the ciphertext, and thus the decryption algorithm recovers the $\mathcal{U}(((\text{PK}'_1, m_1), \dots, (\text{PK}'_{n_x}, m_{n_x})), (F, \epsilon, \dots, \epsilon)) = F(\epsilon, (\text{PK}'_i, m_i)_i)$ as desired.

4.1 Feasibility of MPFE for General Circuits

In this section we discuss the feasibility of general MPFE. Since the framework of MPFE is very general and supports all existing constructions in the literature (that we are aware of), feasibility varies widely depending on the properties desired, such as whether the setup algorithm must be local, centralized or interactive, whether and encryption keys must be public or private and such others. Since the case of centralized setup has been most widely studied in the literature and it is evident that feasibility results for local and interactive setup also apply for a centralized setup, we focus on these below.

Local Setup. In multi-party FE schemes, local setup is highly desirable, since it overcomes the key escrow problem which is one of the main drawbacks of present day FE constructions. In this setting, each user i runs the setup algorithm independently, and obtains her own public key PK_i and master secret key MSK_i . No communication or co-ordination is required between the users. In the symmetric key setting, the user encrypts her input using her master key MSK_i , while in the public key setting, anyone may encrypt using the public key(s) PK_i .

Constructions of multi-party FE with local setup have been notoriously hard to build. For general circuits

and in the symmetric key setting, the primitive of adhoc multi-input functional encryption (aMIFE) recently proposed by Agrawal et al. [ACF⁺20] comes close to a general feasibility result in our model. Although there are important differences between aMIFE and MPFE, we show in Appendix A, that the construction of aMIFE provided in [ACF⁺20] is general enough to provide a feasibility result for MPFE in the symmetric key setting, with local setup.

Interactive Setup. A standard MIFE scheme can be modified to support distributed keys by replacing the setup algorithm with an interactive protocol. This makes use of a function delayed, rerunnable two round MPC protocol similar to the construction of aMIFE provided by [ACF⁺20]. Note that the limitation of interactive setup is mitigated by the fact that it must only be run once. This transformation also works in the public key setting. Please see Appendix A for details.

5 Multi-Authority ABE \circ IPFE for LSSS Access Structures

In this section, we build a multi-authority attribute-based inner product functional encryption (MA-AB-IPFE) scheme for linear secret sharing schemes (LSSS) access structures from simple assumptions over bilinear maps. Although our general framework of MPFE described in Section 3 is clearly powerful enough to capture this new functionality, we provide the precise syntax and definition here for ease of exposition.

Syntax. A MA-AB-IPFE scheme for predicate class $\mathcal{C} = \{C_n : \mathcal{X}_n \rightarrow \{0, 1\}\}_{n \in \mathbb{N}}$ and inner product message space $\mathcal{U} = \{\mathcal{U}_\ell\}_{\ell \in \mathbb{N}}$ consists of the following PPT algorithms:

GSetup(1^λ) \rightarrow PP. On input the security parameter λ , the setup algorithm outputs public parameters PP.

LSetup(PP, $1^n, 1^\ell, i$) \rightarrow (PK, MSK). On input the public parameters PP, attribute length n , message space index ℓ , and authority's index $i \in [n]$, the authority setup algorithm outputs a pair of master public-secret key (PK, MSK) for the i -th authority.

KeyGen(MSK $_j$, GID, b, \mathbf{u}) \rightarrow SK $_{j, \text{GID}, b, \mathbf{u}}$. The key generation algorithm takes as input the authority master secret key MSK $_j$, global identifier GID, an attribute bit $b \in \{0, 1\}$, and key vector $\mathbf{u} \in \mathcal{U}_\ell$. It outputs a partial secret key SK $_{j, \text{GID}, b, \mathbf{u}}$.

Enc($\{\text{PK}_i\}_{i \in [n]}$, C, \mathbf{v}) \rightarrow CT. The encryption algorithm takes as input the list of public keys $\{\text{PK}_i\}_i$, predicate circuit C , and a message vector $\mathbf{v} \in \mathcal{U}_\ell$, and outputs a ciphertext CT.

Dec($\{\text{SK}_{i, \text{GID}, \mathbf{x}, \mathbf{u}}\}_{i \in [n]}$, CT) $\rightarrow m/\perp$. On input a list of n partial secret keys $\{\text{SK}_{i, \text{GID}, \mathbf{x}, \mathbf{u}}\}_i$ and a ciphertext CT, the decryption algorithm either outputs a message m (corresponding to the inner product value) or a special string \perp (to denote decryption failure).

Correctness. A MA-AB-IPFE scheme is said to be correct if for all $\lambda, n, \ell \in \mathbb{N}$, $C \in \mathcal{C}_n$, $\mathbf{u}, \mathbf{v} \in \mathcal{U}_\ell$, $\mathbf{x} \in \mathcal{X}_n$, GID, if $C(\mathbf{x}) = 1$, the following holds:

$$\Pr \left[\begin{array}{l} \text{Dec}(\text{SK}, \text{CT}) = \langle \mathbf{u}, \mathbf{v} \rangle : \\ \text{PP} \leftarrow \text{GSetup}(1^\lambda) \\ \forall i \in [n] : (\text{PK}_i, \text{MSK}_i) \leftarrow \text{LSetup}(\text{PP}, 1^n, 1^\ell, i) \\ \forall j \in [n] : \text{SK}_{j, \text{GID}, x_j, \mathbf{u}} \leftarrow \text{KeyGen}(\{\text{PK}_i\}_i, \text{MSK}_j, \text{GID}, x_j, \mathbf{u}) \\ \text{CT} \leftarrow \text{Enc}(\{\text{PK}_i\}_i, C, \mathbf{v}), \text{SK} = \{\text{SK}_{i, \text{GID}, x_i, \mathbf{u}}\}_i \end{array} \right] = 1.$$

Security. In terms of security, a MA-AB-IPFE provides powerful notion of encrypted message vector indistinguishability where the adversary is allowed to corrupt the key generation authorities and also make key queries for message vector distinguishing key vectors (as long as the attribute does not satisfy the encrypted predicate). Below we provide the selective security variant of the corresponding property.⁹

⁹In this work, we only focus on standard semantic security, but one could also amplify to its CCA counterpart by relying on the generic CPA-to-CCA amplification techniques [KW19].

Definition 5.1 (Selective MA-AB-IPFE security with static corruptions). A MA-AB-IPFE scheme is selectively secure with static corruptions if for every stateful admissible PPT adversary \mathcal{A} , there exists a negligible function $\text{negl}(\cdot)$ such that for all $\lambda \in \mathbb{N}$, the following holds

$$\Pr \left[\begin{array}{l} \mathcal{A}^{O(\text{key}, \cdot, \cdot, \cdot)}(\{\text{PK}_i\}_{i \in [n] \setminus S^*}, \text{CT}) = b : \\ \text{PP} \leftarrow \text{GSetup}(1^\lambda) \\ (1^n, 1^\ell, S^*, C, (\mathbf{v}_0, \mathbf{v}_1), \{\text{PK}_i\}_{i \in S^*}) \leftarrow \mathcal{A}(1^\lambda, \text{PP}) \\ \forall i \in [n] \setminus S^* : (\text{PK}_i, \text{MSK}_i) \leftarrow \text{LSetup}(\text{PP}, 1^n, 1^\ell, i) \\ b \leftarrow \{0, 1\}, \text{CT} \leftarrow \text{Enc}(\{\text{PK}_i\}_{i \in [n]}, C, \mathbf{v}_b) \\ \text{key} = \{(\text{PK}_i, \text{MSK}_i)\}_{i \in [n] \setminus S^*} \end{array} \right] \leq \frac{1}{2} + \text{negl}(\lambda),$$

where the oracle $O(\text{key}, \cdot, \cdot, \cdot)$ has the master key for honest authorities hardwired. The oracle on input a tuple of a global identifier GID , an authority index $j \in [n] \setminus S^*$, and an attribute-key vector pair (b, \mathbf{u}) , responds with a partial secret key computed as $\text{SK}_{j, \text{GID}, b, \mathbf{u}} \leftarrow \text{KeyGen}(\text{MSK}_j, \text{GID}, b, \mathbf{u})$. Note that the adversary is only allowed to submit key queries for non-corrupt authorities (i.e., $j \notin S^*$). Also, the adversary \mathcal{A} is admissible as long as every secret key query made by \mathcal{A} to the key generation oracle O satisfies the condition that — (1) either $\langle \mathbf{u}, \mathbf{v}_0 \rangle = \langle \mathbf{u}, \mathbf{v}_1 \rangle$, or (2) C does not accept any input \mathbf{x} such that $x_j = b$ for $(b, j) \in Q_{\text{GID}}$ where Q_{GID} contains the attribute bits queries for GID ¹⁰.

5.1 Construction

Let Gen be a composite-order bilinear group generator. Also, let \mathbb{G} and \mathbb{G}_T be the source and target groups, respectively. Additionally, we rely on a hash function $H : \{0, 1\}^* \rightarrow \mathbb{G}$ that maps global identities GID to elements of \mathbb{G} and we later model it as a random oracle in the proof. Below we provide our MA-AB-IPFE scheme based on composite-order bilinear maps for the predicates described as an access policy for a linear secret sharing scheme.

$\text{GSetup}(1^\lambda) \rightarrow \text{PP}$. The setup algorithm samples a bilinear group as follows

$$(p_1, p_2, p_3, \mathbb{G}, \mathbb{G}_T, e(\cdot, \cdot)) \leftarrow \text{Gen}(1^\lambda, 3).$$

It samples a random generator $g_1 \in \mathbb{G}_1$, and sets the global public parameters as $\text{PP} = (g_1, N = p_1 p_2 p_3, \mathbb{G}, \mathbb{G}_T, e(\cdot, \cdot))$.

(Notation. Here and throughout, we use the ‘bracket’ notation for representing group elements. Where $[1]_1 := g_1$, and $[1]_{T,1} := e(g_1, g_1)$.)

$\text{LSetup}(\text{PP}, 1^n, 1^\ell, i) \rightarrow (\text{PK}, \text{MSK})$. The algorithm samples two random vectors $\alpha, \mathbf{w} \leftarrow \mathbb{Z}_N^\ell$, and sets the authority public-secret key pair as $\text{PK} = (\text{PP}, [\alpha]_{T,1}, [\mathbf{w}]_1)$ and $\text{MSK} = (\alpha, \mathbf{w})$. (Here and throughout, note that $[\mathbf{w}]_1$ and similar terms can be computed as $g_1^{\mathbf{w}}$.)

$\text{KeyGen}(\text{MSK}_j, \text{GID}, b, \mathbf{u}) \rightarrow \text{SK}_{j, \text{GID}, b, \mathbf{u}}$. It parses the authority key as described above. If $b = 0$, it sets the secret key as empty string. Otherwise, it first hashes the GID to create a masking term $[\mu] \in \mathbb{G}$ as $[\mu] = H(\text{GID})$. It then outputs the secret key as

$$\text{SK}_{j, \text{GID}, b, \mathbf{u}} = [\langle \alpha, \mathbf{u} \rangle]_1 \cdot [\mu \cdot \langle \mathbf{w}, \mathbf{u} \rangle].$$

Note that since the vectors $\mathbf{u}, \mathbf{w}, \alpha$ are known to the algorithm in the clear, thus the above key term can be computed efficiently.

¹⁰Note that in general this could be a non-falsifiable condition to check if S^* is $\omega(\log \lambda)$ and the predicate class contains general non-monotonic functions.

$\text{Enc}(\{\text{PK}_i\}_{i \in [n]}, (\mathbf{A}, \rho), \mathbf{v}) \rightarrow \text{CT}$. The encryption algorithm first parses the keys PK_i as $(\text{PP}, [\boldsymbol{\alpha}_i]_{T,1}, [\mathbf{w}_i]_1)$, and the predicate contains an $m_1 \times m_2$ access matrix \mathbf{A} with function ρ mapping the rows to the attribute positions. It samples a $m_2 \times \ell$ matrix \mathbf{S} and $(m_2 - 1) \times \ell$ matrix \mathbf{T}' uniformly at random as $\mathbf{S} \leftarrow \mathbb{Z}_N^{m_2 \times \ell}$ and $\mathbf{T}' \leftarrow \mathbb{Z}_N^{(m_2-1) \times \ell}$. It sets a $m_2 \times \ell$ matrix \mathbf{T} , and arranges two $m_1 \times \ell$ matrices $\boldsymbol{\Delta}$ and $\boldsymbol{\Gamma}$ as

$$\mathbf{T} = \begin{pmatrix} \mathbf{0}^\top \\ \mathbf{T}' \end{pmatrix}, \quad \boldsymbol{\Delta} = \begin{pmatrix} \boldsymbol{\alpha}_{\rho(1)}^\top \\ \vdots \\ \boldsymbol{\alpha}_{\rho(m_1)}^\top \end{pmatrix}, \quad \boldsymbol{\Gamma} = \begin{pmatrix} \mathbf{w}_{\rho(1)}^\top \\ \vdots \\ \mathbf{w}_{\rho(m_1)}^\top \end{pmatrix}.$$

That is, the matrix \mathbf{T} contains all zeros in the first row and is random otherwise. It also samples a random vector as $\mathbf{r} \leftarrow \mathbb{Z}_N^{m_1}$, and computes the ciphertext $\text{CT} = (C_0, C_1, C_2, C_3)$ as:

$$\begin{aligned} C_0 &= [\mathbf{s}_1 + \mathbf{v}]_{T,1}, & C_1 &= [\mathbf{A} \cdot \mathbf{S} + \boldsymbol{\Delta} \odot (\mathbf{r} \otimes \mathbf{1}^\top)]_{T,1}, \\ C_2 &= [\mathbf{r}]_1, & C_3 &= [\mathbf{A} \cdot \mathbf{T} + \boldsymbol{\Gamma} \odot (\mathbf{r} \otimes \mathbf{1}^\top)]_1. \end{aligned}$$

Here the vector \mathbf{s}_1 is the first column vector of matrix \mathbf{S}^\top (that is, $\mathbf{s}_1 = \mathbf{S}^\top \cdot \mathbf{e}_1$ where \mathbf{e}_1 is the first fundamental basis vector of $\mathbb{Z}_N^{m_2}$).

$\text{Dec}(\{\text{SK}_{i, \text{GID}, x_i, \mathbf{u}}\}_{i \in [n]}, \text{CT}) \rightarrow M$. It parses the secret key and ciphertext as described above. Let (\mathbf{A}, ρ) be the access policy associated with the ciphertext, and \mathbf{u} be the key vector associated with the partial secret keys. (This could either be explicitly added to the ciphertext and secret keys above, or passed as an auxiliary input.)

The decryptor first computes the LSSS reconstruction vector \mathbf{z} such that $\mathbf{z}^\top \cdot \mathbf{A} = \mathbf{e}_1^\top = (1, 0, \dots, 0)$. The decryptor then arranges the key terms as

$$K = \begin{pmatrix} \text{SK}_{\rho(1), \text{GID}, x_{\rho(1)}, \mathbf{u}} \\ \vdots \\ \text{SK}_{\rho(m_1), \text{GID}, x_{\rho(m_1)}, \mathbf{u}} \end{pmatrix}$$

and recovers the inner product message value M by computing the discrete log of the following the following:

$$[M]_{T,1} = \frac{\langle C_0, \mathbf{u} \rangle}{(\mathbf{z}^\top \cdot C_1 \cdot \mathbf{u})} \cdot \frac{\mathbf{z}^\top \cdot e(K, C_2)}{e(H(\text{GID}), \mathbf{z}^\top \cdot C_3 \cdot \mathbf{u})}$$

where the matrix vector operations involving group elements and exponents are performed by first raising the exponent of each term (component-by-component) for performing multiplication in the exponent, and then followed by multiplication of the resulting encodings to simulate addition being performed in the exponent. Also, the operation $e(K, C_2)$ performs the pairing operation element-by-element for each element of the vector.

5.2 Correctness and Security

In this section, we provide the proofs of correctness and security. Our security proof consists of a description of a sequence of hybrid games, and prove indistinguishability of successive games.

Correctness. The proof of correctness follows from the correctness of the LSSS scheme, and underlying algebraic structure. Below we briefly highlight the main points.

First, note that for an access policy (\mathbf{A}, ρ) , if an attribute vector $\mathbf{x} \in \{0, 1\}^n$ satisfies it, then there exists an LSSS reconstruction vector \mathbf{z} as defined in decryption that satisfies the property that $\mathbf{z}^\top \cdot \mathbf{A} = \mathbf{e}_1^\top = (1, 0, \dots, 0)$. Now consider the partial secret keys associated with attribute vector \mathbf{x} and key vector \mathbf{u} for identifier GID , $\text{SK}_{i, \text{GID}, x_i, \mathbf{u}}$ for $i \in [n]$. Observe that $\text{SK}_{i, \text{GID}, x_i, \mathbf{u}}$ is empty whenever $x_i = 0$, otherwise it contains the group element $[\langle \boldsymbol{\alpha}_i, \mathbf{u} \rangle]_1 \cdot [\mu \cdot \langle \mathbf{w}_i, \mathbf{u} \rangle]$ where $[\mu] = H(\text{GID})$.

Now let us look at the term $[M]_{T,1}$, the decryptor computes. It contains the following terms which can be simplified as follows–

$$\begin{aligned}
\langle C_0, \mathbf{u} \rangle &= \langle [\mathbf{s}_1 + \mathbf{v}]_{T,1}, \mathbf{u} \rangle = [\langle \mathbf{s}_1 + \mathbf{v}, \mathbf{u} \rangle]_{T,1}, \\
\mathbf{z}^\top \cdot C_1 \cdot \mathbf{u} &= \mathbf{z}^\top \cdot [\mathbf{A} \cdot \mathbf{S} + \mathbf{\Delta} \odot (\mathbf{r} \otimes \mathbf{1}^\top)]_{T,1} \cdot \mathbf{u} \\
&= [\mathbf{z}^\top \cdot \mathbf{A} \cdot \mathbf{S} \cdot \mathbf{u} + \mathbf{z}^\top \cdot (\mathbf{\Delta} \odot (\mathbf{r} \otimes \mathbf{1}^\top)) \cdot \mathbf{u}]_{T,1} \\
&= [\mathbf{e}_1^\top \cdot \mathbf{S} \cdot \mathbf{u} + \mathbf{z}^\top \cdot (\mathbf{\Delta} \odot (\mathbf{r} \otimes \mathbf{1}^\top)) \cdot \mathbf{u}]_{T,1} \\
&= [\mathbf{e}_1^\top \cdot \mathbf{S} \cdot \mathbf{u} + \mathbf{z}^\top \cdot ((\mathbf{\Delta} \cdot \mathbf{u}) \odot \mathbf{r})]_{T,1} \\
\\
\mathbf{z}^\top \cdot e(K, C_2) &= \mathbf{z}^\top \cdot \begin{pmatrix} e(\text{SK}_{\rho(1), \text{GID}, x_{\rho(1)}, \mathbf{u}}, [r_1]_1) \\ \vdots \\ e(\text{SK}_{\rho(m_1), \text{GID}, x_{\rho(m_1)}, \mathbf{u}}, [r_{m_1}]_1) \end{pmatrix} \\
&= \mathbf{z}^\top \cdot \begin{pmatrix} e([\langle \boldsymbol{\alpha}_{\rho(1)}, \mathbf{u} \rangle]_1 \cdot [\mu \cdot \langle \mathbf{w}_{\rho(1)}, \mathbf{u} \rangle], [r_1]_1) \\ \vdots \\ e([\langle \boldsymbol{\alpha}_{\rho(m_1)}, \mathbf{u} \rangle]_1 \cdot [\mu \cdot \langle \mathbf{w}_{\rho(m_1)}, \mathbf{u} \rangle], [r_{m_1}]_1) \end{pmatrix} \\
&= \mathbf{z}^\top \cdot [(\mathbf{\Delta} \cdot \mathbf{u} + \mu \cdot \boldsymbol{\Gamma} \cdot \mathbf{u}) \odot \mathbf{r}]_{T,1} \\
&= [\mathbf{z}^\top \cdot ((\mathbf{\Delta} \cdot \mathbf{u}) \odot \mathbf{r}) + \mathbf{z}^\top \cdot ((\mu \cdot \boldsymbol{\Gamma} \cdot \mathbf{u}) \odot \mathbf{r})]_{T,1} \\
\\
\mathbf{z}^\top \cdot C_3 \cdot \mathbf{u} &= \mathbf{z}^\top \cdot [\mathbf{A} \cdot \mathbf{T} + \boldsymbol{\Gamma} \odot (\mathbf{r} \otimes \mathbf{1}^\top)]_1 \cdot \mathbf{u} \\
&= [\mathbf{e}_1^\top \cdot \mathbf{T} \cdot \mathbf{u} + \mathbf{z}^\top \cdot ((\boldsymbol{\Gamma} \cdot \mathbf{u}) \odot \mathbf{r})]_1 \\
e(H(\text{GID}), \mathbf{z}^\top \cdot C_3 \cdot \mathbf{u}) &= [\mu \cdot \mathbf{e}_1^\top \cdot \mathbf{T} \cdot \mathbf{u} + \mu \cdot \mathbf{z}^\top \cdot ((\boldsymbol{\Gamma} \cdot \mathbf{u}) \odot \mathbf{r})]_{T,1}
\end{aligned}$$

Combining all these together, we get that

$$\begin{aligned}
[M]_{T,1} &= \frac{[\langle \mathbf{s}_1 + \mathbf{v}, \mathbf{u} \rangle]_{T,1}}{[\mathbf{e}_1^\top \cdot \mathbf{S} \cdot \mathbf{u} + \mathbf{z}^\top \cdot ((\mathbf{\Delta} \cdot \mathbf{u}) \odot \mathbf{r})]_{T,1}} \cdot \frac{[\mathbf{z}^\top \cdot ((\mathbf{\Delta} \cdot \mathbf{u}) \odot \mathbf{r}) + \mathbf{z}^\top \cdot ((\mu \cdot \boldsymbol{\Gamma} \cdot \mathbf{u}) \odot \mathbf{r})]_{T,1}}{[\mu \cdot \mathbf{e}_1^\top \cdot \mathbf{T} \cdot \mathbf{u} + \mu \cdot \mathbf{z}^\top \cdot ((\boldsymbol{\Gamma} \cdot \mathbf{u}) \odot \mathbf{r})]_{T,1}} \\
&= [\langle \mathbf{s}_1 + \mathbf{v}, \mathbf{u} \rangle - \mathbf{e}_1^\top \cdot \mathbf{S} \cdot \mathbf{u} - \mu \cdot \mathbf{e}_1^\top \cdot \mathbf{T} \cdot \mathbf{u}]_{T,1} \\
&= [\langle \mathbf{v}, \mathbf{u} \rangle]_{T,1}
\end{aligned}$$

where the last equality follows from the facts that $\mathbf{e}_1^\top \cdot \mathbf{S} = \mathbf{s}_1^\top$, and $\mathbf{e}_1^\top \cdot \mathbf{T} = \mathbf{0}^\top$. Therefore, whenever the policy (\mathbf{A}, ρ) is satisfied, then $[M]_{T,1}$ gets simplified to $[\mathbf{u}^\top \mathbf{v}]_{T,1}$ which can be recovered by computing the discrete log. Thus, correctness follows.

Security. Below we provide a proof of security under standard assumptions over composite-order bilinear groups.

Theorem 5.2. If the assumptions 2.4, 2.5, 2.6, and 2.7 hold over the bilinear group generator Gen , then the scheme described above is a selectively secure MA-AB-IPFE for LSSS access structures as per Definition 5.1.

As discussed in Section 1.3, in addition of the type of key queries allowed under MA-ABE, we need to answer key queries for attribute-key vector pairs $(j, \text{GID}, b, \mathbf{u})$ such that $\langle \mathbf{u}, \mathbf{v}_0 - \mathbf{v}_1 \rangle = 0$. At a high level, our idea is to sample the authority secret keys $\boldsymbol{\alpha}_i, \mathbf{w}_i$ such that the reduction could honestly compute the secret keys for all attribute-key vector pairs $(j, \text{GID}, b, \mathbf{u})$ such that \mathbf{A} accepts the list of attribute bits queried for a

given identifier GID and $\langle \mathbf{u}, \mathbf{v}_0 \rangle = \langle \mathbf{u}, \mathbf{v}_1 \rangle$. On the other hand, to answer key queries for attribute-key vector pairs $(j, \text{GID}, b, \mathbf{u})$ such that \mathbf{A} does *not* accept the list of attribute bits queried for a given identifier GID , we switch such secret keys to their semi-functional counterparts.

Although this intuition seems to work at a high level, it does not work as described above. To that end, we switch all the accepting secret keys to their semi-functional counterparts as well, but ensure that the challenge ciphertext components that they interact with are only nominally semi-functional. Here the last part is necessary because of two reasons: first, making the entire challenge ciphertext semi-functional will affect decryption w.r.t. accepting keys which will be distinguishable for the adversary; second, it is unclear how to sample the challenge ciphertext in which only one component is semi-functional while other are regular sub-encryptions due to the fact that these different ciphertext components are significantly correlated. We get around this barrier by ensuring that the challenge ciphertext is sampled as what we call a “partial semi-functional ciphertext” (which has full-nominally semi-functional components along with standard semi-functional components).

Proof. We start by sketching the sequence of games where the first game corresponds to the original MA-AB-IPFE security game. First, we switch the way we sample the honest authority’s public-secret keys wherein instead of sampling the random vectors α_i and \mathbf{w}_i for authority i directly, we sample two random vectors $\tilde{\alpha}_i$ and $\tilde{\mathbf{w}}_i$ of same dimensions and implicitly set $\alpha_i^\top = \tilde{\alpha}_i^\top \mathbf{F}$ and $\mathbf{w}_i^\top = \tilde{\mathbf{w}}_i^\top \mathbf{F}$ where $\mathbf{F} \in \mathbb{Z}_N^{\ell \times \ell}$ is a random full rank matrix such that $\mathbf{F}(\mathbf{v}_0 - \mathbf{v}_1) = \mathbf{e}_1$ where \mathbf{e}_1 is the first vector in the canonical basis of \mathbb{Z}_N^ℓ .

Next, we switch the way the random oracle queries are answered wherein we map identities GID to random elements of \mathbb{G}_{p_1} instead of \mathbb{G} . This is followed by switching the challenge ciphertext to come from a special ciphertext distribution which we refer to as a *partial* semi-functional ciphertext, where the intuition is that only the projection of the ℓ -dimension message vector onto a one-dimensional subspace is encrypted under a semi-functional mode while the remaining $(\ell - 1)$ -dimensional projection are encrypted under the full-nominally semi-functional encryption mode. Basically, we split the challenge message vector such that now we encrypt the vector $\beta(\mathbf{v}_0 - \mathbf{v}_1)$ as a semi-functional ciphertext (for random challenge bit β) and vector \mathbf{v}_1 as a full-nominally semi-functional ciphertext and homomorphically combines both of them to create the final challenge ciphertext. Finally, we switch the partial secret keys one-by-one to be semi-functional irrespective of whether \mathbf{A} accept the list of attribute bits queried for a given identifier GID or not. An important property of the full-nominally semi-functional ciphertexts is that they can be decrypted by all types of semi-functional keys, while having the capability to be jointly sampled with a semi-functional ciphertext with correlated randomness.

The idea here is that since the adversary never receives any secret key for attribute-key vector pairs $(j, \text{GID}, b, \mathbf{u})$ where both \mathbf{A} accepts the list of attribute bits queried for a given identifier GID and $\langle \mathbf{u}, \mathbf{v}_0 - \mathbf{v}_1 \rangle \neq 0$, thus the reduction algorithm can perfectly simulate the above games using the dual system encryption paradigm [Wat09] as used in [LW11]. We want to point out that as in [LW11], we will assume a one-use restriction on the attributes throughout the proof, that is the row labeling function ρ of the challenge ciphertext is injective.

Description of games

Game 0. This corresponds to the original MA-AB-IPFE security game. Let Q denote the total number of key queries made by adversary.

Game 1. This is same as the previous game, except for each non-corrupt authority, i.e. $i \notin S^*$, the challenger:

- samples a uniformly random vectors $\tilde{\alpha}_i, \tilde{\mathbf{w}}_i \leftarrow \mathbb{Z}_N^\ell$,
- samples a random *orthogonal* matrix $\mathbf{F} \in \mathbb{Z}_N^{\ell \times \ell}$ subject to the constraint that $\mathbf{F}(\mathbf{v}_0 - \mathbf{v}_1) = \mathbf{e}_1$, where $\mathbf{v}_0, \mathbf{v}_1$ are the challenge message vectors and $\mathbf{e}_1 = (1, 0, \dots, 0)^\top$ (i.e., the first canonical basis vector of \mathbb{Z}_N^ℓ), and

— sets $\alpha_i = \mathbf{F}^\top \tilde{\alpha}_i, \mathbf{w}_i = \mathbf{F}^\top \tilde{\mathbf{w}}_i$ instead of sampling it uniformly at random.

Re-writing the keys and challenge ciphertext. With the above change in sampling the parameters, we below re-write how all the keys and ciphertexts in the scheme can be rewritten.

$$\begin{aligned} \text{PK}_i &= (\text{PP}, \mathbf{F}^\top \cdot [\tilde{\alpha}_i]_{T,1}, \mathbf{F}^\top \cdot [\tilde{\mathbf{w}}_i]_1), \quad \text{SK}_{i,\text{GID},1,\mathbf{u}} = [\langle \tilde{\alpha}_i, \mathbf{F} \cdot \mathbf{u} \rangle]_1 \cdot [\mu \cdot \langle \tilde{\mathbf{w}}_i, \mathbf{F} \cdot \mathbf{u} \rangle]. \\ \tilde{\mathbf{S}} &\leftarrow \mathbb{Z}_N^{m_2 \times \ell}, \quad \tilde{\mathbf{T}} = \begin{pmatrix} & \mathbf{0}^\top \\ \tilde{\mathbf{T}}' & \leftarrow \mathbb{Z}_N^{(m_2-1) \times \ell} \end{pmatrix}, \quad \tilde{\mathbf{\Delta}} = \begin{pmatrix} \tilde{\alpha}_{\rho(1)}^\top \\ \vdots \\ \tilde{\alpha}_{\rho(m_1)}^\top \end{pmatrix}, \quad \tilde{\mathbf{\Gamma}} = \begin{pmatrix} \tilde{\mathbf{w}}_{\rho(1)}^\top \\ \vdots \\ \tilde{\mathbf{w}}_{\rho(m_1)}^\top \end{pmatrix}. \\ \mathbf{S} &= \tilde{\mathbf{S}} \cdot \mathbf{F}, \quad \mathbf{T} = \tilde{\mathbf{T}} \cdot \mathbf{F}, \quad \mathbf{\Delta} = \tilde{\mathbf{\Delta}} \cdot \mathbf{F}, \quad \mathbf{\Gamma} = \tilde{\mathbf{\Gamma}} \cdot \mathbf{F}. \\ C_0 &= [\mathbf{F}^\top \tilde{\mathbf{S}}^\top \mathbf{e}_1 + \beta(\mathbf{v}_1 - \mathbf{v}_0) + \mathbf{v}_0]_{T,1}, \quad C_1 = [\mathbf{A} \cdot \tilde{\mathbf{S}} \cdot \mathbf{F} + \tilde{\mathbf{\Delta}} \cdot \mathbf{F} \odot (\mathbf{r} \otimes \mathbf{1}^\top)]_{T,1}, \\ &= \mathbf{F}^\top \cdot [\tilde{\mathbf{S}}^\top \mathbf{e}_1 - \beta \mathbf{e}_1 + \mathbf{F} \cdot \mathbf{v}_0]_{T,1}, \quad = [\mathbf{A} \cdot \tilde{\mathbf{S}} + \tilde{\mathbf{\Delta}} \odot (\mathbf{r} \otimes \mathbf{1}^\top)]_{T,1} \cdot \mathbf{F}, \\ C_2 &= [\mathbf{r}]_1, \quad C_3 = [\mathbf{A} \cdot \tilde{\mathbf{T}} \cdot \mathbf{F} + \tilde{\mathbf{\Gamma}} \cdot \mathbf{F} \odot (\mathbf{r} \otimes \mathbf{1}^\top)]_1, \\ & \quad = [\mathbf{A} \cdot \tilde{\mathbf{T}} + \tilde{\mathbf{\Gamma}} \odot (\mathbf{r} \otimes \mathbf{1}^\top)]_1 \cdot \mathbf{F}, \end{aligned}$$

The above re-writing of the keys and ciphertext will be crucial later in the proof when dividing the type of queries and moving the ciphertext to *partial* semi-functional.

Game 2. This is same as the previous game, except the challenger answers the random oracle queries as follows:

— for fresh random oracle query, it samples a random vector $\mu_{\text{GID}} \leftarrow \mathbb{Z}_N$ and responding with $[\mu_{\text{GID}}]_1 = g_1^{\mu_{\text{GID}}}$. Therefore, the partial secret keys can be simplified as:

$$\text{SK}_{i,\text{GID},1,\mathbf{u}} = [\langle \tilde{\alpha}_i + \mu_{\text{GID}} \cdot \tilde{\mathbf{w}}_i, \mathbf{F} \cdot \mathbf{u} \rangle]_1.$$

Defining partial semi-functional ciphertexts and secret keys. Similar to [LW11], we define the space of semi-functional ciphertexts and keys for our scheme. The main differences are that rather than turning the entire ciphertext to a semi-functional ciphertext, we only switch one dimensional subspace out of the ℓ dimensions of the ciphertext to be semi-functional and remaining $\ell - 1$ dimensional projection of the ciphertext space are turned full-nominally semi-functional, and these components are then algebraically manipulated given the vector space transformation matrix \mathbf{F} by performing group operations on the encodings. Looking ahead, a full-nominally semi-functional ciphertext is more general than a nominally semi-functional ciphertext in that the latter can not be decrypted by one type of semi-functional keys whereas a full-nominally semi-functional ciphertext can be decrypted by all types of semi-functional keys. The difference between a regular and a full-nominally semi-functional ciphertext is only in that a regular ciphertext has no blinding factors, whereas a full-nominally semi-functional ciphertext has blinding factors in all remaining subgroups but they are structured such that decryption is unaffected. At a high level, the reason for defining a full-nominally semi-functional ciphertext is that algebraical manipulation of a semi-functional ciphertext and regular ciphertext is not possible due to lack of blinding factors in a regular ciphertext. Therefore, introducing structured blinding factors in a regular ciphertext make it decryptable with respect to all semi-functional keys but they also contain blinding factors thereby making algebraic manipulation possible.

For the remainder of the proof and descriptions of all types of semi-functional ciphertexts and keys, we use the following notation. Let $\text{bad} \subset [m_1]$ denote the set of rows for which the corresponding attributes

Figure 5.1: Sampling partial semi-functional ciphertexts.

The challenger samples the following ciphertext matrices

$$\tilde{C}_2 = [\mathbf{r}]_1, \quad \tilde{C}_3 = [\mathbf{A} \cdot \tilde{\mathbf{T}} + \tilde{\mathbf{\Gamma}} \odot (\mathbf{r} \otimes \mathbf{1}^\top)]_1,$$

where \mathbf{r} and $\tilde{\mathbf{T}}$ are sampled as described previously. It then samples two random matrices $\mathbf{T}^*, \hat{\mathbf{T}}^* \leftarrow \mathbb{Z}_N^{m_2 \times \ell}$, and computes the semi-functional term of C_3 as

$$\tilde{C}_3^* = [\mathbf{A} \cdot \mathbf{T}^*]_2 \odot [\mathbf{A} \cdot \hat{\mathbf{T}}^*]_3 \odot [\tilde{\mathbf{\Gamma}} \odot ((\mathbf{r} \odot \mathbb{I}^*) \otimes \mathbf{1}^\top)]_{\{2,3\}}$$

where encodings in subgroups 2 and 3 are computed using a fixed generator of these subgroups sampled at setup time. Finally, it sets the C_2 and C_3 components of the challenge ciphertext as

$$C_2 = \tilde{C}_2 \odot [\mathbf{r} \odot \mathbb{I}^*]_{\{2,3\}}, \quad C_3 = (\tilde{C}_3 \odot \tilde{C}_3^*) \cdot \mathbf{F}. \quad (5.1)$$

We say the ciphertext is full-nominally semi-functional for column $j \in [\ell]$ if the first row of matrices \mathbf{T}^* and $\hat{\mathbf{T}}^*$ has 0 at the j -th index (that is, $T_{1,j}^* = \hat{T}_{1,j}^* = 0$). In other words, this means that the secret being shared in column j is 0.

And, we say that the ciphertext is full-nominally semi-functional if it is full-nominally semi-functional for all but the first column. That is, the first rows of matrices \mathbf{T}^* and $\hat{\mathbf{T}}^*$ is the first basis vector in the canonical basis of \mathbb{Z}_N^ℓ .

belong to corrupted authorities. That is, $\text{bad} = \{i \in [m_1] : \rho(i) \in S^*\}$. Also, let $\mathbb{I}^* \in \{0, 1\}^{m_1}$ be the vector such that i -th component $\mathbb{I}_i^* = 0$ if $i \in \text{bad}$, otherwise $\mathbb{I}_i^* = 1$.

Partial semi-functional ciphertexts. The C_0 and C_1 components are sampled exactly as before. Here we switch how the C_2 and C_3 components are sampled. The procedure is described in detail in Fig. 5.1.

Semi-functional keys. When we refer to the key for identity GID , we regard as the set of all keys $\text{SK}_{i, \text{GID}, 1, \mathbf{u}}$ for the attributes i belonging to non-corrupt authorities requested by the attacker throughout the game. Recall that the adversary is also allowed to make key queries for accepting attributes for a fixed GID given that $\langle \mathbf{u}, \mathbf{v}_0 - \mathbf{v}_1 \rangle = 0$. Now we sample all the secret keys as semi-functional keys irrespective of whether $\langle \mathbf{u}, \mathbf{v}_0 - \mathbf{v}_1 \rangle = 0$ or not. One might think that this will be problematic since a semi-functional key for an accepting attribute will no longer decrypt the semi-functional ciphertext from above. However, this is precisely avoided by making the ciphertext to be full-nominally semi-functional, and ensuring that the accepting secret keys only interact with the nominally semi-functional component of the ciphertext. Recall that regular keys are sampled as

$$H(\text{GID}) = [\mu_{\text{GID}}]_1, \quad \text{SK}_{i, \text{GID}, 1, \mathbf{u}} = [\langle \tilde{\alpha}_i, \mathbf{F} \cdot \mathbf{u} \rangle]_1 \cdot [\langle \mu_{\text{GID}} \cdot \tilde{\mathbf{w}}_i, \mathbf{F} \cdot \mathbf{u} \rangle]_1. \quad (5.2)$$

The important aspect of these keys are that they *do not* depend on the first component of vectors $\tilde{\alpha}_i$ and $\tilde{\mathbf{w}}_i$, that is $\tilde{\alpha}_{i,1}$ and $\tilde{w}_{i,1}$ are not needed for generating $\text{SK}_{i, \text{GID}, 1, \mathbf{u}}$ whenever the key corresponds to accepting attributes. This is because whenever $\langle \mathbf{u}, \mathbf{v}_0 - \mathbf{v}_1 \rangle = 0$ (that is, the condition which must be true for accepting key queries), we get that $\mathbf{F} \cdot \mathbf{u}$ is orthogonal to the first basis vector in the canonical basis of \mathbb{Z}_N^ℓ . That is, $\mathbf{F} \cdot \mathbf{u}$ is of the form $(0, *, \dots, *)^\top$.

Now we sample the secret keys one of two ways. We refer to these as type 1 and 2 semi-functional keys as in [LW11]. A type-1 secret key contains blinding factors in the second subgroup, whereas type-2 secret key contains blinding factors in third subgroup such that they do not a nominally semi-functional ciphertext component does not get affected by type-1 keys.

A type-1 semi-functional secret key is computed as follows

$$H(\text{GID}) = [\mu_{\text{GID}}]_{\{1,2\}}, \quad \text{SK}_{i,\text{GID},1,\mathbf{u}} = [\langle \tilde{\boldsymbol{\alpha}}_i, \mathbf{F} \cdot \mathbf{u} \rangle]_1 \cdot [\langle \mu_{\text{GID}} \cdot \tilde{\mathbf{w}}_i, \mathbf{F} \cdot \mathbf{u} \rangle]_{\{1,2\}}. \quad (5.3)$$

Similarly, a type-2 semi-functional secret key is computed as follows

$$H(\text{GID}) = [\mu_{\text{GID}}]_{\{1,3\}}, \quad \text{SK}_{i,\text{GID},1,\mathbf{u}} = [\langle \tilde{\boldsymbol{\alpha}}_i, \mathbf{F} \cdot \mathbf{u} \rangle]_1 \cdot [\langle \mu_{\text{GID}} \cdot \tilde{\mathbf{w}}_i, \mathbf{F} \cdot \mathbf{u} \rangle]_{\{1,3\}}. \quad (5.4)$$

Game 3. This is the same as the previous game, except that the challenge ciphertext is partial semi-functional as described in Fig. 5.1. That is, the challenger samples the first two components of the challenge ciphertext C_0 and C_1 as before that is

$$C_0 = \mathbf{F}^\top \cdot [\tilde{\mathbf{S}}^\top \mathbf{e}_1 - \beta \mathbf{e}_1 + \mathbf{F} \cdot \mathbf{v}_0]_{T,1}, \quad C_1 = [\mathbf{A} \cdot \tilde{\mathbf{S}} + \tilde{\boldsymbol{\Delta}} \odot (\mathbf{r} \otimes \mathbf{1}^\top)]_{T,1} \cdot \mathbf{F}$$

where β is the random challenge bit. Whereas the last two components C_2 and C_3 are sampled as in Eq. (5.1).

Note. The above Game 3 is same as Game 4.0.2 that we describe below.

Game 4.q.1. This is same as Game 3, except the challenger answers the key queries as follows:

- for the first $q - 1$ key queries, on attribute-key vector pairs $(j, \text{GID}, b, \mathbf{u})$, it computes the partial key as a semi-functional key of ‘type-2’ as in Eq. (5.4).
- for the q -th key query, on attribute-key vector pair $(j, \text{GID}, b, \mathbf{u})$, it computes the partial key as a semi-functional key of ‘type-1’ as in Eq. (5.3).
- remaining key queries are answered exactly as in Game 3. That is, it computes the partial key honestly as in Eq. (5.2).

Game 4.q.2. This is same as Game 4.q.1, except the challenger answers the key queries as follows:

- for the q -th key query, on attribute-key vector pair $(j, \text{GID}, b, \mathbf{u})$, it computes the partial key as a semi-functional key of ‘type-2’ as in Eq. (5.4). (That is, the q -th key is switched from type 1 to type 2 as well. All keys in this game are either a type-2 semi-functional key, or honestly sampled keys. Also, note that all the accepting partial keys do not depend on the first component of the key vectors $\tilde{\boldsymbol{\alpha}}_i$ and $\tilde{\mathbf{w}}_i$.)

Game 5. This is the same as Game 4.Q.2, except that the challenge ciphertext is a partial semi-functional ciphertext where the C_0 component encrypts a random vector along the direction $\mathbf{v}_1 - \mathbf{v}_0$. That is,

$$C_0 = \mathbf{F}^\top \cdot [\tilde{\mathbf{S}}^\top \mathbf{e}_1 - \kappa \mathbf{e}'_1 + \mathbf{F} \cdot \mathbf{v}_0]_{T,1}$$

where κ is a random exponent, and $\mathbf{e}_1, \mathbf{e}'_1$ are first basis vectors in the canonical basis of $\mathbb{Z}_N^{m_2}$ and \mathbb{Z}_N^ℓ , respectively.

Note that in the last hybrid game, the challenger’s responses are independent of the challenge bit β , thus to complete the proof we only need to show that each adjacent game is indistinguishable.

Indistinguishability of games

To complete the proof we need to show that adjacent games are indistinguishable. The high level strategy behind the reductions is similar to [LW11] which is to rely on appropriate type of subgroup hiding assumption coupled with the one-use restriction on the attributes so that information-theoretic security of the secret sharing scheme could be used to move from nominally semi-functional ciphertexts to semi-functional ciphertexts.

However, due to the fact that we can not make all subspaces of the challenge ciphertext semi-functional, thus we turn only the first subspace to be semi-functional while remaining subspaces are made nominally semi-functional. The precise reductions are similar to those provided in [LW11], thus we only provide a high level sketch below. A detailed description will be in the full version.

For any adversary \mathcal{A} and game X , we denote by $\text{Adv}_s^{\mathcal{A}}(\lambda)$, the probability that \mathcal{A} wins in game s .

Lemma 5.3. For any PPT adversary \mathcal{A} , we have that $\text{Adv}_0^{\mathcal{A}}(\lambda) - \text{Adv}_1^{\mathcal{A}}(\lambda) \leq \text{negl}(\lambda)$ for some negligible function $\text{negl}(\cdot)$.

Proof. This follows directly from the fact that, for any invertible matrix $\mathbf{F} \in \mathbb{Z}_N^{\ell \times \ell}$, the following distributions are identical.

$$\left\{ \{\boldsymbol{\alpha}_i, \mathbf{w}_i\}_{i \in [n]} : \forall i \in [n], \boldsymbol{\alpha}_i, \mathbf{w}_i \leftarrow \mathbb{Z}_N^\ell \right\} \equiv \left\{ \{\mathbf{F}^\top \cdot \tilde{\boldsymbol{\alpha}}_i, \mathbf{F}^\top \cdot \tilde{\mathbf{w}}_i\}_{i \in [n]} : \forall i \in [n], \tilde{\boldsymbol{\alpha}}_i, \tilde{\mathbf{w}}_i \leftarrow \mathbb{Z}_N^\ell \right\}.$$

Since N is a composite with large factors, thus with all but negligible probability \mathbf{F} will be invertible as otherwise we could factor N thereby breaking security of the composite-order bilinear maps. And, since \mathbf{F} is invertible with all but negligible probability, thus the lemma follows. \square

Lemma 5.4. If assumption 2.4 holds over the group generator Gen , then for any PPT adversary \mathcal{A} , we have that $\text{Adv}_1^{\mathcal{A}}(\lambda) - \text{Adv}_2^{\mathcal{A}}(\lambda) \leq \text{negl}(\lambda)$ for some negligible function $\text{negl}(\cdot)$.

Proof. The proof of this lemma is similar to that of [LW11, Lemma 7]. That is, the reduction algorithm simply uses the challenge group element T , which lies either in \mathbb{G} or \mathbb{G}_1 , to answer all the random oracle queries (made directly or indirectly via the partial key queries). If $T \in \mathbb{G}$, then the reduction simulates Game 1, otherwise it simulates Game 2. Thus, the lemma follows. \square

Lemma 5.5. If assumption 2.4 holds over the group generator Gen , then for any PPT adversary \mathcal{A} , we have that $\text{Adv}_2^{\mathcal{A}}(\lambda) - \text{Adv}_3^{\mathcal{A}}(\lambda) \leq \text{negl}(\lambda)$ for some negligible function $\text{negl}(\cdot)$.

Proof. The proof of this lemma is similar to that of [LW11, Lemma 8], except now the reduction algorithm needs to generate a much larger ciphertext (ℓ [LW11]-like ciphertexts but non-trivially correlated). Ideally, if the correlations between the ℓ sub-ciphertexts were mild, and could be solved using reusability of correlated components then the proof would be more straightforward. But since one of these ℓ sub-ciphertexts must not be decryptable by the queried secret keys, while all other ciphertexts must be, thus the correlations are significant.

However, despite the technical challenges, subgroup decision assumption is still sufficient for our purposes. At a high level, the reduction algorithm simply uses the challenge group element T , which lies either in \mathbb{G} or \mathbb{G}_1 , to sample the C_2 and C_3 components of the ciphertext corresponding to rows which are non-corrupt, while for corrupt rows, only C_3 depends on the challenge T . This portion of reduction can be carried out similar to [LW11, Lemma 8] by sampling the vectors and matrices of appropriate dimensions directly over \mathbb{Z}_N and implicitly setting certain exponents such as \mathbf{r} . The main difference appears in the step that unlike [LW11, Lemma 8], where the authors directly argued semi-functionality of the ciphertext due to the fact that all keys are rejecting, we can not rely on the same strategy as we allow the adversary to make accepting queries as well as long as $\langle \mathbf{u}, \mathbf{v}_0 - \mathbf{v}_1 \rangle = 0$. But we note that by sampling the subspace transformation matrix \mathbf{F} and applying homomorphically on top of the challenge ciphertext, we could pinpoint a single subspace (which is along the first canonical basis vector in our case) where we switch the ciphertext to be semi-functional from the adversary's perspective, while all other sub-ciphertext components (that is, in the subspace orthogonal to the first canonical basis vector) are nominally semi-functional from the adversary's perspective as well. In a little more detail, the information-theoretic security argument for arguing that the secret shares are well distributed is only applied for the first subspace, but not the rest. \square

Lemma 5.6. If assumption 2.5 holds over the group generator Gen , then for any PPT adversary \mathcal{A} and $q \in \{1, \dots, Q\}$, we have that $\text{Adv}_{4.(q-1).2}^{\mathcal{A}}(\lambda) - \text{Adv}_{4.q.1}^{\mathcal{A}}(\lambda) \leq \text{negl}(\lambda)$ for some negligible function $\text{negl}(\cdot)$.

Proof. The proof of this lemma is similar to that of [LW11, Lemma 9] but with changes similar to that for Lemma 5.5. Here again we rely on the fact that neither the reduction or the attacker can distinguish the type-1 semi-functional key from a regular key. This is because for all but the first subspace the ciphertext is still nominally semi-functional both from the reduction and attacker’s perspective; whereas the first subspace appears like a semi-functional ciphertext to \mathcal{A} but for the challenger it is still nominally semi-functional. Now as in the previous lemma, we only apply the information-theoretic secret sharing guarantee in the first sub-ciphertext component. \square

Lemma 5.7. If assumption 2.6 holds over the group generator Gen , then for any PPT adversary \mathcal{A} and $q \in \{1, \dots, Q\}$, we have that $\text{Adv}_{4,q,1}^{\mathcal{A}}(\lambda) - \text{Adv}_{4,q,2}^{\mathcal{A}}(\lambda) \leq \text{negl}(\lambda)$ for some negligible function $\text{negl}(\cdot)$.

Proof. The proof of this lemma is identical to that of [LW11, Lemma 10]. \square

Lemma 5.8. If assumption 2.7 holds over the group generator Gen , then for any PPT adversary \mathcal{A} , we have that $\text{Adv}_{4,Q,2}^{\mathcal{A}}(\lambda) - \text{Adv}_5^{\mathcal{A}}(\lambda) \leq \text{negl}(\lambda)$ for some negligible function $\text{negl}(\cdot)$.

Proof. The proof of this lemma is similar to that of [LW11, Lemma 11], except now the strategy is to switch from encryption of bit β to a random element in the first subspace of the challenge ciphertext, while keeping the remaining sub-ciphertext portions to be identically distributed. Since the reduction algorithm receives the generators for all subgroups from the challenger, thus it could simulate the distribution of nominally semi-functional ciphertext components honestly. And , it uses the remaining components of the bilinear challenge to simulate the first sub-ciphertext component by using the strategy similar to in [LW11, Lemma 11]. This completes the proof. \square

\square

6 Function-Hiding DDFE for Inner Products

In this section, we present our function-hiding decentralized dynamic inner product functional encryption (IP-DDFE) scheme. As described in the introduction, we first obtain a function-hiding multi-client inner product functional encryption (IP-MCFE) scheme, and then lift it to a function-hiding IP-DDFE scheme in a non-black box manner. We first define necessary notions to describe our IP-MCFE and IP-DDFE scheme.

Syntax of MCFE. Let \mathcal{F} be a function family such that, for all $f \in \mathcal{F}$, $f : \mathcal{M}_1 \times \dots \times \mathcal{M}_n \rightarrow \mathcal{Z}$. Let \mathcal{L} be a label space. An MCFE scheme for \mathcal{F} and \mathcal{L} consists of four algorithms.

Setup($1^\lambda, 1^n$): It takes a security parameter 1^λ and a number 1^n of slots, and outputs a public parameter PK, encryption keys $\{\text{EK}_i\}_{i \in [n]}$, a master secret key MSK. The other algorithms implicitly take PK.

KeyGen(MSK, f): It takes MSK and $f \in \mathcal{F}$, and outputs a secret key SK.

Enc($i, \text{EK}_i, x_i, \text{lab}$): It takes MSK, an index $i \in [n]$, $x_i \in \mathcal{M}_i$, and a label lab and outputs a ciphertext CT_i .

Dec($\text{CT}_1, \dots, \text{CT}_n, \text{SK}$): It takes $\text{CT}_1, \dots, \text{CT}_n$ and SK, and outputs a decryption value $d \in \mathcal{Z}$ or a symbol \perp .

Correctness. An MCFE scheme is correct if it satisfies the following condition. For all $\lambda, n \in \mathbb{N}$, $(x_1, \dots, x_n) \in \mathcal{M}_1 \times \dots \times \mathcal{M}_n$, $f \in \mathcal{F}$, $\text{lab} \in \mathcal{L}$, we have

$$\Pr \left[d = f(x_1, \dots, x_n) : \begin{array}{l} \text{PK}, \{\text{EK}_i\}, \text{MSK} \leftarrow \text{Setup}(1^\lambda) \\ \text{CT}_i \leftarrow \text{Enc}(i, \text{EK}_i, x_i, \text{lab}) \\ \text{SK} \leftarrow \text{KeyGen}(\text{MSK}, f) \\ d = \text{Dec}(\text{CT}_1, \dots, \text{CT}_n, \text{SK}) \end{array} \right] = 1.$$

Security. We basically adopt the security definition for MCFE in [ABKW19] and extend it to function-hiding security. We also introduce a selective variant because our final goal is IP-DDFE with selective security, and selectively secure IP-MCFE is sufficient for the security analysis of our IP-DDFE scheme.

Definition 6.1 (Function-hiding security of MCFE). An MCFE scheme is Leak_y -xx-yy-function-hiding ($\text{xx} \in \{\text{sel}, \text{sta}, \text{adt}\}$, $\text{yy} \in \{\text{any}, \text{pos}\}$) if for every stateful PPT adversary \mathcal{A} , there exists a negligible function $\text{negl}(\cdot)$ such that for all $\lambda \in \mathbb{N}$, the following holds

$$\Pr \left[\beta \leftarrow \mathcal{A}^{\text{QCor}(), \text{QEnc}(), \text{QKeyGen}()}(\text{PK}) : \begin{array}{l} \beta \leftarrow \{0, 1\} \\ \text{PK}, \{\text{EK}_i\}, \text{MSK} \leftarrow \text{Setup}(1^\lambda) \end{array} \right] \leq \frac{1}{2} + \text{negl}(\lambda)$$

where $\text{QCor}(i)$ outputs EK_i , $\text{QEnc}(i, x_i^0, x_i^1, \text{lab})$ outputs $\text{Enc}(i, \text{EK}_i, x_i^\beta, \text{lab})$, and $\text{QKeyGen}(f^0, f^1)$ outputs $\text{KeyGen}(\text{MSK}, f^\beta)$. Let $q_{c,i,\text{lab}}$ be the numbers of queries of the forms of $\text{QEnc}(i, *, *, \text{lab})$. Let \mathcal{HS} be the set of parties on which the adversary has not queried QCor at the end of the game, and $\mathcal{CS} = [n] \setminus \mathcal{HS}$. Then, the adversary's queries must satisfy the following conditions.

- For $i \in \mathcal{CS}$, the queries $\text{QEnc}(i, x_i^0, x_i^1, \text{lab})$ and $\text{QKeyGen}(f^0, f^1)$ must satisfy $x_i^0 = x_i^1$ and $\text{Leak}_y(i, f^0) = \text{Leak}_y(i, f^1)$, respectively.¹¹
- There are no sequences $(x_1^0, \dots, x_n^0, f^0, \text{lab})$ and $(x_1^1, \dots, x_n^1, f^1, \text{lab})$ that satisfy all the conditions:
 - For all $i \in [n]$, $[\text{QEnc}(i, x_i^0, x_i^1, \text{lab})$ is queried] or $[x_i^0 = x_i^1 \in \mathcal{M}_i$ and $i \in \mathcal{CS}]$.
 - $\text{QKeyGen}(f^0, f^1)$ are queried.
 - $f^0(x_1^0, \dots, x_n^0) \neq f^1(x_1^1, \dots, x_n^1)$.
- When $\text{xx} = \text{sta}$: the adversary cannot query QCor after querying QEnc or QKeyGen even once.
- When $\text{xx} = \text{sel}$: the adversary must make all queries in one shot. That is, first it outputs $(\mathcal{CS}, \{i, x_i^0, x_i^1, \text{lab}\}, \{f^0, f^1\})$ and obtains the response: $(\{\text{EK}_i\}_{i \in \mathcal{CS}}, \{\text{Enc}(i, \text{EK}_i, x_i^\beta, \text{lab})\}, \{\text{KeyGen}(\text{MSK}, f^\beta)\})$.
- When $\text{yy} = \text{pos}$: for each $\text{lab} \in \mathcal{L}$, either $q_{c,i,\text{lab}} > 0$ for all $i \in [n]$ or $q_{c,i,\text{lab}} = 0$ for all $i \in [n]$.

Syntax of DDFE. We define the syntax of DDFE. Note that we use an identifier $i \in \mathcal{ID}$ to specify each party while they use PK for identifier in the original definition [CDSG+20]. We describe the syntax of DDFE in the context of MPFE and change some expressions from the original definition. For instance, we use MSK instead of SK for secret keys of each party, public/private inputs for Enc and KeyGen instead of using empty keys, and so on.

Let $\mathcal{ID}, \mathcal{K}, \mathcal{M}$ be an ID space, a key space, and a message space, respectively. \mathcal{K}, \mathcal{M} consist of a public part and a private part, that is, $\mathcal{K} = \mathcal{K}_{\text{pri}} \times \mathcal{K}_{\text{pub}}, \mathcal{M} = \mathcal{M}_{\text{pri}} \times \mathcal{M}_{\text{pub}}$. Let f be a function such that $f : \bigcup_{i \in \mathbb{N}} (\mathcal{ID} \times \mathcal{K})^i \times \bigcup_{i \in \mathbb{N}} (\mathcal{ID} \times \mathcal{M})^i \rightarrow \mathcal{Z}$. An DDFE scheme for f consists of five algorithms.

GSetup($1^\lambda, 1^n$): It takes a security parameter 1^λ and outputs a public parameter PP . The other algorithms implicitly take PP .

LSetup(PP): It takes PP and outputs local public parameter PK_i and a master secret key MSK_i . The following three algorithms implicitly take PK_i .

KeyGen($\text{MSK}_i, k = (k_{\text{pri}}, k_{\text{pub}})$): It takes MSK_i and $k \in \mathcal{K}$, and outputs a secret key SK_i .

Enc($\text{MSK}_i, m = (m_{\text{pri}}, m_{\text{pub}})$): It takes MSK_i and $m \in \mathcal{M}$, and outputs a ciphertext CT_i .

Dec($\{\text{SK}_i\}_{i \in \mathcal{U}_K}, \{\text{CT}_i\}_{i \in \mathcal{U}_M}$): It takes $\{\text{SK}_i\}_{i \in \mathcal{U}_K}, \{\text{CT}_i\}_{i \in \mathcal{U}_M}$ and outputs a decryption value $d \in \mathcal{Z}$ or a symbol \perp .

¹¹The leakage function captures information that EK_i reveals from SK .

Correctness. An DDFE scheme for f is correct if it satisfies the following condition. For all $\lambda \in \mathbb{N}$, $\{i, k_i\}_{i \in \mathcal{U}_K} \in \bigcup_{i \in \mathbb{N}} (\mathcal{ID} \times \mathcal{K})^i$, $\{i, m_i\}_{i \in \mathcal{U}_M} \in \bigcup_{i \in \mathbb{N}} (\mathcal{ID} \times \mathcal{M})^i$, we have

$$\Pr \left[\begin{array}{l} d = f(\{i, k_i\}_{i \in \mathcal{U}_K}, \{i, m_i\}_{i \in \mathcal{U}_M}) : \\ \text{PP} \leftarrow \text{GSetup}(1^\lambda) \\ \text{PK}_i, \text{MSK}_i \leftarrow \text{LSetup}(\text{PP}) \\ \text{CT}_i \leftarrow \text{Enc}(\text{MSK}_i, m_i) \\ \text{SK}_i \leftarrow \text{KeyGen}(\text{MSK}_i, k_i) \\ d = \text{Dec}(\{\text{SK}_i\}_{i \in \mathcal{U}_K}, \{\text{CT}_i\}_{i \in \mathcal{U}_M}) \end{array} \right] = 1.$$

Note that \mathcal{U}_K and \mathcal{U}_M can be multisets in general, while IP-DDFE only considers the case where \mathcal{U}_K and \mathcal{U}_M are sets. We assume that \mathbb{N} contains 0 here and $(\mathcal{ID} \times \mathcal{K})^0 = \{i, k_i\}_{i \in \emptyset} = \emptyset$. That is, \mathcal{U}_K and \mathcal{U}_M can be an empty set, which corresponds to the case where Dec does not take secret keys/ciphertexts as input.

Security. We naturally extend the security definition for DDFE in [CDSG+20] to the function-hiding setting as follows.

Definition 6.2 (Function-hiding security of DDFE). An DDFE scheme is xx-yy -function-hiding ($\text{xx} \in \{\text{sel}, \text{adt}\}$, $\text{yy} \in \{\text{sym}, \text{asym}\}$) if for every stateful PPT adversary \mathcal{A} , there exists a negligible function $\text{negl}(\cdot)$ such that for all $\lambda \in \mathbb{N}$, the following holds

$$\Pr \left[\beta \leftarrow \mathcal{A}^{\text{QHonestGen}(), \text{QCor}(), \text{QEnc}(), \text{QKeyGen}()}(\text{PP}) : \begin{array}{l} \beta \leftarrow \{0, 1\} \\ \text{PP} \leftarrow \text{GSetup}(1^\lambda) \end{array} \right] \leq \frac{1}{2} + \text{negl}(\lambda)$$

where $\text{QHonestGen}(i)$ runs $\text{PK}_i, \text{MSK}_i \leftarrow \text{LSetup}(\text{PP})$ and returns PK_i . For i such that $\text{QHonestGen}(i)$ was queried, the adversary can make the following queries: $\text{QCor}(i)$ outputs MSK_i , $\text{QEnc}(i, m^0, m^1)$ outputs $\text{Enc}(\text{MSK}_i, m^\beta)$, and $\text{QKeyGen}(i, k^0, k^1)$ outputs $\text{KeyGen}(\text{MSK}_i, k^\beta)$. Note that k^β and m^β consist of the private elements $k_{\text{pri}}^\beta, m_{\text{pri}}^\beta$ and the public elements $k_{\text{pub}}, m_{\text{pub}}$, respectively (we always require that $k_{\text{pub}}^0 = k_{\text{pub}}^1 = k_{\text{pub}}$ and $m_{\text{pub}}^0 = m_{\text{pub}}^1 = m_{\text{pub}}$ as the public elements are not hidden in SK or CT). Let \mathcal{S} be the set of parties on which $\text{QHonestGen}(i)$ is queried, \mathcal{HS} be the set of parties on which the adversary has not queried QCor at the end of the game, and $\mathcal{CS} = \mathcal{S} \setminus \mathcal{HS}$. Then, the adversary's queries must satisfy the following conditions.

- There are no sequences $(\{i, k_i^0\}_{i \in \mathcal{U}_K}, \{i, m_i^0\}_{i \in \mathcal{U}_M})$ and $(\{i, k_i^1\}_{i \in \mathcal{U}_K}, \{i, m_i^1\}_{i \in \mathcal{U}_M})$ that satisfy all the conditions:
 - For all $i \in \mathcal{U}_K$, $[\text{QKeyGen}(i, k_i^0, k_i^1)$ is queried] or $[k_i^0 = k_i^1 \in \mathcal{K}$ and $i \in \mathcal{CS}_K]$.
 - For all $i \in \mathcal{U}_M$, $[\text{QEnc}(i, m_i^0, m_i^1)$ is queried] or $[m_i^0 = m_i^1 \in \mathcal{M}$ and $i \in \mathcal{CS}_M]$.
 - $f(\{i, k_i^0\}_{i \in \mathcal{U}_K}, \{i, m_i^0\}_{i \in \mathcal{U}_M}) \neq f(\{i, k_i^1\}_{i \in \mathcal{U}_K}, \{i, m_i^1\}_{i \in \mathcal{U}_M})$.
- When $\text{xx} = \text{sel}$: the adversary first generates a set \mathcal{S} of honest users in one shot. After that it makes the corruption, key generation, encryption queries in one shot to obtain $\{\text{MSK}_i\}, \{\text{KeyGen}(\text{MSK}_i, k^\beta)\}, \{\text{Enc}(\text{MSK}_i, m^\beta)\}$.
- When $\text{yy} = \text{sym}$: for $i \in \mathcal{CS}$, the queries $\text{QKeyGen}(i, k^0, k^1)$ and $\text{QEnc}(i, m^0, m^1)$ must satisfy $k^0 = k^1$ and $m^0 = m^1$, respectively¹².

Definition 6.3 (Inner Product Functional Encryption (IPFE)). Let $\Pi = (p, \mathbb{G}_1, \mathbb{G}_2, \mathbb{G}_T, e, g_1, g_2)$ be bilinear groups. IPFE for Π is a class of FE where $\mathcal{M} = \mathbb{G}_1^N$, and function $f \in \mathcal{F}$ is represented by $[\mathbf{y}]_2 \in \mathbb{G}_2^N$ where $\mathbf{y} \in \mathbb{Z}_p^N$ and defined as $f([\mathbf{x}]_1) = [\langle \mathbf{x}, \mathbf{y} \rangle]_T$. We say IPFE is function-hiding if it has both message and function privacy as per Section A.1.2.

¹²The symmetric setting captures the case where MSK_i can be used to not only encrypt/key generation but also decryption/decoding of CT_i/SK_i .

Definition 6.4 (IP-MCFE). Let $B \in \mathbb{N}$ be a bound of the infinity norm of vectors. IP-MCFE is a class of MCFE where $\mathcal{M}_i = [-B, B]^{\mathbb{N}}$, $\mathcal{Z} = \mathbb{Z}$, and $\mathcal{L} = \{0, 1\}^*$. The function f is represented by $\mathbf{y} \in [-B, B]^{n\mathbb{N}}$ and defined as $f(\mathbf{x}_1, \dots, \mathbf{x}_n) = \langle (\mathbf{x}_1 || \dots || \mathbf{x}_n), \mathbf{y} \rangle$.

Definition 6.5 (IP-DDFE). Let $B \in \mathbb{N}$ be a bound of the infinity norm of vectors. IP-DDFE is a class of DDFE where $\mathcal{ID} = \{0, 1\}^*$, $\mathcal{K}_{\text{pri}} = \mathcal{M}_{\text{pri}} = [-B, B]^{\mathbb{N}}$, $\mathcal{K}_{\text{pub}} = \mathcal{M}_{\text{pub}} = 2^{\mathcal{ID}} \times \mathcal{L}$, $\mathcal{Z} = \mathbb{Z}$ for label space $\mathcal{L} = \{0, 1\}^*$. The function f is defined as, for $\{k_i = (\mathbf{y}_i, \mathcal{U}_{K,i}, \text{lab}_{K,i})\}_{i \in \mathcal{U}'_K}$ and $\{m_i = (\mathbf{x}_i, \mathcal{U}_{M,i}, \text{lab}_{M,i})\}_{i \in \mathcal{U}'_M}$,

$$f(\{i, k_i\}_{i \in \mathcal{U}'_K}, \{i, m_i\}_{i \in \mathcal{U}'_M}) = \begin{cases} \sum_{i \in \mathcal{U}'_K} \langle \mathbf{x}_i, \mathbf{y}_i \rangle & \text{the condition below is satisfied} \\ \perp & \text{otherwise} \end{cases}$$

- $\forall i \in \mathcal{U}'_K, \mathcal{U}'_K = \mathcal{U}'_M = \mathcal{U}_{K,i} = \mathcal{U}_{M,i}$.
- $\exists (\text{lab}_K, \text{lab}_M) \in \mathcal{L}^2, \forall i \in \mathcal{U}'_K, \text{lab}_{K,i} = \text{lab}_K, \text{lab}_{M,i} = \text{lab}_M$.

Definition 6.6 (One key-label restriction for IP-DDFE). We define an additional restriction for the adversary in the security game for IP-DDFE. We say an IP-DDFE scheme is xx-yy-function-hiding under the one key-label restriction if it satisfies Definition 6.2 where the adversary's queries additionally satisfy the following condition: QKeyGen with respect to user $i \in \mathcal{ID}$ and label $\text{lab}_K \in \mathcal{L}$ (the query of the form of QKeyGen($i, *, *, *, \text{lab}_K$)) can be made only once for each pair (i, lab_K) .

Definition 6.7 (All-or-nothing encryption (AoNE)). AoNE is a class of DDFE where $\mathcal{ID} = \{0, 1\}^*$, $\mathcal{M}_{\text{pri}} = \{0, 1\}^L$, $\mathcal{M}_{\text{pub}} = 2^{\mathcal{ID}} \times \mathcal{L}$, $\mathcal{K} = \emptyset$, $\mathcal{Z} = \{0, 1\}^*$. The function f is defined as, for $\mathcal{U}'_K \in 2^{\mathcal{ID}}$ and $\{m_i = (x_i, \mathcal{U}_{M,i}, \text{lab}_{M,i})\}_{i \in \mathcal{U}'_M}$,

$$f(\{i\}_{i \in \mathcal{U}'_K}, \{i, m_i\}_{i \in \mathcal{U}'_M}) = \begin{cases} \{x_i\}_{i \in \mathcal{U}'_M} & \text{the condition below is satisfied} \\ \perp & \text{otherwise} \end{cases}$$

- $\forall i \in \mathcal{U}'_M, \mathcal{U}'_M = \mathcal{U}_{M,i}$.
- $\exists \text{lab}_M \in \mathcal{L}, \forall i \in \mathcal{U}'_M, \text{lab}_{M,i} = \text{lab}_M$.

This means that KeyGen is unnecessary, and Dec works without taking secret keys as input in AoNE (recall that \mathcal{U}'_K can be an empty set).

Chotard et al. showed that sel-sym-IND-secure AoNE can be generically constructed from identity-based encryption [CDSG⁺20]¹³. We also use pseudorandom functions and non-interactive key exchange with quite simple requirements, which can be realized by the original Diffie-Hellman key exchange. We formally define it in Section 2.3.

6.1 Construction of Function-Hiding IP-MCFE

We first construct a function-hiding IP-MCFE scheme for a step to a function-hiding IP-DDFE scheme, while this is the first instantiation of function-hiding IP-MCFE, which is of independent interest. Let $\Pi = (p, \mathbb{G}_1, \mathbb{G}_2, \mathbb{G}_T, e, g_1, g_2)$ be bilinear groups. Let iFE = (iSetup, iKeyGen, iEnc, Dec) be a function-hiding IPFE scheme and $H : \mathcal{L} \rightarrow \mathbb{G}_1$ be a hash function modeled as a random oracle. The construction of function hiding IP-MCFE is provided in Figure 6.1.

¹³In AoNE, there are no secret keys and thus the IND-security defined in [CDSG⁺20] is exactly the same as function-hiding security in our paper.

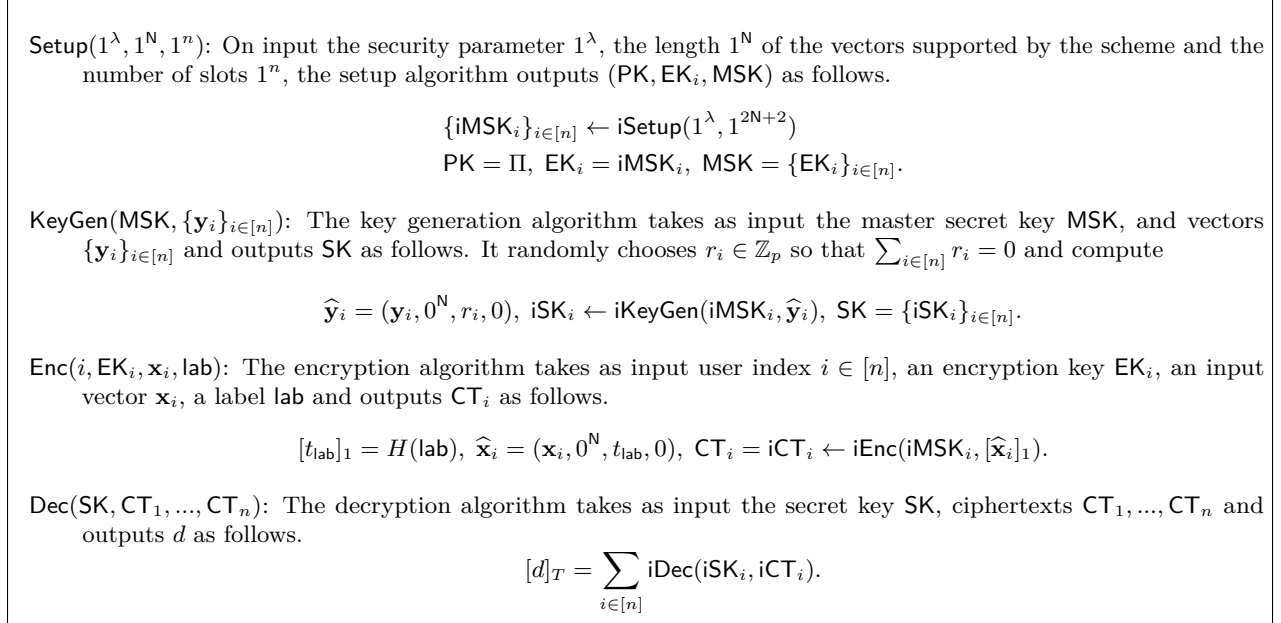


Figure 6.1: Function-Hiding IP-MCFE

Correctness and Security. For correctly generated $(\text{SK}, \text{CT}_1, \dots, \text{CT}_n)$ for $\{\mathbf{y}_i, \mathbf{x}_i\}$, we have

$$\sum_{i \in [n]} \text{iDec}(\text{iSK}_i, \text{iCT}_i) = \left[\sum_{i \in [n]} \langle \widehat{\mathbf{x}}_i, \widehat{\mathbf{y}}_i \rangle \right]_T = \left[\sum_{i \in [n]} \langle \mathbf{x}_i, \mathbf{y}_i \rangle \right]_T.$$

In our scheme, EK_i has a power to decode both CT_i and SK_i since EK_i is a part of MSK . This is captured as the function Leak_y below.

Theorem 6.8. If the SXDH assumption holds in \mathbb{G}_1 and iFE is function-hiding, then our IP-MCFE scheme is Leak_y -sel-pos-function-hiding in the random oracle model, where $\text{Leak}_y(i, \{\mathbf{y}_i\}_{i \in [n]}) = \mathbf{y}_i$.

Proof. Let $\mathcal{HS} \subseteq [n]$ be the set of uncorrupted parties and $\mathcal{CS} = [n] \setminus \mathcal{HS}$. Let $\mathcal{LS} \subset \mathcal{L}$ be a set of labels queried by the adversary. Let q_k be the maximum number of key generation queries. We prove the theorem via a series of hybrids, which are defined as follows. Basically we denote a variable v used in the ℓ -th encryption query for label lab in slot i by $v_{\text{lab}, i}^\ell$ but often omit ℓ and lab when we uniformly handle the variable in all queries. We use similar omission for key generation query.

G^β : The original game. Especially, the challenger sets $\widehat{\mathbf{x}}_i = (\mathbf{x}_i^\beta, 0^N, t_{\text{lab}}, 0)$, $\widehat{\mathbf{y}}_i = (\mathbf{y}_i^\beta, 0^N, r_i, 0)$ for the reply to the encryption query and the key generation query, respectively.

H_0 : This hybrid is the same as G^β except that the challenger sets $\widehat{\mathbf{x}}_i = (\mathbf{x}_i^\beta, \mathbf{x}_i^0, t_{\text{lab}}, 0)$ for the reply to $\text{QEnc}(i, \mathbf{x}_i^0, \mathbf{x}_i^1, \text{lab})$ with $i \in \mathcal{HS}$. The indistinguishability between G^β and H_0 directly follows from the function-hiding property of iFE , since $(\widehat{\mathbf{x}}_i, \widehat{\mathbf{y}}_i)$ in G^β and that in H_0 are the same for all queries of $i \in \mathcal{HS}$.

$H_j (j \in [q_k])$: In this hybrid, the challenger changes $\widehat{\mathbf{y}}_i$ in the reply to $\text{QKeyGen}(\{\mathbf{y}_i^0\}, \{\mathbf{y}_i^1\})$ from H_0 . That is, for $i \in \mathcal{HS}$, it defines $\widehat{\mathbf{y}}_i = (0^N, \mathbf{y}_i^0, r_i, 0)$ for the first j queries and $\widehat{\mathbf{y}}_i = (\mathbf{y}_i^\beta, 0^N, r_i, 0)$ for the last $q_k - j$ queries. The indistinguishability between H_{j-1} and H_j for $j \in [q_k]$ is proven in Lemma 6.9.

H_f : This hybrid is the same as H_{q_k} except that the challenger set $\widehat{\mathbf{x}}_i = (0^N, \mathbf{x}_i^0, t_{\text{lab}}, 0)$ instead of $(\mathbf{x}_i^\beta, \mathbf{x}_i^0, t_{\text{lab}}, 0)$ for all encryption queries with $i \in \mathcal{HS}$. The indistinguishability between H_{q_k} and H_f directly follows from the function-hiding property of iFE, since $\langle \widehat{\mathbf{x}}_i, \widehat{\mathbf{y}}_i \rangle$ in H_{q_k} and that in H_f are the same for all queries with $i \in \mathcal{HS}$. In this hybrid, the adversary's advantage is 0 since its view is independent of β (recall that $x_i^0 = x_i^1$ and $y_i^0 = y_i^1$ for all $i \in \mathcal{CS}$).

□

Lemma 6.9. If the SXDH assumption holds in \mathbb{G}_1 and iFE is function-hiding, then H_{j-1} and H_j are indistinguishable for $j \in [q_k]$.

Proof. We introduce intermediate hybrids to prove the lemma as follows.

$H_{j-1,1}$: This hybrids is the same as H_{j-1} except that, for $i \in \mathcal{HS}$, the challenger sets

$$\widehat{\mathbf{x}}_i = (\mathbf{x}_i^\beta, \mathbf{x}_i^0, t_{\text{lab}}, \underline{r_i^j t_{\text{lab}} + \langle \mathbf{x}_i^\beta, \mathbf{y}_i^{j,\beta} \rangle}), \widehat{\mathbf{y}}_i = \begin{cases} (0^N, \mathbf{y}_i^0, r_i, 0) & \text{(the first } j-1 \text{ queries)} \\ (0^N, 0^N, 0, 1) & \text{(the } j\text{-th query)} \\ (\mathbf{y}_i^\beta, 0^N, r_i, 0) & \text{(the last } q_k - j \text{ queries)} \end{cases}$$

where r_i^j are the random elements chosen in the j -th key generation query. The indistinguishability between H_{j-1} and $H_{j-1,1}$ directly follows from the security of the function-hiding property of iFE, since $\langle \widehat{\mathbf{x}}_i, \widehat{\mathbf{y}}_i \rangle$ in H_{j-1} and that in $H_{j-1,1}$ are the same for all queries with $i \in \mathcal{HS}$.

$H_{j-1,2}$: This hybrids is the same as $H_{j-1,1}$ except that, for all $\text{lab} \in \mathcal{CS}$, and $i \in \mathcal{HS}$, the challenger first randomly choose $\widehat{t}_{\text{lab},i} \in \mathbb{Z}_p$ so that $\sum_{i \in \mathcal{HS}} \widehat{t}_{\text{lab},i} = -t_{\text{lab}} \sum_{i \in \mathcal{CS}} r_i^j$. Then it sets, for $i \in \mathcal{HS}$,

$$\widehat{\mathbf{x}}_i = (\mathbf{x}_i^\beta, \mathbf{x}_i^0, t_{\text{lab}}, \widehat{t}_{\text{lab},i} + \langle \mathbf{x}_i^\beta, \mathbf{y}_i^{j,\beta} \rangle).$$

The indistinguishability between $H_{j-1,1}$ and $H_{j-1,2}$ can be proven under SXDH as follows. The random self-reducibility of SXDH implies $[\{a_\mu, a_\mu \mathbf{b}\}_{\mu \in [\ell]}]_1 \approx_c [\{a_\mu, \mathbf{u}_\mu\}_{\mu \in [\ell]}]_1$ where $a_\mu \leftarrow \mathbb{Z}_p$, $\mathbf{b} = (b_1, \dots, b_h) \leftarrow \mathbb{Z}_p^h$, $\mathbf{u}_\mu = (u_{\mu,1}, \dots, u_{\mu,h}) \leftarrow \mathbb{Z}_p^h$, $\ell = |\mathcal{CS}|$ and $h = |\mathcal{HS}| - 1$. Without loss of generality, we can assume $\mathcal{HS} = \{1, \dots, h+1\}$ and $\mathcal{CS} = \{h+2, \dots, n\}$. We implicitly define $t_{\text{lab}} = a_{\text{L}(\text{lab})}$, $r_i^j = b_i$ (for $i \in [h]$), $r_{h+1}^j = -\sum_{i \in [h]} b_i - \sum_{i \in \mathcal{CS}} r_i^j$, $\widehat{t}_{\text{lab},i} = u_{\text{L}(\text{lab}),i}$, where $\text{L} : \mathcal{CS} \rightarrow [\ell]$ is some bijective function. Then, we can construct a reduction that chooses $r_i^j \leftarrow \mathbb{Z}_p$ for $i \in \mathcal{CS}$ and simulates vectors in queries as

$$\widehat{\mathbf{x}}_i = \begin{cases} (\mathbf{x}_i^\beta, \mathbf{x}_i^0, a_{\text{L}(\text{lab})}, z_{\beta, \text{L}(\text{lab}),i} + \langle \mathbf{x}_i^\beta, \mathbf{y}_i^{j,\beta} \rangle) & (i \in [h]) \\ (\mathbf{x}_i^\beta, \mathbf{x}_i^0, a_{\text{L}(\text{lab})}, -\sum_{i' \in [h]} z_{\beta, \text{L}(\text{lab}),i'} - a_{\text{L}(\text{lab})} \sum_{i' \in \mathcal{CS}} r_{i'}^j + \langle \mathbf{x}_i^\beta, \mathbf{y}_i^{j,\beta} \rangle) & (i = h+1) \\ (\mathbf{x}_i^\beta, 0, a_{\text{L}(\text{lab})}, 0) & (i \in \mathcal{CS}) \end{cases}$$

$$\widehat{\mathbf{y}}_i = \begin{cases} (0^N, \mathbf{y}_i^0, r_i, 0) & \text{(the first } j-1 \text{ queries for } i \in \mathcal{HS}) \\ (0^N, 0^N, 0, 1) & \text{(the } j\text{-th query for } i \in \mathcal{HS}) \\ (\mathbf{y}_i^\beta, 0^N, r_i^j, 0) & \text{(the } j\text{-th query for } i \in \mathcal{CS}) \\ (\mathbf{y}_i^\beta, 0^N, r_i, 0) & \text{(the last } q_k - j \text{ queries for } i \in \mathcal{HS} \text{ and all for } i \in \mathcal{CS}) \end{cases}$$

where $z_{0,\mu,i} = a_\mu b_i$ and $z_{1,\mu,i} = u_{\mu,i}$.

$H_{j-1,3}$: This hybrids is the same as $H_{j-1,2}$ except that, for all $i \in \mathcal{HS}$, the challenger sets $\widehat{\mathbf{x}}_i$ in encryption queries as

$$\widehat{\mathbf{x}}_i = (\mathbf{x}_i^\beta, \mathbf{x}_i^0, t_{\text{lab}}, \widehat{t}_{\text{lab},i} + \underline{\langle \mathbf{x}_i^0, \mathbf{y}_i^{j,0} \rangle}).$$

This change is information-theoretic, which can be proven as follows. Let $(i, \mathbf{x}_{\text{lab},i}^{\ell,0}, \mathbf{x}_{\text{lab},i}^{\ell,1}, \text{lab})$ be the ℓ -th encryption query of the form of $(i, *, *, \text{lab})$. Due to the query condition, we have

$$\Delta_{\text{lab},i}^\beta = \mathbf{x}_{\text{lab},i}^{\ell,\beta} \mathbf{y}_i^{j,\beta} - \mathbf{x}_{\text{lab},i}^{\ell,0} \mathbf{y}_i^{j,0} = \mathbf{x}_{\text{lab},i}^{1,\beta} \mathbf{y}_i^{j,\beta} - \mathbf{x}_{\text{lab},i}^{1,0} \mathbf{y}_i^{j,0} \quad \text{for all } \ell.$$

Thus, we can define $\widehat{t}_{\text{lab},i} = \widehat{t}'_{\text{lab},i} - \Delta_{\text{lab},i}^\beta$ and obtain

$$\widehat{\mathbf{x}}_i = (\mathbf{x}_i^\beta, \mathbf{x}_i^0, t_{\text{lab}}, \widehat{t}_{\text{lab},i} + \langle \mathbf{x}_{\text{lab},i}^{\ell,\beta}, \mathbf{y}_i^{j,\beta} \rangle) = (\mathbf{x}_i^\beta, \mathbf{x}_i^0, t_{\text{lab}}, \widehat{t}'_{\text{lab},i} + \langle \mathbf{x}_{\text{lab},i}^{\ell,0}, \mathbf{y}_i^{j,0} \rangle).$$

Observe that both $\{\widehat{t}_{\text{lab},i}\}_{i \in \mathcal{HS}}$ and $\{\widehat{t}'_{\text{lab},i}\}_{i \in \mathcal{HS}}$ are secret shares of $-\widehat{t}_{\text{lab}} \sum_{i \in \mathcal{CS}} r_i^j$ since $\sum_{i \in \mathcal{HS}} \Delta_{\text{lab},i}^\beta = 0$ due to the query condition. Hence, they are identically distributed. The indistinguishability between $H_{j-1,3}$ and H_j can be proven similarly to that between H_{j-1} and $H_{j-1,2}$.

□

6.2 Construction of Function-Hiding IP-DDFE

We next construct our function-hiding IP-DDFE scheme. Intuitively, our IP-DDFE scheme instantiates our IP-MCFE scheme in parallel per each party set via a pseudorandom function in a non-black box manner. Nevertheless, in the security proof, we can delete the information of the challenge bit β in a hybrid sequence similarly to the security proof of IP-MCFE.

Let $\text{iFE} = (\text{iSetup}, \text{iKeyGen}, \text{iEnc}, \text{Dec})$ be a function-hiding IPFE scheme with the length of the random tape for $\text{iSetup}(1^\lambda, 1^{2N+2})$ being $p(\lambda, N)$, $\text{AoNE} = (\text{aGSetup}, \text{aLSetup}, \text{aEnc}, \text{aDec})$ be an all-or-nothing encryption scheme, $\text{NIKE} = (\text{nSetup}, \text{nKeyGen}, \text{nSharedKey})$ be a non-interactive key exchange scheme, $\{\text{PRF}_1^K\} : \mathcal{L} \rightarrow \mathbb{Z}_p$, $\{\text{PRF}_2^K\} : 2^{\mathcal{ID}} \rightarrow \{0, 1\}^{p(\lambda, N)}$ be families of pseudorandom functions where \mathcal{ID} denotes an identity space, and $H : 2^{\mathcal{ID}} \times \mathcal{L} \rightarrow \mathbb{G}_1$ is a hash function modeled as a random oracle. Let $\mathcal{K}_1, \mathcal{K}_2$ be key spaces of $\text{PRF}_1, \text{PRF}_2$. We assume that the range of nSharedKey and the key space for PRF_1 are the same, namely, \mathcal{K}_1 . Our construction is provided in Figure 6.2.

Correctness and Security. Thanks to the correctness of AoNE , we have $\widetilde{\text{iCT}}_i = \text{iCT}_i$, $\widetilde{\text{iSK}}_i = \text{iSK}_i$. For all $\text{lab}_K, \{\mathcal{K}_{i,j,1}\}, \mathcal{U}$, we have

$$\sum_{i \in \mathcal{U}} r_i = \sum_{i \in \mathcal{U}} \sum_{\substack{j \in \mathcal{U} \\ i \neq j}} (-1)^{j < i} \text{PRF}_1^{\mathcal{K}_{i,j,1}}(\text{lab}_K) = 0$$

since $\mathcal{K}_{i,j,1} = \mathcal{K}_{j,i,1}$. For all $i \in \mathcal{U}$, iSK_i and iCT_i are generated under the same iMSK_i since they are generated using the same random tape $\text{PRF}_2^{\mathcal{K}_{i,2}}(\mathcal{U})$. Thus, thanks to the correctness of iFE , we have $\sum_{i \in \mathcal{U}} \text{iDec}(\widetilde{\text{iSK}}_i, \widetilde{\text{iCT}}_i) = [\sum_{i \in \mathcal{U}} \langle \widehat{\mathbf{x}}_i, \widehat{\mathbf{y}}_i \rangle]_T = [\sum_{i \in \mathcal{U}} \langle \mathbf{x}_i, \mathbf{y}_i \rangle]_T$.

We show security via the following theorem.

Theorem 6.10. If $\{\text{PRF}_1^K\}, \{\text{PRF}_2^K\}$ are families of pseudorandom functions, NIKE is IND-secure, AoNE is sel-sym-IND-secure, the SXDH assumption holds in \mathbb{G}_1 , and iFE is function-hiding, then our IP-DDFE scheme is sel-sym-function-hiding under the one key-label restriction in the random oracle model.

Proof. Let \mathcal{S} be the set of parties generated by honest-party generation queries. Let $\mathcal{HS} \subseteq \mathcal{S}$ be the set of uncorrupted parties and $\mathcal{CS} = \mathcal{S} \setminus \mathcal{HS}$. Let $\mathcal{LS} \subset \mathcal{L}$ be a set of labels queried by the adversary. We prove the theorem via a series of hybrids, which are defined as follows.

\mathcal{G}^β : The original game. Especially, the challenger sets $\widehat{\mathbf{x}}_i = (\mathbf{x}_i^\beta, 0^N, t, 0)$, $\widehat{\mathbf{y}}_i = (\mathbf{y}_i^\beta, 0^N, r_i, 0)$ for the reply to the encryption query and the key generation query, respectively (the procedures in eq. 6.1 and 6.3).

\mathcal{H}_1 : In this hybrid, the challenger uses random functions $R_{i,2}$ instead of $\text{PRF}_2^{\mathcal{K}_{i,2}}$ for $i \in \mathcal{HS}$ in key generation and encryption queries. The indistinguishability directly follows from the security of the pseudorandom function.

GSetup($1^\lambda, 1^N$): On input the security parameter 1^λ and the length 1^N of the vectors supported by the scheme, the setup algorithm outputs PK as follows.

$$\begin{aligned}\Pi &= (p, \mathbb{G}_1, \mathbb{G}_2, \mathbb{G}_T, e, g_1, g_2) \leftarrow \text{Gen}(1^\lambda) \\ \text{aPP} &\leftarrow \text{aGSetup}(1^\lambda), \text{nPP} \leftarrow \text{nSetup}(1^\lambda), \text{PP} = (\Pi, \text{aPP}, \text{nPP}, 1^N).\end{aligned}$$

LSetup(PP): On input PP, user $i \in \mathcal{ID}$ generates $(\text{PK}_i, \text{MSK}_i)$ via the setup algorithm as follows.

$$\begin{aligned}\text{nPK}_i, \text{nSK}_i &\leftarrow \text{nKeyGen}(\text{nPP}), \text{aPK}_i, \text{aMSK}_i \leftarrow \text{aLSetup}(\text{aPP}), K_{i,2} \leftarrow \mathcal{K}_2 \\ \text{PK}_i &= (\text{nPK}_i, \text{aPK}_i), \text{MSK}_i = (\text{nSK}_i, \text{aMSK}_i, K_{i,2}).\end{aligned}$$

KeyGen(MSK_i, k): The key generation algorithm the master secret key MSK_i , and an input $k = (\mathbf{y}_i, \mathcal{U}_K, \text{lab}_K)$ such that $i \in \mathcal{U}_K$ and outputs SK_i as follows.

$$\begin{aligned}\text{rt}_i &= \text{PRF}_2^{\text{K}_{i,2}}(\mathcal{U}_K), \text{iMSK}_i = \text{iSetup}(1^\lambda, 1^{2N+2}; \text{rt}_i), K_{i,j,1} \leftarrow \text{nSharedKey}(\text{nSK}_i, \text{nPK}_j) \\ r_i &= \sum_{\substack{j \in \mathcal{U}_K \\ i \neq j}} (-1)^{j < i} \text{PRF}_1^{\text{K}_{i,j,1}}(\text{lab}_K), \widehat{\mathbf{y}}_i = (\mathbf{y}_i, 0^N, r_i, 0), \text{iSK}_i \leftarrow \text{iKeyGen}(\text{iMSK}_i, \widehat{\mathbf{y}}_i) \quad (6.1)\end{aligned}$$

$$\text{aCT}_i \leftarrow \text{aEnc}(\text{aMSK}_i, (\text{iSK}_i, \mathcal{U}_K, \text{lab}_K)), \text{SK}_i = (\text{aCT}_i, \mathcal{U}_K, \text{lab}_K). \quad (6.2)$$

Enc(MSK_i, m): The encryption algorithm takes as input the public parameters PK, the master secret key MSK_i , and an input $m = (\mathbf{x}_i, \mathcal{U}_M, \text{lab}_M)$ such that $i \in \mathcal{U}_M$ and outputs CT_i as follows.

$$\begin{aligned}\text{rt}_i &= \text{PRF}_2^{\text{K}_{i,2}}(\mathcal{U}_M), \text{iMSK}_i = \text{iSetup}(1^\lambda, 1^{2N+2}; \text{rt}_i), [t]_1 = H(\mathcal{U}_M, \text{lab}_M) \\ \widehat{\mathbf{x}}_i &= (\mathbf{x}_i, 0^N, t, 0), \text{iCT}_i \leftarrow \text{iEnc}(\text{iMSK}_i, [\widehat{\mathbf{x}}_i]_1) \quad (6.3)\end{aligned}$$

$$\text{aCT}_i \leftarrow \text{aEnc}(\text{aMSK}_i, (\text{iCT}_i, \mathcal{U}_M, \text{lab}_M)), \text{CT}_i = (\text{aCT}_i, \mathcal{U}_M, \text{lab}_M). \quad (6.4)$$

Dec($\{\text{SK}_i\}_{i \in \mathcal{U}_K}, \{\text{CT}_i\}_{i \in \mathcal{U}_M}$): The decryption algorithm takes as input the public parameters PK, secret keys $\{\text{SK}_i\}_{i \in \mathcal{U}_K}$, ciphertexts $\{\text{CT}_i\}_{i \in \mathcal{U}_M}$ such that $\mathcal{U} = \mathcal{U}_K = \mathcal{U}_M$ and outputs d as follows. Perse $\text{SK}_i = (\text{aCT}_i, \mathcal{U}_K, \text{lab}_K)$ and $\text{CT}_i = (\text{aCT}'_i, \mathcal{U}_M, \text{lab}_M)$. Compute

$$\widetilde{\text{iSK}}_i = \text{aDec}(\{\text{aCT}_i\}_{i \in \mathcal{U}}), \widetilde{\text{iCT}}_i = \text{aDec}(\{\text{aCT}'_i\}_{i \in \mathcal{U}}), [d]_T = \sum_{i \in \mathcal{U}} \text{iDec}(\widetilde{\text{iSK}}_i, \widetilde{\text{iCT}}_i).$$

Figure 6.2: Function Hiding IP-DDFE

H_2 : We say a key generation query on $\widehat{k} = (i, \mathbf{y}_i^0, \mathbf{y}_i^1, \mathcal{U}_K, \text{lab}_K)$ is incomplete if there exists $i' \in \mathcal{U}_K$ such that $i' \in \mathcal{HS}$ and the query of the form $\widehat{k} = (i', *, *, \mathcal{U}_K, \text{lab}_K)$ is not made. In this hybrid, for all incomplete key generation queries, aCT_i is changed to the encryption of $(0, \mathcal{U}_K, \text{lab}_K)$ (the procedure in eq. 6.2). The indistinguishability directly follows from the security of AoNE.

H_3 : We say an encryption query on $\widehat{m} = (i, \mathbf{x}_i^0, \mathbf{x}_i^1, \mathcal{U}_M, \text{lab}_M)$ is incomplete if there exists $i' \in \mathcal{U}_M$ such that $i' \in \mathcal{HS}$ and the query of the form $\widehat{m} = (i', *, *, \mathcal{U}_M, \text{lab}_M)$ is not made. In this hybrid, for all incomplete encryption queries, aCT_i is changed to the encryption of $(0, \mathcal{U}_M, \text{lab}_M)$ (the procedure in eq. 6.4). The indistinguishability directly follows from the security of AoNE.

H_f : In this hybrid, for all complete queries, which are not incomplete, the challenger sets $\widehat{\mathbf{x}}_i = (0^N, \mathbf{x}_i^0, t, 0)$, $\widehat{\mathbf{y}}_i = (0^N, \mathbf{y}_i^0, r_i, 0)$ for the reply to the encryption query and the key generation query on $i \in \mathcal{HS}$, respectively (the procedures in eq. 6.1 and 6.3). In this hybrid, the adversary's advantage is 0 since its view is independent of β (recall that $x_i^0 = x_i^1$ and $y_i^0 = y_i^1$ for all $i \in \mathcal{CS}$). The indistinguishability from H_3 is shown in Lemma 6.11.

□

Lemma 6.11. If the SXDH assumption holds in \mathbb{G}_1 and iFE is function-hiding, then \mathbf{H}_3 and \mathbf{H}_f are indistinguishable in the random oracle model.

Proof. We say $\mathcal{U} \in 2^{\mathcal{I}^D}$ is a complete ID set if \mathcal{U} is queried in either complete key generation queries or complete encryption query. Let q_u be the number of complete ID sets in the adversary's queries. We always put the complete ID sets in some order as $\mathcal{U}_1, \dots, \mathcal{U}_{q_u}$, e.g., in ascending order with respect to the size of ID set. To prove the lemma, we introduce hybrids $\widehat{\mathbf{H}}_j$ for $j \in [q'_u]$ between \mathbf{H}_3 and \mathbf{H}_f as follows, where q'_u is the upper bound of q_u .

$\widehat{\mathbf{H}}_j$ ($j \in [q'_u]$): This hybrid is the same as \mathbf{H}_3 except that, for complete queries of the form $\widehat{m} = (i, \mathbf{x}_i^0, \mathbf{x}_i^1, \mathcal{U}_M, \text{lab}_M)$ and $\widehat{k} = (i, \mathbf{y}_i^0, \mathbf{y}_i^1, \mathcal{U}_K, \text{lab}_K)$ such that $i \in \mathcal{HS}$, the challenger set

$$\widehat{\mathbf{x}}_i = \begin{cases} (0^N, \mathbf{x}_i^0, t, 0) & (\mathcal{U}_M \in \{\mathcal{U}_1, \dots, \mathcal{U}_j\}) \\ (\mathbf{x}_i^\beta, 0^N, t, 0) & (\mathcal{U}_M \in \{\mathcal{U}_{j+1}, \dots, \mathcal{U}_{q_u}\}) \end{cases}, \widehat{\mathbf{y}}_i = \begin{cases} (0^N, \mathbf{y}_i^0, r_i, 0) & (\mathcal{U}_K \in \{\mathcal{U}_1, \dots, \mathcal{U}_j\}) \\ (\mathbf{y}_i^\beta, 0^N, r_i, 0) & (\mathcal{U}_K \in \{\mathcal{U}_{j+1}, \dots, \mathcal{U}_{q_u}\}) \end{cases}$$

respectively, where $\mathcal{U}_j = \{\perp\}$ for $j > q_u$.

We can see that $\widehat{\mathbf{H}}_{q'_u} = \mathbf{H}_f$. Thus, what we need to do is to prove $\widehat{\mathbf{H}}_{j-1}$ and $\widehat{\mathbf{H}}_j$ are indistinguishable for $j \in [q'_u]$, where we define $\widehat{\mathbf{H}}_0 = \mathbf{H}_3$.

For the purpose, we define intermediate hybrids $\widehat{\mathbf{H}}_{j-1, \kappa}$ for $\kappa \in [q_k]$ where q_k is the maximum number of key generation queries, and which is defined as follows.

$\widehat{\mathbf{H}}_{j-1, \kappa}$ ($\kappa \in [q_k]$): Let $\{\text{lab}_{K, \mathcal{U}_K}^1, \dots, \text{lab}_{K, \mathcal{U}_K}^\kappa\}$ denote a set of labels used in the complete key query of the form $(*, *, *, \mathcal{U}_K, *)$. This hybrid is the same as $\widehat{\mathbf{H}}_{j-1}$ except that, for complete queries of the form $\widehat{m} = (i, \mathbf{x}_i^0, \mathbf{x}_i^1, \mathcal{U}_M, \text{lab}_M)$ and $\widehat{k} = (i, \mathbf{y}_i^0, \mathbf{y}_i^1, \mathcal{U}_K, \text{lab}_{K, \mathcal{U}_K}^\ell)$ such that $i \in \mathcal{HS}$, the challenger set

$$\widehat{\mathbf{x}}_i = \begin{cases} (0^N, \mathbf{x}_i^0, t, 0) & (\mathcal{U}_M \in \{\mathcal{U}_1, \dots, \mathcal{U}_{j-1}\}) \\ (\mathbf{x}_i^\beta, \mathbf{x}_i^0, t, 0) & (\mathcal{U}_M = \mathcal{U}_j) \\ (\mathbf{x}_i^\beta, 0^N, t, 0) & (\mathcal{U}_M \in \{\mathcal{U}_{j+1}, \dots, \mathcal{U}_{q_u}\}) \end{cases}$$

$$\widehat{\mathbf{y}}_i = \begin{cases} (0^N, \mathbf{y}_i^0, r_i, 0) & (\mathcal{U}_K \in \{\mathcal{U}_1, \dots, \mathcal{U}_{j-1}\} \text{ or } (\mathcal{U}_K = \mathcal{U}_j \text{ and } \ell \leq \kappa)) \\ (\mathbf{y}_i^\beta, 0^N, r_i, 0) & (\mathcal{U}_K \in \{\mathcal{U}_{j+1}, \dots, \mathcal{U}_{q_u}\} \text{ or } (\mathcal{U}_K = \mathcal{U}_j \text{ and } \ell > \kappa)) \end{cases}$$

respectively, where $\mathcal{U}_j = \{\perp\}$ for $j > q_u$.

We can prove $\widehat{\mathbf{H}}_{j-1} \approx_c \widehat{\mathbf{H}}_{j-1, 0}$, $\widehat{\mathbf{H}}_{j-1, \kappa-1} \approx_c \widehat{\mathbf{H}}_{j-1, \kappa}$, $\widehat{\mathbf{H}}_{j-1, q_k} \approx_c \widehat{\mathbf{H}}_j$ similarly to Lemma 6.9 where q_k is the maximum number of key generation queries. However, we need a more careful analysis to prove $\widehat{\mathbf{H}}_{j-1, \kappa-1} \approx_c \widehat{\mathbf{H}}_{j-1, \kappa}$ since we additionally utilize NIKE and PRF₁ to prove them following the strategy in Lemma 6.9.

In the following, we focus only on complete encryption queries of the form $(*, *, *, \mathcal{U}_j, *)$ and complete key generation queries of the form $(*, *, *, \mathcal{U}_j, \text{lab}_{K, \mathcal{U}_j}^\kappa)$, since other queries are not changed between $\widehat{\mathbf{H}}_{j-1, \kappa-1}$ and $\widehat{\mathbf{H}}_{j-1, \kappa}$. Let $\mathcal{U}_j \cap \mathcal{HS} = \{u_1, \dots, u_w\}$ and w' be an upper bound of w . We then define intermediate hybrid $\overline{\mathbf{H}}_\eta$ between $\widehat{\mathbf{H}}_{j-1, \kappa-1}$ and $\widehat{\mathbf{H}}_{j-1, \kappa}$.

$\overline{\mathbf{H}}_\eta$ ($\eta \in [w']$): This hybrid is the same as $\widehat{\mathbf{H}}_{j-1, \kappa-1}$ except that, for complete queries of the form $\widehat{m} = (u_i, \mathbf{x}_{\text{lab}_M, u_i}^{\ell, 0}, \mathbf{x}_{\text{lab}_M, u_i}^{\ell, 1}, \mathcal{U}_j, \text{lab}_M)$ and $\widehat{k} = (u_i, \mathbf{y}_{u_i}^0, \mathbf{y}_{u_i}^1, \mathcal{U}_j, \text{lab}_{K, \mathcal{U}_j}^\kappa)$ such that $i \in \mathcal{U}_j \cap \mathcal{HS}$, the challenger

set

$$\widehat{\mathbf{x}}_{u_i} = \begin{cases} (\mathbf{x}_{\text{lab}_{M,u_i}}^{\ell,\beta}, \mathbf{x}_{\text{lab}_{M,u_i}}^{\ell,0}, t, 0) & (i < w) \\ (\mathbf{x}_{\text{lab}_{M,u_i}}^{\ell,\beta}, \mathbf{x}_{\text{lab}_{M,u_i}}^{\ell,0}, t, \sum_{i' \in [\eta]} \Delta_{\text{lab}_{M,u_{i'}}}^{\beta}) & (i = w) \end{cases}$$

$$\widehat{\mathbf{y}}_{u_i} = \begin{cases} (0^N, \mathbf{y}_{u_i}^0, r_{u_i}, 0) & (i \leq \eta) \\ (\mathbf{y}_{u_i}^{\beta}, 0^N, r_{u_i}, 0) & (\eta < i < w) \\ (\mathbf{y}_{u_i}^{\beta}, 0^N, r_{u_i}, \mathbf{1}) & (i = w) \end{cases}$$

respectively, where $(\mathbf{x}_{\text{lab}_{M,u_i}}^{\ell,0}, \mathbf{x}_{\text{lab}_{M,u_i}}^{\ell,1})$ are the vectors in the ℓ -th encryption query of the form $(u_i, *, *, \mathcal{U}_j, \text{lab}_M)$ and $\Delta_{\text{lab}_{M,u_i}}^{\beta} = \mathbf{x}_{\text{lab}_{M,u_i}}^{1,\beta} \mathbf{y}_{u_i}^{\beta} - \mathbf{x}_{\text{lab}_{M,u_i}}^{1,0} \mathbf{y}_{u_i}^0$.

Due to the query condition of the adversary, we have

$$\Delta_{\text{lab}_{M,u_i}}^{\beta} = \mathbf{x}_{\text{lab}_{M,u_i}}^{\ell,\beta} \mathbf{y}_{u_i}^{\beta} - \mathbf{x}_{\text{lab}_{M,u_i}}^{\ell,0} \mathbf{y}_{u_i}^0 \quad \text{for all } \ell \quad (6.5)$$

$$\sum_{i \in [w']} \Delta_{\text{lab}_{M,u_i}}^{\beta} = 0 \quad (6.6)$$

We can easily prove that $\widehat{H}_{j-1,\kappa-1} \approx_c \overline{H}_0$ and $\overline{H}_{w'} \approx_c \widehat{H}_{j-1,\kappa}$ by the function-hiding security of iFE and eq. 6.6. Thus the remaining task is to prove $\overline{H}_{\eta-1} \approx_c \overline{H}_{\eta}$ for $\eta \in [w']$. To prove this, we first change the way of choosing $K_{u_{\eta},u_w,1}$ as $K_{u_{\eta},u_w,1} \leftarrow \mathcal{K}_1$ instead of $K_{u_{\eta},u_w,1} \leftarrow \text{nSharedKey}(\text{nSK}_{u_{\eta}}, \text{nPK}_{u_w})$, which directly follows from IND-security of NIKE. We next change

$$r_{u_{\eta}} = \sum_{\substack{i \in \mathcal{U}_j \cap \mathcal{HS} \\ i \neq u_{\eta}}} (-1)^{i < u_{\eta}} \text{PRF}_1^{K_{u_{\eta},i,1}}(\text{lab}_{K,\mathcal{U}_j}^{\kappa})$$

to

$$r_{u_{\eta}} = \sum_{\substack{i \in \mathcal{U}_j \cap \mathcal{HS} \\ i \notin \{u_{\eta}, u_w\}}} (-1)^{i < u_{\eta}} \text{PRF}_1^{K_{u_{\eta},i,1}}(\text{lab}_{K,\mathcal{U}_j}^{\kappa}) - s_{u_{\eta},u_w}$$

and

$$r_{u_w} = \sum_{\substack{i \in \mathcal{U}_j \cap \mathcal{HS} \\ i \neq u_w}} (-1)^{i < u_w} \text{PRF}_1^{K_{u_w,i,1}}(\text{lab}_{K,\mathcal{U}_j}^{\kappa})$$

to

$$r_{u_w} = \sum_{\substack{i \in \mathcal{U}_j \cap \mathcal{HS} \\ i \notin \{u_{\eta}, u_w\}}} (-1)^{i < u_w} \text{PRF}_1^{K_{u_w,i,1}}(\text{lab}_{K,\mathcal{U}_j}^{\kappa}) + s_{u_{\eta},u_w}$$

where $s_{u_{\eta},u_w} \leftarrow \mathbb{Z}_p$. This indistinguishability directly follows from the security of pseudorandom function PRF_1 . Then, we change encrypted vectors using the function-hiding security of iFE as follows:

$$\widehat{\mathbf{x}}_{u_i} = \begin{cases} (\mathbf{x}_{\text{lab}_{M,u_i}}^{\ell,\beta}, \mathbf{x}_{\text{lab}_{M,u_i}}^{\ell,0}, t, -ts_{u_{\eta},u_w}) & (i = \eta) \\ (\mathbf{x}_{\text{lab}_{M,u_i}}^{\ell,\beta}, \mathbf{x}_{\text{lab}_{M,u_i}}^{\ell,0}, t, \sum_{i' \in [\eta-1]} \Delta_{\text{lab}_{M,u_{i'}}}^{\beta} + ts_{u_{\eta},u_w}) & (i = w) \end{cases}$$

$$\widehat{\mathbf{y}}_{u_i} = \begin{cases} (\mathbf{y}_{u_i}^{\beta}, 0^N, r_{u_i} + s_{u_{\eta},u_w}, \mathbf{1}) & (i = \eta) \\ (\mathbf{y}_{u_i}^{\beta}, 0^N, r_{u_i} - s_{u_{\eta},u_w}, \mathbf{1}) & (i = w) \end{cases}.$$

The remaining things are the similar to Lemma 6.9:

1. change ts_{u_{η},u_w} to a random element $\widehat{s}_{u_{\eta},u_w}$ by the SXDH assumption;

2. implicitly define $\widehat{s}_{u_\eta, u_w} = \widehat{s}'_{u_\eta, u_w} + \Delta_{\text{lab}_M, u_\eta}^\beta$;
3. change $\widehat{\mathbf{y}}_{u_\eta} = (\mathbf{y}_{u_\eta}^\beta, 0^N, r_{u_\eta} + s_{u_\eta, u_w}, 1)$ to $\widehat{\mathbf{y}}_{u_\eta} = (0^N, \mathbf{y}_{u_\eta}^0, r_{u_\eta} + s_{u_\eta, u_w}, 1)$ by using eq. 6.5 and function-hiding security of iFE;
4. go to \overline{H}_η by rewinding changes of vectors by function-hiding security, the SXDH assumption, PRF, and NIKE.

□

7 Distributed Ciphertext Policy ABE

In this section, we show how to distribute the recent construction of succinct ciphertext policy ABE by Agrawal and Yamada [AY20]. We show that, by incurring a linear overhead in the parameter sizes, we could distribute the ABE construction by obtaining a non-interactive distributed setup and key generation such that as long as at least one authority is still honest, the scheme remains secure.

7.1 Definition

In this section, we define the notion of a distributed ciphertext-policy attribute-based encryption. First, we provide the syntax, and later describe the security definition.

Syntax. A distributed attribute-based encryption for predicate class $\mathcal{C} = \{\mathcal{C}_\ell : \{0, 1\}^\ell \rightarrow \{0, 1\}\}_{\ell \in \mathbb{N}}$ and 1-bit message space consists of the following PPT algorithms:

GSetup(1^λ) \rightarrow PP. On input the security parameter λ , the setup algorithm outputs public parameters PP.

LSetup(PP, $1^n, 1^\ell$) \rightarrow (PK, MSK). On input the public parameters PP, number of authorities n , and attribute length ℓ , the authority setup algorithm outputs a pair of master public-secret key (PK, MSK).

KeyGen(MSK $_i$, GID, \mathbf{x}) \rightarrow SK $_{i, \text{GID}, \mathbf{x}}$. The key generation algorithm takes as input an authority master secret key MSK $_i$, global identifier GID , and an attribute $\mathbf{x} \in \{0, 1\}^\ell$. It outputs a partial secret key SK $_{i, \text{GID}, \mathbf{x}}$.

Enc($\{\text{PK}_i\}_{i \in [n]}$, C, μ) \rightarrow CT. The encryption algorithm takes as input the list of public keys $\{\text{PK}_i\}_i$, predicate circuit C , and a message bit μ , and outputs a ciphertext CT.

Dec($\{\text{SK}_{i, \text{GID}, \mathbf{x}}\}_{i \in [n]}$, CT) \rightarrow μ/\perp . On input a list of n partial secret keys $\{\text{SK}_{i, \text{GID}, \mathbf{x}}\}_i$ and a ciphertext CT, the decryption algorithm either outputs a message bit μ or a special string \perp .

We require such an ABE scheme to satisfy the following properties.

Correctness. A distributed ABE scheme is said to be correct if for all $\lambda, n, \ell \in \mathbb{N}$, $C \in \mathcal{C}_\ell$, $\mathbf{x} \in \{0, 1\}^\ell$, $\mu \in \{0, 1\}$, if $C(\mathbf{x}) = 1$ then the following holds:

$$\Pr \left[\text{Dec}(\{\text{SK}_{i, \text{GID}, \mathbf{x}}\}_i, \text{CT}) = \mu : \begin{array}{l} \text{PP} \leftarrow \text{GSetup}(1^\lambda) \\ \forall i \in [n] : (\text{PK}_i, \text{MSK}_i) \leftarrow \text{LSetup}(\text{PP}, 1^n, 1^\ell) \\ \forall i \in [n] : \text{SK}_{i, \text{GID}, \mathbf{x}} \leftarrow \text{KeyGen}(\text{MSK}_i, \text{GID}, \mathbf{x}) \\ \text{CT} \leftarrow \text{Enc}(\{\text{PK}_i\}_i, C, \mu) \end{array} \right] = 1.$$

Security. For security, we consider standard semantic security but in presence of corrupt authorities. Here we consider a much stronger adversary which can even choose the corrupt authorities' keys on its own.

Definition 7.1 (Distributed ABE security in presence of corrupt authorities with unknown corrupt keys). A distributed ABE scheme is secure with corrupt authorities if for every stateful admissible PPT adversary \mathcal{A} , there exists a negligible function $\text{negl}(\cdot)$ such that for all $\lambda \in \mathbb{N}$, the following holds

$$\Pr \left[\begin{array}{l} \mathcal{A}^{O(\text{key}, \cdot, \cdot, \cdot)}(\text{CT}) = \mu^* : \\ \text{PP} \leftarrow \text{GSetup}(1^\lambda) \\ (1^n, 1^\ell, S^*) \leftarrow \mathcal{A}(1^\lambda, \text{PP}) \\ \forall i \notin S^* : (\text{PK}_i, \text{MSK}_i) \leftarrow \text{LSetup}(\text{PP}, 1^n, 1^\ell) \\ \text{key} = \{\text{MSK}_i\}_{i \notin S^*} \\ (C^*, \{\text{PK}_i\}_{i \in S^*}) \leftarrow \mathcal{A}^{O(\text{key}, \cdot, \cdot, \cdot)}(\text{PP}, \{\text{PK}_i\}_{i \notin S^*}) \\ \mu^* \leftarrow \{0, 1\}, \text{CT} \leftarrow \text{Enc}(\{\text{PK}_i\}_{i \in [n]}, C^*, \mu^*) \end{array} \right] \\ \leq \frac{1}{2} + \text{negl}(\lambda),$$

where the oracle $O(\text{key}, \cdot, \cdot, \cdot)$ has key material hardwired, and on input a tuple of a global identifier GID , an authority index j , and an attribute \mathbf{x} , and responds with a partial secret key computed as $\text{SK}_{j, \text{GID}, \mathbf{x}} \leftarrow \text{KeyGen}(\text{MSK}_j, \text{GID}, \mathbf{x})$. Note that the adversary is only allowed to submit key queries for non-corrupt authorities (i.e., $j \notin S^*$). Also, the adversary \mathcal{A} is admissible as long as every secret key query made by \mathcal{A} to the key generation oracle O satisfies the condition that — (1) either $C^*(\mathbf{x}) = 0$, or (2) adversary does not query at least one non-corrupt authority for the same global identifier, attribute tuple (GID, \mathbf{x}) .

Remark 7.2 (Comparing with decentralized ABE). The notion of distributed ABE differs from the notion of decentralized ABE in the sense that in a distributed ABE scheme each authority controls essentially an almost identical secret share of the master key, whereas in a decentralized (or multi-authority) ABE scheme each authority only controls the master key component for the attributes under its control.

7.2 Preliminaries

In this section, we define some preliminaries that we require for our constructions. Some other relevant preliminaries are provided in Section 2.

7.2.1 Attribute Based Encryption

Let $R = \{R_\lambda : A_\lambda \times B_\lambda \rightarrow \{0, 1\}\}_\lambda$ be a relation where A_λ and B_λ denote “ciphertext attribute” and “key attribute” spaces. An attribute-based encryption (ABE) scheme for R is defined by the following PPT algorithms:

$\text{Setup}(1^\lambda) \rightarrow (\text{PK}, \text{MSK})$: The setup algorithm takes as input the unary representation of the security parameter λ and outputs a master public key PK and a master secret key MSK .

$\text{Enc}(\text{PK}, X, \mu) \rightarrow \text{CT}$: The encryption algorithm takes as input a master public key PK , a ciphertext attribute $X \in A_\lambda$, and a message bit μ . It outputs a ciphertext CT .

$\text{KeyGen}(\text{PK}, \text{MSK}, Y) \rightarrow \text{SK}_Y$: The key generation algorithm takes as input the master public key PK , the master secret key MSK , and a key attribute $Y \in B_\lambda$. It outputs a private key SK_Y .

$\text{Dec}(\text{PK}, \text{CT}, X, \text{SK}_Y, Y) \rightarrow \mu$ or \perp : We assume that the decryption algorithm is deterministic. The decryption algorithm takes as input the master public key PK , a ciphertext CT , ciphertext attribute $X \in A_\lambda$, a private key SK_Y , and private key attribute $Y \in B_\lambda$. It outputs the message μ or \perp which represents that the ciphertext is not in a valid form.

Definition 7.3 (Correctness). An ABE scheme for relation family R is correct if for all $\lambda \in \mathbb{N}$, $X \in A_\lambda$, $Y \in B_\lambda$ such that $R(X, Y) = 1$, and for all messages $\mu \in \text{msg}$,

$$\Pr \left[\begin{array}{l} (\text{PK}, \text{MSK}) \leftarrow \text{Setup}(1^\lambda), \\ \text{SK}_Y \leftarrow \text{KeyGen}(\text{PK}, \text{MSK}, Y), \\ \text{CT} \leftarrow \text{Enc}(\text{PK}, X, \mu) : \\ \text{Decrypt}(\text{PK}, \text{SK}_Y, Y, \text{CT}, X) \neq \mu \end{array} \right] = \text{negl}(\lambda)$$

where the probability is taken over the coins of Setup , KeyGen , and Enc .

Definition 7.4 (Ada-IND security for ABE). For an ABE scheme $\text{ABE} = \{\text{Setup}, \text{Enc}, \text{KeyGen}, \text{Decrypt}\}$ for a relation family $R = \{R_\lambda : A_\lambda \times B_\lambda \rightarrow \{0, 1\}\}_\lambda$ and a message space $\{\text{msg}_\lambda\}_{\lambda \in \mathbb{N}}$ and an adversary A , let us define Ada-IND security game as follows.

1. **Setup phase:** On input 1^λ , the challenger samples $(\text{PK}, \text{MSK}) \leftarrow \text{Setup}(1^\lambda)$ and gives PK to A .
2. **Query phase:** During the game, A adaptively makes the following queries, in an arbitrary order. A can make unbounded many key queries, but can make only single challenge query.
 - (a) **Key Queries:** A chooses an input $Y \in B_\lambda$. For each such query, the challenger replies with $\text{SK}_Y \leftarrow \text{KeyGen}(\text{PK}, \text{MSK}, Y)$.
 - (b) **Challenge Query:** At some point, A submits a pair of equal length messages $(\mu_0, \mu_1) \in (\text{msg})^2$ and the target $X^* \in A_\lambda$ to the challenger. The challenger samples a random bit $b \leftarrow \{0, 1\}$ and replies to A with $\text{CT} \leftarrow \text{Enc}(\text{PK}, X^*, \mu_b)$.

We require that $R(X^*, Y) = 0$ holds for any Y such that A makes a key query for Y in order to avoid trivial attacks.

3. **Output phase:** A outputs a guess bit b' as the output of the experiment.

We define the advantage $\text{Adv}_{\text{ABE}, A}^{\text{Ada-IND}}(1^\lambda)$ of A in the above game as

$$\text{Adv}_{\text{ABE}, A}^{\text{Ada-IND}}(1^\lambda) := |\Pr[\text{Exp}_{\text{ABE}, A}(1^\lambda) = 1 | b = 0] - \Pr[\text{Exp}_{\text{ABE}, A}(1^\lambda) = 1 | b = 1]|.$$

The ABE scheme ABE is said to satisfy Ada-IND security (or simply *adaptive security*) if for any stateful PPT adversary A , there exists a negligible function $\text{negl}(\cdot)$ such that $\text{Adv}_{\text{ABE}, A}^{\text{Ada-IND}}(1^\lambda) \neq \text{negl}(\lambda)$.

We can consider the following stronger version of the security where we require the ciphertext to be pseudorandom.

Definition 7.5 (Ada-INDr security for ABE). We define Ada-INDr security game similarly to Ada-IND security game except that the adversary A chooses single message μ instead of (μ_0, μ_1) at the challenge phase and the challenger returns $\text{CT} \leftarrow \text{Enc}(\text{PK}, X^*, \mu)$ if $b = 0$ and a random ciphertext $\text{CT} \leftarrow \mathcal{CT}$ from a ciphertext space \mathcal{CT} if $b = 1$. We define the advantage $\text{Adv}_{\text{ABE}, A}^{\text{Ada-INDr}}(1^\lambda)$ of the adversary A accordingly and say that the scheme satisfies Ada-INDr security if the quantity is negligible.

We also consider (weaker) selective versions of the above notions, where A specifies its target X^* at the beginning of the game.

Definition 7.6 (Sel-IND security for ABE). We define Sel-IND security game as Ada-IND security game with the exception that the adversary A has to choose the challenge ciphertext attribute X^* before the setup phase but key queries Y_1, Y_2, \dots and choice of (μ_0, μ_1) can still be adaptive. We define the advantage $\text{Adv}_{\text{ABE}, A}^{\text{Sel-IND}}(1^\lambda)$ of the adversary A accordingly and say that the scheme satisfies Sel-INDr security (or simply *selective security*) if the quantity is negligible.

Definition 7.7 (Sel-INDr security for ABE). We define Sel-INDr security game as Ada-INDr security game with the exception that the adversary A has to choose the challenge ciphertext attribute X^* before the setup phase but key queries Y_1, Y_2, \dots and choice of μ can still be adaptive. We define the advantage $\text{Adv}_{\text{ABE}, A}^{\text{Sel-INDr}}(1^\lambda)$ of the adversary A accordingly and say that the scheme satisfies Sel-INDr security if the quantity is negligible.

In the following, we recall definitions of various ABEs by specifying the relation. We start with the standard notions of ciphertext-policy attribute-based encryption (CP-ABE) and key-policy attribute-based encryption (KP-ABE).

CP-ABE for circuits. We define CP-ABE for circuit class $\{\mathcal{C}_\lambda\}_\lambda$ by specifying the relation. Here, \mathcal{C}_λ is a set of circuits with input length $\ell(\lambda)$ and binary output. We define $A_\lambda^{\text{CP}} = \mathcal{C}_\lambda$ and $B_\lambda^{\text{CP}} = \{0, 1\}^\ell$. Furthermore, we define the relation R_λ^{CP} as

$$R_\lambda^{\text{CP}}(C, \mathbf{x}) = \neg C(\mathbf{x}).^{14}$$

KP-ABE for circuits. To define KP-ABE for circuits, we simply swap key and ciphertext attributes in CP-ABE for circuits. More formally, to define KP-ABE for circuits, we define $A_\lambda^{\text{KP}} = \{0, 1\}^\ell$ and $B_\lambda^{\text{KP}} = \mathcal{C}_\lambda$. We also define $R_\lambda^{\text{KP}} : A_\lambda^{\text{KP}} \times B_\lambda^{\text{KP}} \rightarrow \{0, 1\}$ as

$$R_\lambda^{\text{KP}}(\mathbf{x}, C) = \neg C(\mathbf{x}).$$

7.2.2 Lattice Preliminaries

Here, we recall some facts on lattices that are needed for the exposition of our construction. Throughout this section, n, m , and q are integers such that $n = \text{poly}(\lambda)$ and $m \geq n \lceil \log q \rceil$. In the following, let $\text{SampZ}(\gamma)$ be a sampling algorithm for the truncated discrete Gaussian distribution over \mathbf{Z} with parameter $\gamma > 0$ whose support is restricted to $z \in \mathbf{Z}$ such that $|z| \leq \sqrt{n}\gamma$.

Learning with Errors. We introduce then learning with errors (LWE) problem.

Definition 7.8 (The LWE Assumption). Let $n = n(\lambda)$, $m = m(\lambda)$, and $q = q(\lambda) > 2$ be integers and $\chi = \chi(\lambda)$ be a distribution over \mathbf{Z}_q . We say that the $\text{LWE}(n, m, q, \chi)$ hardness assumption holds if for any PPT adversary A we have

$$|\Pr[A(\mathbf{A}, \mathbf{s}^\top \mathbf{A} + \mathbf{x}^\top) \rightarrow 1] - \Pr[A(\mathbf{A}, \mathbf{v}^\top) \rightarrow 1]| \leq \text{negl}(\lambda)$$

where the probability is taken over the choice of the random coins by the adversary A and $\mathbf{A} \leftarrow \mathbf{Z}_q^{n \times m}$, $\mathbf{s} \leftarrow \mathbf{Z}_q^n$, $\mathbf{x} \leftarrow \chi^m$, and $\mathbf{v} \leftarrow \mathbf{Z}_q^m$. We also say that $\text{LWE}(n, m, q, \chi)$ problem is subexponentially hard if the above probability is bounded by $2^{-n^\epsilon} \cdot \text{negl}(\lambda)$ for some constant $0 < \epsilon < 1$ for all PPT A .

As shown by previous works [Reg09, BLP⁺13], if we set $\chi = \text{SampZ}(\gamma)$, the $\text{LWE}(n, m, q, \chi)$ problem is as hard as solving worst case lattice problems such as gapSVP and SIVP with approximation factor $\text{poly}(n) \cdot (q/\gamma)$ for some $\text{poly}(n)$. Since the best known algorithms for 2^k -approximation of gapSVP and SIVP run in time $2^{\tilde{O}(n/k)}$, it follows that the above $\text{LWE}(n, m, q, \chi)$ with noise-to-modulus ratio 2^{-n^ϵ} is likely to be (subexponentially) hard for some constant ϵ .

Trapdoors. Let us consider a matrix $\mathbf{A} \in \mathbf{Z}_q^{n \times m}$. For all $\mathbf{V} \in \mathbf{Z}_q^{n \times m'}$, we let $\mathbf{A}_\gamma^{-1}(\mathbf{V})$ be an output distribution of $\text{SampZ}(\gamma)^{m \times m'}$ conditioned on $\mathbf{A} \cdot \mathbf{A}_\gamma^{-1}(\mathbf{V}) = \mathbf{V}$. A γ -trapdoor for \mathbf{A} is a trapdoor that enables one to sample from the distribution $\mathbf{A}_\gamma^{-1}(\mathbf{V})$ in time $\text{poly}(n, m, m', \log q)$ for any \mathbf{V} . We slightly overload notation and denote a γ -trapdoor for \mathbf{A} by \mathbf{A}_γ^{-1} . We also define the special gadget matrix $\mathbf{G} \in \mathbf{Z}_q^{n \times m}$ as the matrix obtained by padding $\mathbf{I}_n \otimes (1, 2, 4, 8, \dots, 2^{\lceil \log q \rceil})$ with zero-columns. The following properties had been established in a long sequence of works [GPV08, CHKP10, ABB10a, ABB10b, MP12, BLP⁺13].

Lemma 7.9 (Properties of Trapdoors). Lattice trapdoors exhibit the following properties.

¹⁴Here, we follow the standard convention in lattice-based cryptography where the decryption succeeds when $C(\mathbf{x}) = 0$ rather than $C(\mathbf{x}) = 1$.

1. Given \mathbf{A}_τ^{-1} , one can obtain $\mathbf{A}_{\tau'}^{-1}$ for any $\tau' \geq \tau$.
2. Given \mathbf{A}_τ^{-1} , one can obtain $[\mathbf{A}\|\mathbf{B}]_\tau^{-1}$ and $[\mathbf{B}\|\mathbf{A}]_\tau^{-1}$ for any \mathbf{B} .
3. There exists an efficient procedure $\text{TrapGen}(1^n, 1^m, q)$ that outputs $(\mathbf{A}, \mathbf{A}_{\tau_0}^{-1})$ where $\mathbf{A} \in \mathbf{Z}_q^{n \times m}$ for some $m = O(n \log q)$ and is 2^{-n} -close to uniform, where $\tau_0 = \omega(\sqrt{n \log q \log m})$.

Lattice Evaluation. The following is an abstraction of the evaluation procedure in previous LWE based FHE and ABE schemes.

Lemma 7.10 (Fully Homomorphic Computation [GV15]). There exists a pair of deterministic algorithms $(\text{EvalF}, \text{EvalFX})$ with the following properties.

- $\text{EvalF}(\mathbf{B}, F) \rightarrow \mathbf{H}_F$. Here, $\mathbf{B} \in \mathbf{Z}_q^{n \times m\ell}$ and $F : \{0, 1\}^\ell \rightarrow \{0, 1\}$ is a circuit.
- $\text{EvalFX}(F, \mathbf{x}, \mathbf{B}) \rightarrow \widehat{\mathbf{H}}_{F, \mathbf{x}}$. Here, $\mathbf{x} \in \{0, 1\}^\ell$ and $F : \{0, 1\}^\ell \rightarrow \{0, 1\}$ is a circuit with depth d . We have

$$[\mathbf{B} - \mathbf{x} \otimes \mathbf{G}] \widehat{\mathbf{H}}_{F, \mathbf{x}} = \mathbf{B}\mathbf{H}_F - F(\mathbf{x})\mathbf{G} \pmod{q},$$

where we denote $[x_1 \mathbf{G} \parallel \dots \parallel x_k \mathbf{G}]$ by $\mathbf{x} \otimes \mathbf{G}$. Furthermore, we have

$$\|\mathbf{H}_F\|_\infty \leq m \cdot 2^{O(d)}, \quad \|\widehat{\mathbf{H}}_{F, \mathbf{x}}\|_\infty \leq m \cdot 2^{O(d)}.$$

- The running time of $(\text{EvalF}, \text{EvalFX})$ is bounded by $\text{poly}(n, m, \log q, 2^d)$.

The above algorithms are taken from [GV15], which is a variant of similar algorithms proposed by Boneh et al. [BGG⁺14]. The algorithms in [BGG⁺14] work for any polynomial-sized circuit F , but $\|\mathbf{H}_F\|_\infty$ and $\|\widehat{\mathbf{H}}_{F, \mathbf{x}}\|_\infty$ become super-polynomial even if the depth of the circuit is shallow (i.e., logarithmic depth). On the other hand, the above algorithms run in polynomial time only when F is of logarithmic depth, but $\|\mathbf{H}_F\|_\infty$ and $\|\widehat{\mathbf{H}}_{F, \mathbf{x}}\|_\infty$ can be polynomially bounded. The latter property is crucial for our purpose.

7.2.3 KP-ABE Scheme by Boneh et al. [BGG⁺14].

We will use a variant of the KP-ABE scheme proposed by Boneh et al. [BGG⁺14] as a building block of our construction of CP-ABE. We call the scheme BGG^+ and provide the description of the scheme in the following. We focus on the case where the policies associated with secret keys are limited to circuits with logarithmic depth rather than arbitrary polynomially bounded depth, so that we can use the evaluation algorithm due to Gorbunov and Vinayagamurthy [GV15] (see Lemma 7.10). This allows us to bound the noise growth during the decryption by a polynomial factor, which is crucial for our application.

The scheme supports the circuit class $\mathcal{C}_{\ell(\lambda), d(\lambda)}$, which is a set of all circuits with input length $\ell(\lambda)$ and depth at most $d(\lambda)$ with arbitrary $\ell(\lambda) = \text{poly}(\lambda)$ and $d(\lambda) = O(\log \lambda)$.

Setup(1^λ): On input 1^λ , the setup algorithm defines the parameters $n = n(\lambda)$, $m = m(\lambda)$, noise distribution noise over \mathbf{Z} , τ_0 , τ , and $B = B(\lambda)$ as specified later. It then proceeds as follows.

1. Sample $(\mathbf{A}, \mathbf{A}_{\tau_0}^{-1}) \leftarrow \text{TrapGen}(1^n, 1^m, q)$ such that $\mathbf{A} \in \mathbf{Z}_q^{n \times m}$.
2. Sample random matrix $\mathbf{B} = (\mathbf{B}_1, \dots, \mathbf{B}_\ell) \leftarrow (\mathbf{Z}_q^{n \times m})^\ell$ and a random vector $\mathbf{u} \leftarrow \mathbf{Z}_q^n$.
3. Output the master public key $\text{PK} = (\mathbf{A}, \mathbf{B}, \mathbf{u})$ and the master secret key $\text{MSK} = \mathbf{A}_{\tau_0}^{-1}$.

KeyGen(PK, MSK, F): The key generation algorithm takes as input the master public key PK , the master secret key MSK , and a circuit $F \in \mathcal{F}_\lambda$ and proceeds as follows.

1. Compute $\mathbf{H}_F = \text{EvalF}(\mathbf{B}, F)$ and $\mathbf{B}_F = \mathbf{B}\mathbf{H}_F$.
2. Compute $[\mathbf{A}\|\mathbf{B}_F]_\tau^{-1}$ from $\mathbf{A}_{\tau_0}^{-1}$ and sample $\mathbf{r} \in \mathbf{Z}^{2m}$ as $\mathbf{r} \leftarrow [\mathbf{A}\|\mathbf{B}_F]_\tau^{-1}(\mathbf{u})$.

3. Output the secret key $\text{SK}_F := \mathbf{r}$.

$\text{Enc}(\text{PK}, \mathbf{x}, \mu)$: The encryption algorithm takes as input the master public key PK , an attribute $\mathbf{x} \in \{0, 1\}^\ell$, and a message $\mu \in \{0, 1\}$ and proceeds as follows.

1. Sample $\mathbf{s} \leftarrow \mathbf{Z}_q^n$, $e_1 \leftarrow \text{noise}$, $\mathbf{e}_2 \leftarrow \text{noise}^m$, and $\mathbf{S}_{i,b} \leftarrow \{-1, 1\}^{m \times m}$ for $i \in [\ell]$ and $b \in \{0, 1\}$. Then, set $\mathbf{e}_{i,b} := \mathbf{S}_{i,b}^\top \mathbf{e}_2$ for $i \in [\ell]$ and $b \in \{0, 1\}$.
2. Compute

$$\begin{aligned} \psi_1 &:= \mathbf{s}^\top \mathbf{u} + e_1 + \mu \lceil q/2 \rceil \in \mathbf{Z}_q, \quad \psi_2^\top := \mathbf{s}^\top \mathbf{A} + \mathbf{e}_2^\top \in \mathbf{Z}_q^m, \\ \psi_{i,b}^\top &:= \mathbf{s}^\top (\mathbf{B} - x_i \mathbf{G}) + \mathbf{e}_{i,b}^\top \in \mathbf{Z}_q^m \text{ for all } i \in [\ell], \quad b \in \{0, 1\}. \end{aligned}$$

3. Output the ciphertext $\text{CT}_{\mathbf{x}} := (\psi_1, \psi_2, \{\psi_{i,x_i}\}_{i \in [\ell]})$, where x_i is the i -th bit of \mathbf{x} .

$\text{Decrypt}(\text{PK}, \text{SK}_{\mathbf{x}}, \mathbf{x}, F, \text{CT}_F)$: The decryption algorithm takes as input the master public key PK , a secret key SK_F for a circuit F , and a ciphertext $\text{CT}_{\mathbf{x}}$ for an attribute \mathbf{x} and proceeds as follows.

1. Parse $\text{CT}_{\mathbf{x}} \rightarrow (\psi_1 \in \mathbf{Z}_q, \psi_2 \in \mathbf{Z}_q^m, \{\psi_{i,x_i} \in \mathbf{Z}_q^m\}_{i \in [\ell]})$, and $\text{SK}_F \in \mathbf{Z}^{2m}$. If any of the component is not in the corresponding domain or $F(\mathbf{x}) = 1$, output \perp .
2. Concatenate $\{\psi_{i,x_i}\}_{i \in [\ell]}$ to form $\psi_3^\top = (\psi_{1,x_1}^\top, \dots, \psi_{\ell,x_\ell}^\top)$.
3. Compute

$$\psi' := \psi_1 - [\psi_2^\top \parallel \psi_3^\top] \mathbf{r}.$$

4. Output 0 if $\psi' \in [-B, B]$ and 1 if $[-B + \lceil q/2 \rceil, B + \lceil q/2 \rceil]$.

Remark 7.11. We note that the encryption algorithm above computes redundant components $\{\psi_{i,\neg x_i}\}_{i \in [\ell]}$ in the second step, which are discarded in the third step. However, due to this redundancy, the scheme has the following special structure that will be useful for us. Namely, the first and the second steps of the encryption algorithm can be executed without knowing \mathbf{x} . Only the third step of the encryption algorithm needs the information of \mathbf{x} , where it chooses $\{\psi_{i,x_i}\}_{i \in [\ell]}$ from $\{\psi_{i,b}\}_{i \in [\ell], b \in \{0,1\}}$ depending on each bit of \mathbf{x} and then output the former terms along with ψ_1 and ψ_2 .

There, the encryption algorithm, who takes as input a circuit C that specifies the policy and does not know the corresponding input \mathbf{x} , executes the first two steps of the above encryption algorithm. This is possible since these two steps do not need the knowledge of \mathbf{x} .

Parameters and Security. We choose the parameters for the scheme as follows:

$$\begin{aligned} m &= n^{1.1} \log q, & q &= 2^{\Theta(\lambda)}, & \chi &= \text{SampZ}(3\sqrt{n}), \\ \tau_0 &= n \log q \log m, & \tau &= m^{3.1} \ell \cdot 2^{O(d)}, & B &= n^2 m^2 \tau \cdot 2^{O(d)}. \end{aligned}$$

The parameter n will be chosen depending on whether we need **Sel-INDr** security or **Ada-INDr** security for the scheme. If it suffices to have **Sel-INDr** security, we set $n = \lambda^c$ for some constant $c > 1$. If we need **Ada-INDr** security, we have to enlarge the parameter to be $n = (\ell\lambda)^c$ in order to compensate for the security loss caused by the complexity leveraging.

We remark that if we were to use the above ABE scheme stand-alone, we would have been able to set q polynomially bounded as in [GV15]. The reason why we set q exponentially large is that we combine the scheme with bilinear maps of order q to lift the ciphertext components to the exponent so that they are “hidden” as in [AY20]. In order to use the security of the bilinear map, we set the group order q to be exponentially large.

The following theorem summarizes the security and efficiency properties of the construction. There are two parameter settings depending on whether we assume subexponential hardness of LWE or not.

Theorem 7.12 (Adapted from [GV15, BGG⁺14]). Assuming hardness of $\text{LWE}(n, m, q, \chi)$ with $\chi = \text{SampZ}(3\sqrt{n})$ and $q = O(2^{m^{1/\epsilon}})$ for some constant $\epsilon > 1$, the above scheme satisfies Sel-INDr security. Assuming *subexponential* hardness of $\text{LWE}(n, m, q, \chi)$ with the same parameters, the above scheme satisfies Ada-INDr security with respect to the ciphertext space $\mathcal{CT} := \mathbf{Z}_q^{m(\ell+1)+1}$

7.3 Construction

Here we provide our extension of [AY20] to distribute the setup and key generation process among non-interacting a-priori fixed number of authorities. Below we follow the [AY20] scheme syntactically and describe our modified construction. As in [AY20], we also work over asymmetric bilinear groups. Also, let $\mathbb{G}_1, \mathbb{G}_2$ and \mathbb{G}_T be the source and target groups, respectively. Our construction additionally relies on a hash function $H : \{0, 1\}^\lambda \times \{0, 1\}^\ell \rightarrow \mathbb{G}_2$ that is later modelled as a random oracle. Below we provide our ABE scheme following the notation used in [AY20, Section 3] almost verbatim for presentation purposes. For completeness, we recall the notation used later in Section 7.2.2 and Section 2.2.

Our construction can deal with any circuit class $\mathcal{F} = \{\mathcal{F}_\lambda\}_\lambda$ that is subclass of $\{\mathcal{C}_{\ell(\lambda), d(\lambda)}\}_\lambda$ with arbitrary $\ell(\lambda) \leq \text{poly}(\lambda)$ and $d(\lambda) = O(\log \lambda)$, where $\mathcal{C}_{\ell(\lambda), d(\lambda)}$ is a set of circuits with input length $\ell(\lambda)$ and depth at most $d(\lambda)$. Also, all algorithms are provided the public parameters as an additional input. We do not explicitly write for ease of exposition.

GSetup(1^λ): On input 1^λ , the setup algorithm samples a group description $\mathbf{G} = (q, \mathbb{G}_1, \mathbb{G}_2, \mathbb{G}_T, e, [1]_1, [1]_2)$, and sets the public parameters as $\text{PP} = \mathbf{G}$.

LSetup($\text{PP}, 1^N, 1^\ell$): On input the bilinear group description $\text{PP} = \mathbf{G}$, number of authorities N , attribute length ℓ , the authority setup algorithm defines the parameters $n = n(\lambda)$, $m = m(\lambda)$, noise distribution noise over \mathbf{Z} , τ_0, τ , and $B = B(\lambda)$ as specified in Section 7.2.3. It then sets $L := (2\ell + 1)m + 2$ and proceeds as follows.

1. Sample $\mathbf{w} \leftarrow (\mathbf{Z}_q^*)^L$ and compute $[\mathbf{w}]_1$.
2. Output $\text{PK} = [\mathbf{w}]_1$ and $\text{MSK} = \mathbf{w}$.

KeyGen($\text{MSK}_j, \text{GID}, \mathbf{x}$): The key generation algorithm takes as input the authority master secret key $\text{MSK}_j = \mathbf{w}^{(j)}$, global identifier GID , and an attribute $\mathbf{x} \in \{0, 1\}^\ell$ with $x_1 = 1$ and proceeds as follows.

1. Let $\mathbf{1} := (1, \dots, 1)^\top \in \mathbf{Z}_q^m$ and $\mathbf{0} := (0, \dots, 0)^\top \in \mathbf{Z}_q^m$. Set

$$\begin{aligned} \phi_0 &= \mathbf{1} \in \mathbf{Z}_q, & \phi_1 &= \mathbf{1} \in \mathbf{Z}_q, & \phi_2 &:= \mathbf{1} \in \mathbf{Z}_q^m, \\ \phi_{i,b} &:= \begin{cases} \mathbf{1} \in \mathbf{Z}_q^m & \text{if } b = x_i \\ \mathbf{0} \in \mathbf{Z}_q^m & \text{if } b \neq x_i \end{cases} & \text{for } i \in [\ell] \text{ and } b \in \{0, 1\}. \end{aligned} \quad (7.1)$$

2. Vectorize $(\phi_0, \phi_1, \phi_2, \{\phi_{i,b}\}_{i,b})$ to form a vector $\mathbf{d} \in \mathbf{Z}_q^L$ by concatenating each entry of the vectors in a predetermined order.
3. Compute $h = H(\text{GID}, \mathbf{x}) \in \mathbb{G}_2$.
4. Compute $h^{\mathbf{d} \otimes \mathbf{w}^{(j)}} \in \mathbb{G}_2^L$ using $h \in \mathbb{G}_2$ and $\mathbf{w}^{(j)}$ in the master key.
5. Output $\text{SK}_{j, \text{GID}, \mathbf{x}} = h^{\mathbf{d} \otimes \mathbf{w}^{(j)}}$.

Enc($\{\text{PK}_i\}_{i \in [N]}, F, \mu$): The encryption algorithm takes as input a list of authority master public keys PK_i for $i \in [N]$, the circuit F , and a message $\mu \in \{0, 1\}$ and proceeds as follows.

1. Sample fresh BGG⁺ scheme:
 - (a) Sample $(\mathbf{A}, \mathbf{A}_{\tau_0}^{-1}) \leftarrow \text{TrapGen}(1^n, 1^m, q)$ such that $\mathbf{A} \in \mathbf{Z}_q^{n \times m}$.

- (b) Sample random matrix $\mathbf{B} = (\mathbf{B}_1, \dots, \mathbf{B}_\ell) \leftarrow (\mathbf{Z}_q^{n \times m})^\ell$ and a random vector $\mathbf{u} \leftarrow \mathbf{Z}_q^n$.
2. Compute BGG^+ function key for circuit F :
- (a) Compute $\mathbf{H}_F = \text{EvalF}(\mathbf{B}, F)$ and $\mathbf{B}_F = \mathbf{B}\mathbf{H}_F$.
- (b) Compute $[\mathbf{A} \parallel \mathbf{B}_F]_\tau^{-1}$ from $\mathbf{A}_{\tau_0}^{-1}$ and sample $\mathbf{r} \in \mathbf{Z}^{2m}$ as $\mathbf{r} \leftarrow [\mathbf{A} \parallel \mathbf{B}_F]_\tau^{-1}(\mathbf{u})$.
3. Compute BGG^+ ciphertext for all possible inputs:
- (a) Sample $\mathbf{s} \leftarrow \mathbf{Z}_q^n$, $e_1 \leftarrow \text{noise}$, $\mathbf{e}_2 \leftarrow \text{noise}^m$, and $\mathbf{S}_{i,b} \leftarrow \{-1, 1\}^{m \times m}$ for $i \in [\ell]$ and $b \in \{0, 1\}$. Then, set $\mathbf{e}_{i,b} := \mathbf{S}_{i,b}^\top \mathbf{e}_2$ for $i \in [\ell]$ and $b \in \{0, 1\}$.
- (b) Compute

$$\begin{aligned} \psi_0 &:= 1 \in \mathbf{Z}_q, & \psi_1 &:= \mathbf{s}^\top \mathbf{u} + e_1 + \mu \lceil q/2 \rceil \in \mathbf{Z}_q, \\ \psi_2^\top &:= \mathbf{s}^\top \mathbf{A} + \mathbf{e}_2^\top \in \mathbf{Z}_q^m, \\ \psi_{i,b}^\top &:= \mathbf{s}^\top (\mathbf{B}_i - b\mathbf{G}) + \mathbf{e}_{i,b}^\top \in \mathbf{Z}_q^m & \text{for } i \in [\ell] \text{ and } b \in \{0, 1\}. \end{aligned} \quad (7.2)$$

4. N -out-of- N secret share the BGG^+ ciphertexts:
- (a) Vectorize $(\psi_0, \psi_1, \psi_2, \{\psi_{i,b}\}_{i,b})$ to form a vector $\mathbf{c} \in \mathbf{Z}_q^L$ by concatenating each entry of the vectors in a predetermined order (that aligns with the one used in the key generation algorithm).
- (b) Sample $N - 1$ uniformly random vectors $\tilde{\mathbf{c}}^{(j)} \leftarrow \mathbf{Z}_q^L$ for $j \in [N - 1]$.
- (c) Compute vector $\tilde{\mathbf{c}}^{(N)}$ as $\tilde{\mathbf{c}}^{(N)} = \mathbf{c} + \sum_{j \in [N-1]} \tilde{\mathbf{c}}^{(j)}$.
5. Encode the secret shared ciphertexts in exponent of bilinear group:
- (a) Sample $\gamma \leftarrow \mathbf{Z}_q^*$.
- (b) For every $j \in [N]$, compute $[\gamma \tilde{\mathbf{c}}^{(j)} \odot \mathbf{w}^{(j)}]_1 \in \mathbb{G}_1^L$ from γ , $\tilde{\mathbf{c}}^{(j)}$, and $[\mathbf{w}^{(j)}]_1$ in PK_j .
6. Output $\text{CT}_F = (\text{CT}_0 = (\mathbf{A}, \mathbf{B}), \{\text{CT}_1^{(j)} = [\gamma \tilde{\mathbf{c}}^{(j)} \odot \mathbf{w}^{(j)}]_1\}_{j \in [N]}, \text{CT}_2 = \mathbf{r})$.

$\text{Dec}(\{\text{SK}_{j, \text{GID}, \mathbf{x}}\}_{j \in [N]}, \mathbf{x}, F, \text{CT}_F)$: The decryption algorithm takes as input the N partial secret keys $\text{SK}_{j, \text{GID}, \mathbf{x}}$ for an attribute \mathbf{x} and authority index $j \in [N]$, and the ciphertext CT_F for a circuit F and proceeds as follows.

- Parse $\text{CT}_F \rightarrow (\text{CT}_0 = (\mathbf{A} \in \mathbf{Z}_q^{n \times m}, \mathbf{B} \in \mathbf{Z}_q^{n \times m \ell}), \{\text{CT}_1^{(j)} \in \mathbb{G}_1^L\}_j, \text{CT}_2 \in \mathbf{Z}^{2m})$ and $\text{SK}_{\mathbf{x}} \in \mathbb{G}_2^L$. If any of the component is not in the corresponding domain or $F(\mathbf{x}) = 1$, output \perp .
- Unmask and reconstruct the secret-shared BGG^+ ciphertexts corresponding to \mathbf{x} by using each partial secret key:
Compute the partial unmasked ciphertexts $[\mathbf{v}^{(j)}]_T := \text{CT}_1^{(j)} \odot \text{SK}_{j, \text{GID}, \mathbf{x}}$, and reconstruct them as $[\mathbf{v}]_T := \prod_{j \in [N]} [\mathbf{v}^{(j)}]_T$. Next, de-vectorize $[\mathbf{v}]_T$ to obtain

$$[v_0]_T \in \mathbf{G}_T, [v_1]_T \in \mathbf{G}_T, [v_2]_T \in \mathbf{G}_T^m, [v_{i,b}]_T \in \mathbf{G}_T^m, \text{ for } i \in [\ell], b \in \{0, 1\}.$$

- Evaluate circuit F on BGG^+ ciphertexts in the exponent:
Compute $\widehat{\mathbf{H}}_{F, \mathbf{x}} = \text{EvalF}(F, \mathbf{x}, \mathbf{B})$.
- Perform BGG^+ decryption in the exponent:
Form $[\mathbf{v}_{\mathbf{x}}^\top]_T = [\mathbf{v}_{1, x_1}^\top, \dots, \mathbf{v}_{\ell, x_\ell}^\top]_T$ and $\text{CT}_2^\top = (\mathbf{r}_1^\top \in \mathbf{Z}_q^m, \mathbf{r}_2^\top \in \mathbf{Z}_q^m)$. Then compute

$$[v']_T := [v_1 - (\mathbf{v}_2^\top \mathbf{r}_1 + \mathbf{v}_{\mathbf{x}}^\top \widehat{\mathbf{H}}_{F, \mathbf{x}} \mathbf{r}_2)]_T$$

from $[v_1]_T, [v_2]_T, [\mathbf{v}_{\mathbf{x}}]_T, \mathbf{r}_1, \mathbf{r}_2$, and $\widehat{\mathbf{H}}_{F, \mathbf{x}}$.

5. Recover exponent via brute force if $F(\mathbf{x}) = 0$:
Find $\eta \in [-B, B] \cup [-B + \lceil q/2 \rceil, B + \lceil q/2 \rceil]$ such that $[v_0]_T^\eta = [v']_T$ by brute-force search. If there is no such η , output \perp . To speed up the operation, one can employ the baby-step giant-step algorithm.
6. Output 0 if $\eta \in [-B, B]$ and 1 if $[-B + \lceil q/2 \rceil, B + \lceil q/2 \rceil]$.

Correctness. The correctness of our scheme follows almost directly from the correctness of the [AY20] CP-ABE scheme. The only difference being that in [AY20] the unmasking step directly reveals the BGG^+ ciphertexts in the exponent, whereas in the above construction the unmasking step reveals the secret shares of the BGG^+ ciphertext, which we first combine to recover the BGG^+ ciphertext and then decrypt as in [AY20]. More concretely, observe that for every $j \in [N]$,

$$\text{CT}_1^{(j)} \odot \text{SK}_{j, \text{GID}, \mathbf{x}} = h^{\gamma \cdot \tilde{\mathbf{c}}^{(j)} \odot \mathbf{d}_{\mathbf{x}}},$$

where $h = H(\text{GID}, \mathbf{x})$, and $\mathbf{d}_{\mathbf{x}}$ is as defined in key generation. Thus, we have that

$$\prod_{j \in [N]} \text{CT}_1^{(j)} \odot \text{SK}_{j, \text{GID}, \mathbf{x}} = h^{\gamma \cdot \sum_{j \in [N]} (\tilde{\mathbf{c}}^{(j)} \odot \mathbf{d}_{\mathbf{x}})} = h^{\gamma \cdot \mathbf{c} \odot \mathbf{d}_{\mathbf{x}}}.$$

Given this, rest of the correctness proof is identical to that provided in [AY20].

Proof of Security

Here we prove security of our construction. Formally, we prove the following.

Theorem 7.13. If the underlying key-policy ABE scheme BGG^+ satisfies semantic security with pseudorandom ciphertexts, then our distributed ciphertext-policy ABE scheme is semantically secure (Definition 7.1) in the generic group model.

The proof structure is similar to that of [AY20], except now the reduction algorithm needs to additionally answer key queries for GID -attribute vector pairs (GID, \mathbf{x}) for which $F^*(\mathbf{x}) = 1$ (where F^* is the challenge circuit) for some non-corrupt/honest authorities as well. Recall that the security must hold even if the attacker queries all but one honest authorities on such accepting input attribute vectors \mathbf{x} . Despite such additional key queries, it turns out the [AY20] proof could be extended to handle them. Intuitively, this is because even given such partial accepting keys one could show that the adversary is not able to make any “bad” zero-test queries as defined in [AY20]. Below we sketch the hybrid games along the lines of [AY20]. We refer the reader to [AY20, Section 4] for more details.

Game 0. This corresponds to the distributed ABE security game as described in Definition 7.1. Let Q_{kq} denote the total number of queries made by the adversary, Q_{zt} denote the number of zero-test queries, and Q_{ro} denote the number of RO queries. Here the challenger simulates both the random oracle as well as generic group oracle. (That is, it gives handles to the group elements each time, and when answering a RO query, it again answers with a handle since hash H maps elements to \mathbb{G}_2 .)

Game 1. This is same as the previous game, except the way in which the challenger answers/resolves each call to the hash function/random oracle:

- on input (GID, \mathbf{x}) , the challenger chooses a random exponent $\delta_{\text{GID}, \mathbf{x}} \leftarrow \mathbf{Z}_q$, stores $\delta_{\text{GID}, \mathbf{x}}$ for answering future queries on same input, and responds with $[\delta_{\text{GID}, \mathbf{x}}]_1$. (Intuitively, the challenger samples a random exponent per input to simulate the randomness term δ in the [AY20] secret keys.)

Game 2. This is same as the previous game, except the challenger samples the honest authority master keys $\mathbf{w}^{(j)}$ (for $j \notin S^*$, where S^* is the set of corrupt users), random exponents $\delta_{\text{GID},\mathbf{x}}$ (for all key queries (GID, \mathbf{x}), and RO queries too), ciphertext components $\mathbf{A}, \mathbf{B}, \mathbf{u}, \gamma, \mu, \mathbf{c}$, and $\tilde{\mathbf{c}}^{(j)}$ (for $j \in [N]$) at the beginning of the game. Since all these terms are independent of the challenge circuit F^* , thus the game is well defined. (This is similar to that done in Game 1 in [AY20, Section 4].)

Game 3. This is same as the previous game, except the challenger (partially) switches to the symbolic group model and replaces the terms $\{\mathbf{w}^{(j)}\}_{j \notin S^*}$, $\{\delta_{\text{GID},\mathbf{x}}\}_{\text{GID},\mathbf{x}}$, γ , \mathbf{c} and $\{\tilde{\mathbf{c}}^{(j)}\}_j$ with their respective formal variables $\{\mathbf{W}^{(j)}\}_{j \notin S^*}$, $\{\Delta_{\text{GID},\mathbf{x}}\}_{\text{GID},\mathbf{x}}$, Γ , \mathbf{C} and $\{\tilde{\mathbf{C}}^{(j)}\}_j$. Let $\mathbf{W}_j = (W_1^{(j)}, \dots, W_L^{(j)})$ and $\tilde{\mathbf{C}}^{(j)} = (\tilde{C}_1^{(j)}, \dots, \tilde{C}_L^{(j)})$. As a result, all handles given to adversary \mathcal{A} refer to elements in the ring

$$\mathbb{T} := \mathbf{Z}_q[\{W_i^{(j)}, 1/W_i^{(j)}\}_{j \notin S^*, i \in [L]}, \{\Delta_{\text{GID},\mathbf{x}}\}_{\text{GID},\mathbf{x}}, \Gamma, \{\tilde{C}_i^{(j)}\}_{j \in [N], i \in [L]}],$$

where $\{1/W_i^{(j)}\}_{j,i}$ are needed to represent the components in the secret keys. However, whenever it performs a zero-test then it checks if the underlying ring element evaluates to 0 by substituting formal variables with corresponding elements in \mathbf{Z}_q . (This is similar to that done in Game 2 in [AY20, Section 4]. However, the only difference is that since the adversary corrupts some of the key authorities, thus it knows/chooses the key vectors $\mathbf{w}^{(j)}$ on its own for $j \in S^*$. Recall that [AY20] this was not the case.)

Note. As in [AY20], we extend the definition of sets $S_{T,1}$ and $S_{T,2}$ to our above distributed ABE construction. Briefly, the difference being that now the number of products being unmatching positions increases by a large amount due to cross terms between ciphertext and key components that correspond to different authorities.

Game 4. This is same as the previous game, except the challenger treats $\{\mathbf{W}_j\}_{j \notin S^*}$, $\{\Delta_{\text{GID},\mathbf{x}}\}_{\text{GID},\mathbf{x}}$, and Γ as formal variables rather than elements in \mathbf{Z}_q even when answering zero-test queries. (This is similar to that done in Game 3 in [AY20, Section 4].)

Game 5. This is same as the previous game, except now the challenger aborts the game and enforces the adversary to output a random bit when there exists $i \in [L]$ and $j \notin S^*$ such that $\tilde{c}_i^{(j)} = 0$, where $\tilde{\mathbf{c}}^{(j)} = (\tilde{c}_1^{(j)}, \dots, \tilde{c}_L^{(j)})^\top$ is sampled as in the previous game. (Note that this departs slightly from Game 4 in [AY20, Section 4]. This is due to the fact that [AY20] is a single-authority ABE system, whereas ours is a distributed scheme.)

Game 6. This is same as the previous game, except now the challenger changes how it answers the zero-test queries slightly. In particular, when adversary makes a zero-test query, then the challenger interprets it as a ring element of the form $\sum_{Z \in S_{T,1}} a_Z Z + \sum_{Z \in S_{T,2}} a_Z Z$ (where sets $S_{T,1}$ and $S_{T,2}$ are defined as appropriate extensions of that in [AY20] as discussed after description of Game 3 above). And, if there exists a $Z \in S_{T,1}$ such that its coefficient $a_Z \neq 0$, then the challenger returns 0, otherwise it answers the query as before. (This is similar to that done in Game 5 in [AY20, Section 4].)

Game 7. This is same as the previous game, except now the challenger answers *all* the zero-test queries completely over the ring (defined over the formal variables). Namely, when adversary makes a zero-test query for a handle corresponding to a ring element $f \in \mathbb{T}$, then the challenger returns 0 if f is not already a zero element over the ring \mathbb{T} , that is $f \neq 0$ over \mathbb{T} . In other words, the challenger returns 0 if there exists a $Z \in S_{T,1} \cup S_{T,2}$ such that its coefficient $a_Z \neq 0$. Note that $\{\tilde{\mathbf{c}}^{(j)}\}_{j \notin S^*}$ is not used in this game, thus the challenger does not have to sample them any more. (This is similar to that done in Game 6 in [AY20, Section 4].)

Indistinguishability of games

We complete the proof by showing that adjacent games are indistinguishable. For any adversary \mathcal{A} and game X , we denote by $\text{Adv}_S^{\mathcal{A}}(\lambda)$, the probability that \mathcal{A} wins in game S .

Lemma 7.14. For any adversary \mathcal{A} , we have that $\text{Adv}_0^{\mathcal{A}}(\lambda) = \text{Adv}_1^{\mathcal{A}}(\lambda)$.

Proof. This follows directly from the fact that H is modelled as a programmable random oracle. \square

Lemma 7.15. For any adversary \mathcal{A} , we have that $\text{Adv}_1^{\mathcal{A}}(\lambda) = \text{Adv}_2^{\mathcal{A}}(\lambda)$.

Proof. Since this is only a conceptual change where the challenger pre-computes ciphertext terms, the lemma immediately follows. \square

Lemma 7.16. For any adversary \mathcal{A} , we have that $\text{Adv}_2^{\mathcal{A}}(\lambda) = \text{Adv}_3^{\mathcal{A}}(\lambda)$.

Proof. This lemma again follows from the fact that this game just constitutes a syntactic difference, where the only difference is in how queries are answered. But since the terms are sampled identically in both games, thus the distributions are identical. Therefore, the lemma follows. \square

Lemma 7.17. For any adversary \mathcal{A} , we have that $\text{Adv}_3^{\mathcal{A}}(\lambda) - \text{Adv}_4^{\mathcal{A}}(\lambda) \leq Q_{\text{zt}}(N \cdot L + 3)^2/q$.

Proof. The proof of this lemma is identical to that of [AY20, Lemma 4.4], and follows by an application of Schwartz-Zippel lemma and union bound. The main difference is that the number of authority random variables \mathbf{W} grew from L to $N \cdot L$, thus the degree of the polynomial on which Schwartz-Zippel is applied gets raised by a factor of $N - |S^*| \leq N$, where S^* is the set of corrupt authorities. \square

Lemma 7.18. For any adversary \mathcal{A} , we have that $\text{Adv}_4^{\mathcal{A}}(\lambda) - \text{Adv}_5^{\mathcal{A}}(\lambda) \leq N \cdot L/q$.

Proof. This is a statistical property, and follows from the combination of following two facts. First, the challenger samples all but the last encoding vector $\tilde{\mathbf{c}}^{(j)}$ uniformly at random. Thus, by a simple union bound we get that the probability $\tilde{c}_i^{(j)} \neq 0$ for all $i \in [L]$ and $j \in [N - 1]$ is at most $(N - 1)L/q$. Now, the vector $\tilde{\mathbf{c}}^{(N)}$ contains the BGG^+ ciphertext components which are either fixed to be 1 or well distributed over \mathbf{Z}_q (assuming LWE).¹⁵ Therefore, the lemma follows by combining these. \square

Lemma 7.19. For any adversary \mathcal{A} , we have that $\text{Adv}_5^{\mathcal{A}}(\lambda) = \text{Adv}_6^{\mathcal{A}}(\lambda)$.

Proof. The proof of this lemma is similar to that of [AY20, Lemma 4.6], and intuitively follows from the fact that all monomials in $S_{T,1}$ are distinct (even if one substitutes ciphertext formal variables $\{\tilde{C}_i^{(j)}\}_{j,i}$ with its actual value $\{\tilde{c}_i^{(j)}\}_{j,i}$ and ignore the difference between the coefficients of the monomials), thus if $a_Z \neq 0$ for some $Z \in S_{T,1}$, then that term will never get cancelled, thus not lead to a successful zero-test query. Hence, the lemma follows. \square

Lemma 7.20. If the underlying key-policy ABE scheme BGG^+ satisfies semantic security with pseudorandom ciphertexts and $\log q = \omega(\log \lambda)$, then for every PPT adversary \mathcal{A} , there exists a negligible function $\text{negl}(\cdot)$ such that $\text{Adv}_6^{\mathcal{A}}(\lambda) - \text{Adv}_7^{\mathcal{A}}(\lambda) \leq \text{negl}(\lambda)$.

Proof. In [AY20], the authors proposed a hybrid based approach for proving the single-authority variant of the above lemma. Here we provide a different case based analysis. Our proof structure could also be used in [AY20] to potentially provide a simpler of the same.

The proof of this lemma is based on the ideas used in the proof of [AY20, Lemma 4.7], but there are a few differences. Very briefly, the main difference is that here we could have more than one non-corrupt/honest authorities, thus an adversary could potentially make key queries to all but one honest authorities on GID -attribute pairs (GID, \mathbf{x}) where $F^*(\mathbf{x}) = 1$. (Note that an adversary is allowed to query some of the

¹⁵This part is similar to the proof of [AY20, Lemma 4.5].

honest authorities on accepting attributes, as long as it does not query all the honest authorities thereby simply being able to decrypt the challenge ciphertext, since we want to disallow an adversary utilizing partial secret keys for two different accepting attributes.)

First, observe that Games 6 and 7 differ only when \mathcal{A} makes a zero-test query for a handle corresponding to $f \in \mathbb{T}$ that can be represented as

$$f(\{W_i^{(j)}\}_{j \notin S^*, i \in [L]}, \{\Delta_{\text{GID}, \mathbf{x}}\}_{\text{GID}, \mathbf{x}}, \Gamma, \{\tilde{C}_i^{(j)}\}_{j \in [N], i \in [L]}) = \sum_{Z \in S_{T,2}} a_Z Z \quad (7.3)$$

and satisfies $f \neq 0$ over \mathbb{T} , but

$$f(\{W_i^{(j)}\}_{j \notin S^*, i \in [L]}, \{\Delta_{\text{GID}, \mathbf{x}}\}_{\text{GID}, \mathbf{x}}, \Gamma, \{\tilde{C}_i^{(j)}\}_{j \in [N], i \in [L]}) = 0. \quad (7.4)$$

That is, only substituting the formal variables corresponding to $\tilde{C}_i^{(j)}$ to their actual values makes the ring element go to 0. Following [AY20], we call such a query *bad*. Now note that the set $S_{T,2}$ contains the following terms:

$$\begin{aligned} S_{T,2} = & \left\{ \Gamma \tilde{C}_i^{(j)} \Delta_{\text{GID}, \mathbf{x}} : \begin{array}{l} (j, \text{GID}, \mathbf{x}) \in \mathcal{Q}_{\text{key}}, i \in [L], d_{\mathbf{x}, i} = 1 \\ \text{where } \mathbf{d}_{\mathbf{x}} = (d_{\mathbf{x}, 1}, \dots, d_{\mathbf{x}, L}) \text{ and} \\ \mathcal{Q}_{\text{key}} \text{ denotes the set of key queries by } \mathcal{A} \end{array} \right\} \\ \cup & \left\{ \Gamma \tilde{C}_i^{(j)} \Delta_{\text{GID}, \mathbf{x}} : \begin{array}{l} j \in S^*, i \in [L], (\text{GID}, \mathbf{x}) \in \mathcal{Q}_{\text{ro}}, \\ \text{where } \mathcal{Q}_{\text{ro}} \text{ denotes the set of RO queries by } \mathcal{A} \end{array} \right\} \end{aligned}$$

Suppose \mathcal{A} always queries RO on (GID, \mathbf{x}) whenever it makes a key query to some honest authority on (GID, \mathbf{x}) . This can be assumed without loss of generality. Now one could additionally split the above set $S_{T,2}$ as per the GID-attribute pairs (GID, \mathbf{x}) as:

$$\begin{aligned} S_{T,2}^{(\text{GID}, \mathbf{x})} = & \left\{ \Gamma \tilde{C}_i^{(j)} \Delta_{\text{GID}, \mathbf{x}} : \begin{array}{l} i \in [L], d_{\mathbf{x}, i} = 1, \forall j \text{ s.t. } (j, \text{GID}, \mathbf{x}) \in \mathcal{Q}_{\text{key}} \\ \text{where } \mathbf{d}_{\mathbf{x}} = (d_{\mathbf{x}, 1}, \dots, d_{\mathbf{x}, L}) \text{ and} \\ \mathcal{Q}_{\text{key}} \text{ denotes the set of key queries by } \mathcal{A} \end{array} \right\} \\ \cup & \left\{ \Gamma \tilde{C}_i^{(j)} \Delta_{\text{GID}, \mathbf{x}} : j \in S^*, i \in [L] \right\} \end{aligned}$$

First, note that any *bad* query $f \in \mathbb{T}$ can now be represented as

$$\begin{aligned} f(\{W_i^{(j)}\}_{j \notin S^*, i \in [L]}, \{\Delta_{\text{GID}, \mathbf{x}}\}_{\text{GID}, \mathbf{x}}, \Gamma, \{\tilde{C}_i^{(j)}\}_{j \in [N], i \in [L]}) \\ = \sum_{(\text{GID}, \mathbf{x}) \in \mathcal{Q}_{\text{ro}}} \sum_{Z \in S_{T,2}^{(\text{GID}, \mathbf{x})}} a_Z Z. \end{aligned}$$

Now we further classify *bad* queries into “types”. We say that f is a Type- (GID, \mathbf{x}) bad query if:

$$\sum_{Z \in S_{T,2}^{(\text{GID}, \mathbf{x})}} a_Z Z \neq 0 \quad \text{and} \quad \sum_{Z \in S_{T,2}^{(\text{GID}, \mathbf{x})}} a_Z Z(\{\tilde{C}_i^{(j)}\}_{j \in [N], i \in [L]}) = 0,$$

where $Z(\{\tilde{C}_i^{(j)}\}_{j, i})$ denotes $Z(\{W_i^{(j)}\}_{j, i}, \{\Delta_{\text{GID}, \mathbf{x}}\}_{\text{GID}, \mathbf{x}}, \Gamma, \{\tilde{C}_i^{(j)}\}_{j, i}) \in \mathbb{T}$ above.

Next, we claim the following.

Claim 7.21. If $f \in \mathbb{T}$ is a “bad” query, then there must exist $(\text{GID}, \mathbf{x}) \in \mathcal{Q}_{\text{ro}}$ such that f is a “Type- (GID, \mathbf{x}) bad” query.

Proof. The proof of this claim follows along the lines of indistinguishability proof of hybrids 5.3 and 5.4 in [AY20]. The idea is that if f is a bad query, then there must exist a $(\text{GID}, \mathbf{x}) \in \mathcal{Q}_{\text{ro}}$ satisfying $\sum_{Z \in S_{T,2}^{(\text{GID}, \mathbf{x})}} a_Z Z \neq 0$. Furthermore, we have that

$$\sum_{Z \in S_{T,2}^{(\text{GID}, \mathbf{x})}} a_Z Z(\{\tilde{C}_i^{(j)}\}_{j \in [N], i \in [L]}) = - \sum_{(\text{GID}', \mathbf{x}') \neq (\text{GID}, \mathbf{x})} \sum_{Z \in S_{T,2}^{(\text{GID}', \mathbf{x}')}} a_Z Z(\{\tilde{C}_i^{(j)}\}_{j \in [N], i \in [L]}).$$

However, the above is impossible unless the left hand side equals to 0 since any monomial in $S_{T,2}^{(\text{GID}, \mathbf{x})}$ never appears in $S_{T,2}^{(\text{GID}', \mathbf{x}')}$ for $(\text{GID}', \mathbf{x}') \neq (\text{GID}, \mathbf{x})$ even if we replace $\{\tilde{C}_i^{(j)}\}$ with $\{\tilde{c}_i^{(j)}\}$ and ignore the difference between the coefficients of the monomials. This is due to the fact that there is a $\delta_{\text{GID}, \mathbf{x}}$ variable on the left side that never appears on the right side. Thus, f must be Type-(GID, \mathbf{x}) bad query. \square

To complete the proof of this lemma, we just need to show that adversary \mathcal{A} never makes any “Type-(GID, \mathbf{x}) bad” query. To argue this, we prove the following:

Claim 7.22. For every GID-attribute pair $(\text{GID}, \mathbf{x}) \in \mathcal{Q}_{\text{ro}}$, zero-test query number $t \in [Q_{\text{zt}}]$, if BGG^+ satisfies semantic security with pseudorandom ciphertexts and $\log q = \omega(\log \lambda)$, then the probability that adversary \mathcal{A} 's first bad query is the t -th zero-test query and it is a “Type-(GID, \mathbf{x}) bad” query is at most $\text{negl}'(\lambda)$ for some negligible function $\text{negl}'(\cdot)$.

Proof. Now the GID-attribute pairs (GID, \mathbf{x}) queried by \mathcal{A} can be further divided into two categories — (1) $F^*(\mathbf{x}) = 0$ (that is, attribute does not satisfy the predicate), or (2) there exists an index $j \notin S^*$ such that \mathcal{A} does not make a key generation query of the form $(j, \text{GID}, \mathbf{x})$ (that is, it does not query at least one honest authority for its key share for GID-attribute pair (GID, \mathbf{x})). (Refer to Definition 7.1 for admissible key queries.)

First, we argue that in case (2) as per above categorization, the claim is correct. This follows from the observation that if \mathcal{A} makes its first bad query on t -th zero-test query and it is a “Type-(GID, \mathbf{x}) bad” query, and \mathcal{A} did not make a key query on (GID, \mathbf{x}) to some honest authority say j^* , then in that case the every monomial $Z \in S_{T,2}^{(\text{GID}, \mathbf{x})}$ does not depend on $\tilde{\mathbf{c}}^{(j^*)}$. Now since $\tilde{\mathbf{c}}^{(j)}$ are sampled uniformly at random with the only constraint that $\sum_j \tilde{\mathbf{c}}^{(j)} = \mathbf{c}$, thus the joint distributions of $\tilde{\mathbf{c}}^{(j)}$ for $j \neq j^*$ (i.e., excluding $\tilde{\mathbf{c}}^{(j^*)}$) is identical to that of a random distribution over \mathbf{Z}_q^L . By using uniformity of $\tilde{\mathbf{c}}^{(j)}$ for $j \neq j^*$, we can conclude the argument that this event can happen only with probability at most $1/q$ since f is represented as a linear combination of $\{\Gamma \tilde{C}_i^{(j)} \Delta_{\text{GID}, \mathbf{x}}\}_{i,j \neq j^*}$ and all entries of $\{\tilde{c}_i^{(j)}\}_{i \in [L], j \neq j^*}$ are chosen uniformly at random. This completes the argument for case (2).

Otherwise, in case (1) as per above categorization, we have that $F^*(\mathbf{x}) = 0$. To argue negligible probability of this event, we define an intermediate hybrid game as follows. (This part of the proof is similar to that of [AY20, Lemma 4.7-4.8].) In the intermediate hybrid game, the challenger samples c_i as in the original scheme only for $i \in [L]$ such that $d_{\mathbf{x},i} = 1$, where $\mathbf{d}_{\mathbf{x}} = (d_{\mathbf{x},1}, \dots, d_{\mathbf{x},L})$ and $\mathbf{d}_{\mathbf{x}}$ is as defined in the construction for attribute \mathbf{x} . The game is still well-defined since the only place in the game where we need the information of \mathbf{c} is to check that t -th zero-test query is a “Type-(GID, \mathbf{x}) bad” query, and we only need c_i for $i \in [L]$ such that $d_{\mathbf{x},i} = 1$. Basically given these, the reduction can sample all the $\tilde{c}_i^{(j)}$ terms appropriately. (Here appropriately means that for $j \in S^*$, it samples all of them to be uniformly random, while for $j \notin S^*$ it only needs to simulate half of the $\tilde{c}_i^{(j)}$ terms depending on \mathbf{x} which are computable given half of the c_i terms.)

We start by claiming that this intermediate hybrid is identically distributed to the previous game. This follows from the fact that $S_{T,2}^{(\text{GID}, \mathbf{x})}$ only contains terms $\Gamma \tilde{C}_i^{(j)} \Delta_{\text{GID}, \mathbf{x}}$ for $j \notin S^*, i \in [L], d_{\mathbf{x},i} = 1$. Thus, since c_i is N -out-of- N secret shared into $\tilde{c}_i^{(j)}$, thus the challenger only needs to sample c_i for $i \in [L]$ such that $d_{\mathbf{x},i} = 1$ in both the original game as well as the intermediate hybrid game. This proves indistinguishability of the intermediate hybrid game with the original game.

Next, once the challenger only needs to sample c_i for $i \in [L]$ such that $d_{\mathbf{x},i} = 1$ and given that $F^*(\mathbf{x}) = 0$, thus we could switch all these c_i terms to be sampled uniformly at random instead by relying semantic security (with pseudorandom ciphertexts) of BGG^+ . This reduction is identical to that of [AY20, Lemma 4.8]. Finally, once c_i for $i \in [L]$ such that $d_{\mathbf{x},i} = 1$ are sampled uniformly at random, then as in the proof of case (2) previously, we could show that a bad query happens with probability at most $1/q$. This completes the argument in case (1) too, and hence this completes the proof. \square

Finally, using the above claim with a simple union bound over the total number of zero-test queries and key queries, we get that Games 6 and 7 are negligibly far apart. Since, the adversary's advantage is Game 7 is clearly zero as it contains no information about the challenge bit, thus the theorem follows.

□

8 Decentralized ABE \circ IPFE with Policy Hiding

In this section, we extend the decentralized inner-product predicate encryption by Michalevsky and Joye [MJ18] to provide an inner-product functional encryption capability. We observe that the primitive of decentralized attribute based, inner product FE is a natural and meaningful composition of decentralized attribute based encryption (DABE) and inner product functional encryption (IPFE), which is suggested by casting both existing primitives as special cases of our model.

For the construction, our high-level idea is to exploit the underlying algebraic structure and use the linear ciphertext and key homomorphism properties already satisfied by their construction, reminiscing the inner-product functional encryption schemes from (linearly homomorphic) PKE in [ABDCP15, ALS16, BJK15].

8.1 Definition

In this section, we define the notion of decentralized attribute-based inner-product function encryption. First, we recall the syntax, and later describe the security definition.

Syntax. A decentralized attribute-based inner-product function encryption for predicate class $\mathcal{C} = \{\mathcal{C}_n : \mathcal{X}_n \rightarrow \{0, 1\}\}_{n \in \mathbb{N}}$ and inner product message space $\mathcal{U} = \{\mathcal{U}_\ell\}_{\ell \in \mathbb{N}}$ consists of the following PPT algorithms:

$\text{GSetup}(1^\lambda) \rightarrow \text{PP}$. On input the security parameter λ , the setup algorithm outputs public parameters PP.

$\text{LSetup}(\text{PP}, 1^n, 1^\ell, i) \rightarrow (\text{PK}, \text{MSK})$. On input the public parameters PP, attribute length n , message space index ℓ , and authority's index $i \in [n]$, the authority setup algorithm outputs a pair of master public-secret key (PK, MSK).

$\text{KeyGen}(\{\text{PK}_i\}_{i \in [n]}, \text{MSK}_j, \text{GID}, \mathbf{x}, \mathbf{u}) \rightarrow \text{SK}_{j, \text{GID}, \mathbf{x}, \mathbf{u}}$. The key generation algorithm takes as input the public keys of all the authorities $\{\text{PK}_i\}_i$, an authority master secret key MSK_j , global identifier GID, an attribute $\mathbf{x} \in \mathcal{X}_n$, and key vector $\mathbf{u} \in \mathcal{U}_\ell$. It outputs a partial secret key $\text{SK}_{j, \text{GID}, \mathbf{x}, \mathbf{u}}$.

$\text{Enc}(\{\text{PK}_i\}_{i \in [n]}, C, \mathbf{v}) \rightarrow \text{CT}$. The encryption algorithm takes as input the list of public keys $\{\text{PK}_i\}_i$, predicate circuit C , and a message vector $\mathbf{v} \in \mathcal{U}_\ell$, and outputs a ciphertext CT.

$\text{Dec}(\{\text{SK}_{i, \text{GID}, \mathbf{x}, \mathbf{u}}\}_{i \in [n]}, \text{CT}) \rightarrow m/\perp$. On input a list of n partial secret keys $\{\text{SK}_{i, \text{GID}, \mathbf{x}, \mathbf{u}}\}_i$ and a ciphertext CT, the decryption algorithm either outputs a message m (corresponding to the inner product value) or a special string \perp (to denote decryption failure).

We require such an ABE scheme to satisfy the following properties.

Correctness. A decentralized attribute-based inner-product function encryption (AB-IPFE) scheme is said to be correct if for all $\lambda, n, \ell \in \mathbb{N}$, $C \in \mathcal{C}_n$, $\mathbf{u}, \mathbf{v} \in \mathcal{U}_\ell$, $\mathbf{x} \in \mathcal{X}_n$, GID, if $C(\mathbf{x}) = 1$, the following holds:

$$\Pr \left[\text{Dec}(\text{SK}, \text{CT}) = \langle \mathbf{u}, \mathbf{v} \rangle : \begin{array}{l} \text{PP} \leftarrow \text{GSetup}(1^\lambda) \\ \forall i \in [n] : (\text{PK}_i, \text{MSK}_i) \leftarrow \text{LSetup}(\text{PP}, 1^n, 1^\ell, i) \\ \forall j \in [n] : \text{SK}_{j, \text{GID}, \mathbf{x}, \mathbf{u}} \leftarrow \text{KeyGen}(\{\text{PK}_i\}_i, \text{MSK}_j, \text{GID}, \mathbf{x}, \mathbf{u}) \\ \text{CT} \leftarrow \text{Enc}(\{\text{PK}_i\}_i, C, \mathbf{v}), \text{SK} = \{\text{SK}_{i, \text{GID}, \mathbf{x}, \mathbf{u}}\}_i \end{array} \right] = 1.$$

Security. For security, we consider a combination of semantic security for decentralized ABE systems in presence of corrupt authorities and message vector indistinguishability for inner product functional encryption. Below we provide the selective security variant of the corresponding property.¹⁶

Definition 8.1 (Selective decentralized AB-IPFE security with 1-sided policy hiding in presence of corrupt authorities). A decentralized AB-IPFE scheme is selectively secure with 1-sided policy hiding in presence of corrupt authorities if for every stateful admissible PPT adversary \mathcal{A} , there exists a negligible function $\text{negl}(\cdot)$ such that for all $\lambda \in \mathbb{N}$, the following holds

$$\Pr \left[\begin{array}{l} \mathcal{A}^{O(\text{key}, \cdot, \cdot)}(\text{params}, \text{CT}) = b : \\ \begin{array}{l} (1^n, 1^\ell, S^*, (C_0^*, C_1^*), (\mathbf{v}_0^*, \mathbf{v}_1^*)) \leftarrow \mathcal{A}(1^\lambda), \text{PP} \leftarrow \text{GSetup}(1^\lambda) \\ \forall i \in [n] : (\text{PK}_i, \text{MSK}_i) \leftarrow \text{LSetup}(\text{PP}, 1^n, 1^\ell, i) \\ b \leftarrow \{0, 1\}, \text{CT} \leftarrow \text{Enc}(\{\text{PK}_i\}_{i \in [n]}, C_b^*, \mathbf{v}_b^*) \\ \text{key} = \{(\text{PK}_i, \text{MSK}_i)\}_{i \in [n]} \\ \text{params} = (\{\text{PK}_i\}_{i \in [n]}, \{\text{MSK}_i\}_{i \in S^*}) \end{array} \end{array} \right] \leq \frac{1}{2} + \text{negl}(\lambda),$$

where the oracle $O(\text{key}, \cdot, \cdot)$ has key material hardwired, and on input a tuple of a global identifier GID , an authority index j , and an attribute-key vector pair (\mathbf{x}, \mathbf{u}) , and responds with a partial secret key computed as $\text{SK}_{j, \text{GID}, \mathbf{x}, \mathbf{u}} \leftarrow \text{KeyGen}(\{\text{PK}_i\}_i, \text{MSK}_j, \text{GID}, \mathbf{x}, \mathbf{u})$. Note that the adversary is only allowed to submit key queries for non-corrupt authorities (i.e., $j \notin S^*$). Also, the adversary \mathcal{A} is admissible as long as every secret key query made by \mathcal{A} to the key generation oracle O satisfies the condition that — (1) either $C_0^*(\mathbf{x}) = C_1^*(\mathbf{x}) = 0$, or (2) $C_0^* = C_1^*$ and $\langle \mathbf{u}, \mathbf{v}_0^* \rangle = \langle \mathbf{u}, \mathbf{v}_1^* \rangle$, or (3) $C_0^* = C_1^*$ and adversary does not query at least one non-corrupt authority for the same global identifier, attribute-key vector tuple $(\text{GID}, \mathbf{x}, \mathbf{u})$.

8.2 Construction

Let Gen be a bilinear group generator with asymmetric source groups of prime order p , and k be the parameter used to select the computational hardness assumption as in [MJ18]. Also, let $\mathbb{G}_1, \mathbb{G}_2$ and \mathbb{G}_T be the source and target groups, respectively. Similar to [MJ18], we rely on a hash function $H : \mathbb{G}_2 \times \{0, 1\}^\lambda \times \mathbb{Z}_p^n \times \mathbb{Z}_p^\ell \rightarrow \mathbb{G}_2^{k+1}$ that is later modelled as a random oracle. Below we provide our decentralized inner-product FE scheme based on bilinear maps. We are using a notation similar to that used in [MJ18, Section 3.2] for ease of exposition.

$\text{GSetup}(1^\lambda) \rightarrow \text{PP}$. The setup algorithm samples a bilinear group Π as follows

$$\Pi = (p, \mathbb{G}_1, \mathbb{G}_2, \mathbb{G}_T, e(\cdot, \cdot)) \leftarrow \text{Gen}(1^\lambda).$$

It next samples random generators $g_1 \in \mathbb{G}_1, g_2 \in \mathbb{G}_2$. Additionally it chooses random matrices \mathbf{A}, \mathbf{U} as $\mathbf{A} \leftarrow \mathbb{Z}_p^{(k+1) \times k}$ and $\mathbf{U} \leftarrow \mathbb{Z}_p^{(k+1) \times (k+1)}$, and sets the public parameters as

$$\text{PP} = \left(\Pi, g_1 = [1]_1, g_2 = [1]_2, [\mathbf{A}]_1, [\mathbf{U}^\top \mathbf{A}]_1 \right).$$

$\text{LSetup}(\text{PP}, 1^n, 1^\ell, i) \rightarrow (\text{PK}, \text{MSK})$. The algorithm samples random matrices $\mathbf{W}_i \leftarrow \mathbb{Z}_p^{(k+1) \times (k+1)}$ and $\Delta_i \leftarrow \mathbb{Z}_p^{(k+1) \times \ell}$. It also samples a random exponent $\sigma_i \leftarrow \mathbb{Z}_p$, and sets the authority public-secret key pair as

$$\text{PK} = \left(\text{PP}, [\mathbf{W}_i^\top \mathbf{A}]_1, [\Delta_i^\top \mathbf{A}]_T, y_i = [\sigma_i]_2 \right), \quad \text{MSK} = (\mathbf{W}_i, \Delta_i, \sigma_i).$$

$\text{KeyGen}(\{\text{PK}_i\}_{i \in [n]}, \text{MSK}_j, \text{GID}, \mathbf{x}, \mathbf{u}) \rightarrow \text{SK}_{j, \text{GID}, \mathbf{x}, \mathbf{u}}$. It parses the authority public keys as described above, and first computes a masking value $\mu_j \in \mathbb{Z}_p^{k+1}$ as

$$[\mu_j]_2 = \prod_{i=1}^{j-1} H([\sigma_i \sigma_j]_2, \text{GID}, \mathbf{x}, \mathbf{u}) \odot \prod_{i=j+1}^n H([\sigma_i \sigma_j]_2, \text{GID}, \mathbf{x}, \mathbf{u}),$$

¹⁶In this work, we only focus on standard semantic security, but one could also amplify to its CCA counterpart by relying on recently developed generic techniques [KW19].

where the masking value is implicitly sampled. Next, it also implicitly samples $\mathbf{h} \in \mathbb{Z}_p^{k+1}$ as $[\mathbf{h}]_2 = H(g_2, \text{GID}, \mathbf{x}, \mathbf{u})$. Finally, it outputs the secret key as

$$\text{SK}_{j, \text{GID}, \mathbf{x}, \mathbf{u}} = ([\Delta_j \mathbf{u} - x_j \mathbf{W}_j \mathbf{h} + \mu_j]_2, \mathbf{x}, \mathbf{u}).$$

$\text{Enc}(\{\text{PK}_i\}_{i \in [n]}, \mathbf{y}, \mathbf{v}) \rightarrow \text{CT}$. The encryption algorithm first parses the keys PK_i as defined during setup. It samples a random exponent vector $\mathbf{s} \leftarrow \mathbb{Z}_p^k$, and outputs the ciphertext CT as

$$\text{CT} = \left(C_0 = [\mathbf{A}\mathbf{s}]_1, \left\{ C_i = [(y_i \mathbf{U}^\top + \mathbf{W}_i^\top) \mathbf{A}\mathbf{s}]_1 \right\}_{i \in [n]}, C_m = [\sum_{i=1}^n \Delta_i^\top \mathbf{A}\mathbf{s} + \mathbf{v}]_T \right).$$

$\text{Dec}(\{\text{SK}_{i, \text{GID}, \mathbf{x}, \mathbf{u}}\}_{i \in [n]}, \text{CT}) \rightarrow m$. It parses the secret key and ciphertext as described above. The decryptor first combines the partial keys to obtain key term $K = \odot_{i=1}^n [\Delta_i \mathbf{u} - x_i \mathbf{W}_i \mathbf{h} + \mu_i]_2$. The algorithm then recovers the inner product by computing the discrete log of the following

$$\frac{\langle \mathbf{u}, C_m \rangle}{e\langle C_0, K \rangle \cdot \prod_{i \in [n]} e\langle C_i, [x_i \mathbf{h}]_2 \rangle},$$

where the operation $\langle \mathbf{u}, C_m \rangle$ computes the encoded inner product by first raising \mathbf{u} in the exponent of each term (component-by-component), and then followed by multiplication of the resulting encodings resulting in the inner product being computed in the exponent. Also, the operation $e\langle \cdot, \cdot \rangle$ carries the same inner product operation where the pairing function $e(\cdot, \cdot)$ is being used for computing the product terms instead.

8.3 Correctness and Security

In this section, we provide the correctness proof and a high level description of security games and prove indistinguishability of successive games.

Correctness. The proof of correctness is along the lines of the Michalevsky-Joye [MJ18] scheme. Below we briefly highlight the main points.

Consider n partial secret keys $\text{SK}_{i, \text{GID}, \mathbf{x}, \mathbf{u}} = ([\Delta_i \mathbf{u} - x_i \mathbf{W}_i \mathbf{h} + \mu_i]_2, \mathbf{x}, \mathbf{u})$ for authority indices $i \in [n]$ and key vectors \mathbf{x}, \mathbf{u} . First, note that the key product term K can be simplified as follows:

$$\begin{aligned} K &= \odot_{i=1}^n [\Delta_i \mathbf{u} - x_i \mathbf{W}_i \mathbf{h} + \mu_i]_2 \\ &= \odot_{i=1}^n [\Delta_i \mathbf{u} - x_i \mathbf{W}_i \mathbf{h}]_2 \quad (\text{since } \sum_i \mu_i = \mathbf{0}) \\ &= [\Delta \mathbf{u} - \sum_{i=1}^n x_i \mathbf{W}_i \mathbf{h}]_2 \quad (\text{where } \Delta = \sum_i \Delta_i) \end{aligned}$$

Therefore, combining with the fact that $C_0 = [\mathbf{A}\mathbf{s}]_1$ we get that

$$e\langle C_0, K \rangle = [(\mathbf{u}^\top \Delta^\top - \sum_{i=1}^n x_i \mathbf{h}^\top \mathbf{W}_i^\top) \mathbf{A}\mathbf{s}]_T.$$

Also, we can simplify the terms $e\langle C_i, [x_i \mathbf{h}]_2 \rangle$ and $\prod_i e\langle C_i, [x_i \mathbf{h}]_2 \rangle$ as

$$\begin{aligned} e\langle C_i, [x_i \mathbf{h}]_2 \rangle &= [(x_i y_i \mathbf{h}^\top \mathbf{U}^\top + x_i \mathbf{h}^\top \mathbf{W}_i^\top) \mathbf{A}\mathbf{s}]_T, \\ \prod_{i \in [n]} e\langle C_i, [x_i \mathbf{h}]_2 \rangle &= [(\mathbf{x}^\top \mathbf{y}) \mathbf{h}^\top \mathbf{U}^\top + \sum_{i=1}^n x_i \mathbf{h}^\top \mathbf{W}_i^\top] \mathbf{A}\mathbf{s}]_T. \end{aligned}$$

Lastly, we also get that $\langle \mathbf{u}, C_m \rangle = [\mathbf{u}^\top \Delta^\top \mathbf{A}\mathbf{s} + \mathbf{u}^\top \mathbf{v}]_T$ where $\Delta = \sum_i \Delta_i$. Combining all these terms as done during decryption, we get that

$$\frac{\langle \mathbf{u}, C_m \rangle}{e\langle C_0, K \rangle \cdot \prod_{i \in [n]} e\langle C_i, [x_i \mathbf{h}]_2 \rangle} = [\mathbf{u}^\top \mathbf{v} - (\mathbf{x}^\top \mathbf{y}) \mathbf{h}^\top \mathbf{U}^\top \mathbf{A}\mathbf{s}]_T.$$

Therefore, whenever $\langle \mathbf{x}, \mathbf{y} \rangle = \mathbf{x}^\top \mathbf{y} = 0$, the above gets simplified to $[\mathbf{u}^\top \mathbf{v}]_T$ which can be recovered by computing the discrete log. Thus, correctness follows.

Security. Below we briefly highlight the main ideas behind the security proof.

Theorem 8.2. If the k -linear assumption (assumption 2.3) holds in both the source groups \mathbb{G}_1 and \mathbb{G}_2 over the group generator Gen , then the scheme described above is a selectively secure decentralized AB-IPFE with 1-sided policy hiding for inner product testing predicates as per Definition 8.1 with at least two honest authorities.

The proof structure is similar to that of [MJ18], except now we need additionally to answer key queries for attribute-key vector pairs (\mathbf{x}, \mathbf{u}) such that $\langle \mathbf{x}, \mathbf{y}^* \rangle = 0$ where $\mathbf{y}^*, (\mathbf{v}_0^*, \mathbf{v}_1^*)$ are the challenge key and message vectors (respectively).¹⁷ At a high level, the idea is to sample the central secret key term Δ such that the reduction could honestly compute the secret keys for all attribute-key vector pairs (\mathbf{x}, \mathbf{u}) such that $\langle \mathbf{x}, \mathbf{y}^* \rangle = 0$ and $\langle \mathbf{u}, \mathbf{v}_0^* \rangle = \langle \mathbf{u}, \mathbf{v}_1^* \rangle$, while switching the secret keys to their semi-functional counterparts for attribute-key vector pairs (\mathbf{x}, \mathbf{u}) where $\langle \mathbf{x}, \mathbf{y}^* \rangle \neq 0$.

For ease of exposition, we prove security in an alternate game where the challenger answers each key query with the full (combined) key term K instead of providing partial key shares (i.e., per authority shares). Similar approach was taken in [MJ18, Appendix B], where the idea was to use the fact that by programming the masking terms μ_i 's appropriately, the security proof could be easily extended to the more general case where the challenger responds with partial shares instead of the full key. Since the same idea can be used for our scheme as well, thus we prove security in the simpler model and refer to [MJ18, Appendix B] for more details.

Proof. We start by sketching the sequence of games where the first game corresponds to the original security game. First, we switch the way we sample the master secret key wherein instead of sampling the random master secret matrix Δ directly¹⁸, we sample a random matrix Γ of same dimensions and implicitly set $\Delta = \Gamma \mathbf{R}$ where $\mathbf{R} \in \mathbb{Z}_p^{\ell \times \ell}$ is a random full rank matrix such that $\mathbf{R}(\mathbf{v}_0^* - \mathbf{v}_1^*) = \mathbf{e}_1$ where \mathbf{e}_1 is the first vector in the canonical basis of \mathbb{Z}_p^ℓ . Next, we switch the way the random oracle queries are answered wherein we sample encoding $[\mathbf{h}]_2$ from a random k -dimensional subspace. This is followed by switching all the secret key queries for key vector pairs (\mathbf{x}, \mathbf{u}) to semi-functional whenever $\langle \mathbf{x}, \mathbf{y}_\beta^* \rangle = 0$. Next, we switch the ciphertext to a *partial* semi-functional ciphertext. Here by a partial semi-functional ciphertext we mean that we split the challenge message vector such that now we encrypt the vector $\beta(\mathbf{v}_0^* - \mathbf{v}_1^*)$ as a semi-functional ciphertext (for random challenge bit β) and vector \mathbf{v}_1^* as a normal (non-semi-functional) ciphertext and homomorphically combines both of them to create the final ciphertext. The idea here is that since adversary never receives any secret key for a key vector pair (\mathbf{x}, \mathbf{u}) where both $\langle \mathbf{x}, \mathbf{y}_\beta^* \rangle = 0$ and $\langle \mathbf{u}, \mathbf{v}_0^* - \mathbf{v}_1^* \rangle \neq 0$, thus the reduction algorithm can perfectly simulate the above games using the dual system encryption paradigm [Wat09] as used in [MJ18]. Below we provide a high level description of the security games.

Description of games

Game 0. This corresponds to the modified AB-IPFE security game, where the challenger answers key queries with the combined secret keys instead of partial key shares. Let Q denote the total number of key queries made by adversary.

Game 1. This is same as the previous game, except the challenger:

- samples a uniformly random matrix $\Gamma \leftarrow \mathbb{Z}_p^{(k+1) \times \ell}$,
- samples a random *orthogonal* matrix $\mathbf{R} \in \mathbb{Z}_p^{\ell \times \ell}$ subject to the constraint that $\mathbf{R}(\mathbf{v}_0^* - \mathbf{v}_1^*) = \mathbf{e}_1$, where $\mathbf{v}_0^*, \mathbf{v}_1^*$ are the challenge message vectors and $\mathbf{e}_1 = (1, 0, \dots, 0)^\top$ (i.e., the first canonical basis vector of \mathbb{Z}_p^ℓ), and
- sets $\Delta = \Gamma \mathbf{R}$ instead of sampling it uniformly at random.

¹⁷Note that in the case, the adversary submits two distinct key vectors \mathbf{y}_0^* and \mathbf{y}_1^* as its challenge, then it must be that $\langle \mathbf{x}, \mathbf{y}_0^* \rangle = \langle \mathbf{x}, \mathbf{y}_1^* \rangle = 0$. In this case, we can use the same proof strategy as [MJ18].

¹⁸Recall that we are using Δ to denote the term $\sum_i \Delta_i$ as in the correctness proof.

Game 2. This is same as the previous game, except the challenger answers the random oracle queries as follows:

- during setup, it samples a random matrix $\mathbf{B} \leftarrow \mathbb{Z}_p^{(k+1) \times k}$, and each fresh random oracle query is answered by sampling a random vector $\mathbf{r} \leftarrow \mathbb{Z}_p^k$ and responding with $[\mathbf{h}]_2 = [\mathbf{Br}]_2$.

Game 3.q.1. This is same as Game 2, except the challenger samples random matrices $\mathbf{A}, \mathbf{B} \leftarrow \mathbb{Z}_p^{(k+1) \times k}$ along with vectors $\mathbf{a}^\perp, \mathbf{b}^\perp \in \mathbb{Z}_p^{k+1}$ such that $\mathbf{A}^\top \mathbf{a}^\perp = \mathbf{B}^\top \mathbf{b}^\perp = \mathbf{0}$, and a random exponent $\hat{t} \leftarrow \mathbb{Z}_p$, and the key queries are answered as follows:

- for the first $q - 1$ key queries, on key vector pair (\mathbf{x}, \mathbf{u}) , it checks whether $\langle \mathbf{x}, \mathbf{y}_\beta^* \rangle = 0$ or not. If it is equal to 0, then it computes the key K honestly as

$$K = [\Delta \mathbf{u} - \sum_{i=1}^n x_i \mathbf{W}_i \mathbf{h}]_2.$$

Otherwise, if the predicate is not satisfied, then it computes the key K as

$$K = [\Delta \mathbf{u} + \hat{t} \langle \mathbf{u}, \mathbf{v}_0^* - \mathbf{v}_1^* \rangle \mathbf{a}^\perp - \sum_{i=1}^n x_i \mathbf{W}_i \mathbf{h}]_2,$$

where h in both cases is computed as in Game 2, i.e. by picking a random r and computing $[\mathbf{h}]_2 = [\mathbf{Br}]_2$.

- for the q -th key query (\mathbf{x}, \mathbf{u}) , it again checks whether $\langle \mathbf{x}, \mathbf{y}_\beta^* \rangle = 0$ or not. Now it computes the key K (in both cases) as

$$K = [\Delta \mathbf{u} - \sum_{i=1}^n x_i \mathbf{W}_i \mathbf{h}]_2,$$

except if the predicate is satisfied (that is, $\langle \mathbf{x}, \mathbf{y}_\beta^* \rangle = 0$), then it computes h as in Game 2, i.e. by picking a random r and computing $[\mathbf{h}]_2 = [\mathbf{Br}]_2$; whereas if the predicate is not satisfied, then it computes h by sampling random $\mathbf{r} \leftarrow \mathbb{Z}_p^k$ and $\hat{r} \leftarrow \mathbb{Z}_p$, and setting $[\mathbf{h}]_2 = [\mathbf{Br} + \mathbf{a}^\perp \hat{r}]_2$.

- remaining $Q - q$ key queries are answered exactly as in Game 2. That is, it computes the key K (in both cases) as

$$K = [\Delta \mathbf{u} - \sum_{i=1}^n x_i \mathbf{W}_i \mathbf{h}]_2,$$

where h is computed as $[\mathbf{h}]_2 = [\mathbf{Br}]_2$.

In the above game, the random oracle queries are also answered similarly wherein the challenger does not sample the full secret key, but instead it just computes the hash vector \mathbf{h} depending on the query counter and the type of query.

Game 3.q.2. This is same as the previous game, except the way the challenger answers the q -th query:

- for the q -th key query (\mathbf{x}, \mathbf{u}) , it checks whether $\langle \mathbf{x}, \mathbf{y}_\beta^* \rangle = 0$ or not. If it is equal to 0, then it computes the key K honestly as

$$K = [\Delta \mathbf{u} - \sum_{i=1}^n x_i \mathbf{W}_i \mathbf{h}]_2,$$

where h is computed as $[\mathbf{h}]_2 = [\mathbf{Br}]_2$. Otherwise, if the predicate is not satisfied, then it computes the key K as

$$K = [\Delta \mathbf{u} + \hat{t} \langle \mathbf{u}, \mathbf{v}_0^* - \mathbf{v}_1^* \rangle \mathbf{a}^\perp - \sum_{i=1}^n x_i \mathbf{W}_i \mathbf{h}]_2,$$

where h is computed by sampling random $\mathbf{r} \leftarrow \mathbb{Z}_p^k$ and $\hat{r} \leftarrow \mathbb{Z}_p$, and setting $[\mathbf{h}]_2 = [\mathbf{Br} + \mathbf{a}^\perp \hat{r}]_2$.

Game 3.q.3. This is same as the previous game, except when it answers the q -th query, then it computes h as $[\mathbf{h}]_2 = [\mathbf{Br}]_2$ irrespective of whether the predicate is satisfied or not. (Note that Game 3.0.3 is same as Game 2.)

Game 4. This is the same as Game 3.Q.3, except that the challenge ciphertext is semi-functional. That is, the challenger samples a random vector $\mathbf{z} \leftarrow \mathbb{Z}_p^{k+1}$ and computes the challenge ciphertext as:

$$CT = \left(C_0 = [\mathbf{z}]_1, \left\{ C_i = [(y_{\beta,i}^* \mathbf{U}^\top + \mathbf{W}_i^\top) \mathbf{z}]_1 \right\}_{i \in [n]}, C_m = [\Delta^\top \mathbf{z} + \beta \mathbf{v}_\beta^*]_T \right),$$

where β is the random challenge bit.

Game 5. This is the same as Game 4, except that the challenge ciphertext encrypts a message vector of the form $\mathbf{v}_0^* + \eta(\mathbf{v}_1^* - \mathbf{v}_0^*)$ instead of $\beta \mathbf{v}_\beta^* = \mathbf{v}_0^* + \beta(\mathbf{v}_1^* - \mathbf{v}_0^*)$, where η is sampled uniformly at random. That is,

$$C_m = [\Delta^\top \mathbf{z} + \mathbf{v}_0^* + \eta(\mathbf{v}_1^* - \mathbf{v}_0^*)]_T,$$

where η is a random exponent.

Indistinguishability of games

We complete the proof by showing that adjacent games are indistinguishable. For any adversary \mathcal{A} and game X , we denote by $\text{Adv}_S^{\mathcal{A}}(\lambda)$, the probability that \mathcal{A} wins in game S .

Lemma 8.3. For any (potentially unbounded) adversary \mathcal{A} , we have that $\text{Adv}_0^{\mathcal{A}}(\lambda) = \text{Adv}_1^{\mathcal{A}}(\lambda)$.

Proof. This follows directly from the fact that, for any invertible matrix $\mathbf{R} \in \mathbb{Z}_p^{\ell \times \ell}$, the following distributions are identical.

$$\left\{ \Delta : \Delta \leftarrow \mathbb{Z}_p^{(k+1) \times \ell} \right\} \equiv \left\{ \Gamma \mathbf{R} : \Gamma \leftarrow \mathbb{Z}_p^{(k+1) \times \ell} \right\}.$$

□

Lemma 8.4. If k -linear assumption (assumption 2.3) holds in \mathbb{G}_2 over the group generator Gen , then for any PPT adversary \mathcal{A} , we have that $\text{Adv}_1^{\mathcal{A}}(\lambda) - \text{Adv}_2^{\mathcal{A}}(\lambda) \leq \text{negl}(\lambda)$ for some negligible function $\text{negl}(\cdot)$.

Proof. The proof of this lemma is identical to that of [MJ18, Lemma 1].

□

Lemma 8.5. If k -linear assumption (assumption 2.3) holds in \mathbb{G}_2 over the group generator Gen , then for any PPT adversary \mathcal{A} and $q \in \{1, \dots, Q\}$, we have that $\text{Adv}_{3.(q-1).3}^{\mathcal{A}}(\lambda) - \text{Adv}_{3.q.1}^{\mathcal{A}}(\lambda) \leq \text{negl}(\lambda)$ for some negligible function $\text{negl}(\cdot)$.

Proof. The proof of this lemma is identical to that of [MJ18, Lemma 2].

□

Lemma 8.6. For any (potentially unbounded) adversary \mathcal{A} and $q \in \{1, \dots, Q\}$, we have that $\text{Adv}_{3.q.1}^{\mathcal{A}}(\lambda) - \text{Adv}_{3.q.2}^{\mathcal{A}}(\lambda) \leq \text{negl}(\lambda)$ for some negligible function $\text{negl}(\cdot)$.

Proof. The proof of this lemma is identical to that of [MJ18, Lemma 3].

□

Lemma 8.7. If k -linear assumption (assumption 2.3) holds in \mathbb{G}_2 over the group generator Gen , then for any PPT adversary \mathcal{A} and $q \in \{1, \dots, Q\}$, we have that $\text{Adv}_{3.q.2}^{\mathcal{A}}(\lambda) - \text{Adv}_{3.q.3}^{\mathcal{A}}(\lambda) \leq \text{negl}(\lambda)$ for some negligible function $\text{negl}(\cdot)$.

Proof. The proof of this lemma is identical to that of [MJ18, Lemma 4].

□

Lemma 8.8. If k -linear assumption (assumption 2.3) holds in \mathbb{G}_1 over the group generator Gen , then for any PPT adversary \mathcal{A} , we have that $\text{Adv}_{3.Q.3}^{\mathcal{A}}(\lambda) - \text{Adv}_4^{\mathcal{A}}(\lambda) \leq \text{negl}(\lambda)$ for some negligible function $\text{negl}(\cdot)$.

Proof. The proof of this lemma is identical to that of [MJ18, Lemma 5].

□

Lemma 8.9. For any (potentially unbounded) adversary \mathcal{A} , we have that $\text{Adv}_4^{\mathcal{A}}(\lambda) - \text{Adv}_5^{\mathcal{A}}(\lambda) \leq \text{negl}(\lambda)$ for some negligible function $\text{negl}(\cdot)$.

Proof. The proof of this lemma is similar to that of [MJ18, Lemma 6], except now the secret matrix Δ is set in a more intricate way. First, note that in game 4, we can rewrite the message component of the ciphertext (i.e., C_m) as follows:

$$\begin{aligned} C_m &= [\Delta^\top \mathbf{z} + \beta \mathbf{v}_\beta]_T \\ &= [\mathbf{R}^\top \Gamma^\top \mathbf{z} + \beta(\mathbf{v}_1^* - \mathbf{v}_0^*) + \mathbf{v}_0^*]_T \end{aligned}$$

Since we sampled matrix \mathbf{R} such that it is a random *orthogonal* ($\ell \times \ell$) matrix with the constraint that $\mathbf{R}(\mathbf{v}_0^* - \mathbf{v}_1^*) = \mathbf{e}_1$, thus we have that $\mathbf{R}^\top \mathbf{e}_1 = \mathbf{v}_0^* - \mathbf{v}_1^*$. So, we can write C_m as

$$C_m = [\mathbf{R}^\top \Gamma^\top \mathbf{z} - \beta \mathbf{R}^\top \mathbf{e}_1 + \mathbf{v}_0^*]_T = [\mathbf{R}^\top (\Gamma^\top \mathbf{z} - \beta \mathbf{e}_1) + \mathbf{v}_0^*]_T.$$

Now the idea is to sample matrix Δ as $\Delta = \hat{\Delta} - (\mathbf{e}_1^\top \otimes \mathbf{a}^\perp) \hat{t}$ for a random matrix $\hat{\Delta}$ and random exponent \hat{t} . Here \otimes denotes the tensoring operation.

With this modification, we can apply a proof strategy similar to that in [MJ18, Lemma 6], where we rely on the fact that the public keys and public parameters can be sampled without the knowledge of \hat{t} (i.e., only given the matrix $\hat{\Delta}$ and other public terms). Next, we can show that for every key query (\mathbf{x}, \mathbf{u}) we again do not need \hat{t} , and can only be answered using $\hat{\Delta}$. For this, consider two cases.

First, when $\langle \mathbf{x}, \mathbf{y}_\beta^* \rangle = 0$, then it must be the case that $\langle \mathbf{u}, \mathbf{v}_0^* - \mathbf{v}_1^* \rangle = 0$, therefore $\mathbf{u} \in \text{Span}(\{\mathbf{R}^\top \mathbf{e}_i\}_{i \in [n] \setminus \{1\}})$. This gives that $\Delta \mathbf{u} = \Gamma \mathbf{R} \mathbf{u} \in \Gamma \cdot \text{Span}(\{\mathbf{e}_i\}_{i \in [n] \setminus \{1\}})$, that is it does not depend on the first column of Γ which depends on \hat{t} .

Second, when $\langle \mathbf{x}, \mathbf{y}_\beta^* \rangle \neq 0$, then we get that the first column entry of $\mathbf{R} \mathbf{u}$ corresponds to $\langle \mathbf{u}, \mathbf{v}_0^* - \mathbf{v}_1^* \rangle$, and since the key term contains the term $\Delta \mathbf{u} + \hat{t} \langle \mathbf{u}, \mathbf{v}_0^* - \mathbf{v}_1^* \rangle \mathbf{a}^\perp$, therefore this gets simplified to $\hat{\Delta} \mathbf{u}$. Thus, this can also be computed without the knowledge of \hat{t} .

Finally, we notice that setting \mathbf{z} in the challenge ciphertext as $\mathbf{A} \mathbf{s} + \mathbf{b}^\perp \hat{s}$ (where $\mathbf{s} \leftarrow \mathbb{Z}_p^k$, $\hat{s} \leftarrow \mathbb{Z}_p$), we obtain that C_m completely hides the challenge bit β . This completes the proof. \square

Lemma 8.10. If the k -linear assumption (assumption 2.3) holds over the group generator Gen , then for any PPT adversary \mathcal{A} , we have that $\text{Adv}_5^{\mathcal{A}}(\lambda) \leq \text{negl}(\lambda)$ for some negligible function $\text{negl}(\cdot)$.

Proof. If $\mathbf{y}_0^* = \mathbf{y}_1^*$, then the advantage $\text{Adv}_5^{\mathcal{A}}(\lambda) = 0$ since the challenge ciphertext is independent of the challenge bit β . Whereas if $\mathbf{y}_0^* \neq \mathbf{y}_1^*$, then it must be the case that all the adversary's key queries were for rejecting attributes that is $\langle \mathbf{x}, \mathbf{y}_0^* \rangle = \langle \mathbf{x}, \mathbf{y}_1^* \rangle = 0$ for all vectors \mathbf{x} queried by \mathcal{A} . In the latter case, we have that all the keys queried have been semi-functional functional above, and thus we can make the entire ciphertext semi-functional as well by using a similar hybrid structure as above, and finally by using an argument similar to that of [MJ18, Lemma 7], we can conclude that the challenge ciphertext encrypts a random vector \mathbf{y} instead of \mathbf{y}_β^* thereby completing the proof. \square

\square

9 New Primitives Predicted by MPFE

In this section, we define new primitives that can be seen as special cases of multi-party functional encryption and have natural, compelling applications. We also provide a high level sketch of a generic construction of our new reputation point based encryption described in Section 1.5 from a multi-authority ABE scheme.

Reputation Point Based Encryption: Below we sketch the constructions in two different setup modes which leads to different trade-offs in the type of collusion security and aggregation functionality.

Centralized setup with bounded aggregation. In the centralized setup setting, the central authority on input the upper bound, N , on the number of user key aggregations samples N pairs of MA-ABE authority public-secret key pairs $(\text{PK}_i, \text{MSK}_i)$, and sets the the master public and secret keys as the sequence of these N keys that is $(\text{PK}_i)_i, (\text{MSK}_i)_i$. Now to generate a key for tag T and points v , the authority starts by sampling a random label L and generates N MA-ABE partial decryption keys by running the MA-ABE key generation algorithm with respect to each key MSK_i with global identifier $\text{GID} = T$ and partial attribute vector (T, L, v) . An encryptor simply hides the message under the policy P which on input a sequence $(T_i, L_i, v_i)_i$ checks that all tags match the encrypted tag T , all labels are distinct, and the user points sum to a value larger than the encrypted threshold w . Note that here we are using randomly sampled labels in each user’s key for ensuring that a single user can not combine two components of its own key, while N distinct users can always combine one key from their sequence of N MA-ABE keys with all but negligible probability. At a high level, the above idea seems to work for any MA-ABE scheme but one needs to be careful about the notion of security the underlying MA-ABE scheme must satisfy as now in this system we ask each authority to possibly generate multiple keys for the same global identifier.

Decentralized setup with security in the erasure model. Note that one could try to make the above construction fully decentralized by asking each user to locally sample its own public-secret key pair (PK, MSK) for a MA-ABE scheme, and then the user locally generates a partial predicate key $\text{SK}_{T,v}$ for attribute (T, v) . Now each user must *delete* its master secret key MSK , and it outputs PK as its public key and $\text{SK}_{T,v}$ as its decryption key. An encryptor simply takes as input the list of public keys of all users which it possibly wants to encrypt to, and defines the corresponding policy circuit as follows — either the attribute string is not specified, or it the specified tag matches the encrypted tag T , and the user points sum to a value larger than the encrypted threshold w . To make this work, we need the MA-ABE scheme to support the concept of empty key (or zero-keys) that is the partial decryption key for all zeros attribute is simply the empty string, and the fact that an attacker never corrupts the master secret key of the underlying MA-ABE system for any user (that is, the random coins of every user’s setup are erased). We leave further analysis of above approaches for future work.

Other New Primitives for Future Work. Providing more constructions is beyond the scope of the present work, so we outline syntax of some new primitives here and leave instantiation to future work.

1. *Generalizing MIFE, MCFE and PHFE:* In many cases it is useful to combine ciphertexts from multiple users if a more complex policy is satisfied by their public labels: for instance, we may want to compute on all ciphertexts that were generated during a specified time interval. Here, the labels may be set as the timestamp at which a ciphertext was generated. The above example motivates defining the notion of *partially hiding* multi-input FE, where the input of party P_i is the pair $(\text{lab}_i, \mathbf{x}_i)$ where lab_i is public and \mathbf{x}_i is private. Given encryptions of $(\text{lab}_i, \mathbf{x}_i)$ for $i \in [n]$ and a function key for $f = (f_1, f_2)$, if $f_1(\text{lab}_1, \dots, \text{lab}_n) = 1$, decryption must output $f_2(\mathbf{x}_1, \dots, \mathbf{x}_n)$. Note that this notion generalizes MIFE, MCFE as well as partially hiding FE [GVW12, AJL+19].
2. *Generalizing Partial Key Combinations:* We may permit more meaningful ways of combining keys than those that have been studied before. For instance, in the context of an investigation, multiple key authorities such as hospital, police station and bank may wish to filter suspects in an encrypted database based on different criteria, for which they issue different keys. Let us say that the input \mathbf{x} contains a list of user records. Then, each authority gives a key that takes an input list of records, tests which records satisfy its criteria and erases the rest. Authority i provides a key for function f_i , and given ciphertext for \mathbf{x} , decryption returns $f_1 \circ \dots \circ f_n(\mathbf{x})$.

References

- [ABB10a] Shweta Agrawal, Dan Boneh, and Xavier Boyen. Efficient lattice (H)IBE in the standard model. In *EUROCRYPT*, pages 553–572, 2010.
- [ABB10b] Shweta Agrawal, Dan Boneh, and Xavier Boyen. Lattice basis delegation in fixed dimension and shorter-ciphertext hierarchical IBE. In *CRYPTO*, pages 98–115, 2010.
- [ABDCP15] Michel Abdalla, Florian Bourse, Angelo De Caro, and David Pointcheval. Simple functional encryption schemes for inner products. In *PKC*, 2015.
- [ABG19] Michel Abdalla, Fabrice Benhamouda, and Romain Gay. From single-input to multi-client inner-product functional encryption. In *ASIACRYPT*, 2019.
- [ABKW19] Michel Abdalla, Fabrice Benhamouda, Markulf Kohlweiss, and Hendrik Waldner. Decentralizing inner-product functional encryption. In *PKC*, 2019.
- [ACF⁺20] Shweta Agrawal, Michael Clear, Ophir Frieder, Sanjam Garg, Adam O’Neill, and Justin Thaler. Ad hoc multi-input functional encryption. In *ITCS 2020*, 2020.
- [ACGU20] Michel Abdalla, Dario Catalano, Romain Gay, and Bogdan Ursu. Inner-product functional encryption with fine-grained access control. In *Asiacrypt*, 2020.
- [AJ15] Prabhanjan Ananth and Abhishek Jain. Indistinguishability obfuscation from compact functional encryption. In *CRYPTO 2015*, 2015.
- [AJL⁺19] Prabhanjan Ananth, Aayush Jain, Huijia Lin, Christian Matt, and Amit Sahai. Indistinguishability obfuscation without multilinear maps: io from lwe, bilinear maps, and weak pseudorandomness. In *Crypto*, 2019.
- [ALS16] Shweta Agrawal, Benoit Libert, and Damien Stehle. Fully secure functional encryption for linear functions from standard assumptions, and applications. In *Crypto*, 2016.
- [AY20] Shweta Agrawal and Shota Yamada. Optimal broadcast encryption from pairings and lwe. In *Proc. of EUROCRYPT*, 2020.
- [BCFG17] Carmen Elisabetta Zaira Baltico, Dario Catalano, Dario Fiore, and Romain Gay. Practical functional encryption for quadratic functions with applications to predicate encryption. In *CRYPTO*, 2017.
- [BCG⁺17] Zvika Brakerski, Nishanth Chandran, Vipul Goyal, Aayush Jain, Amit Sahai, and Gil Segev. Hierarchical functional encryption. In *ITCS 2017*, 2017.
- [BF01] Dan Boneh and Matthew K. Franklin. Identity-based encryption from the weil pairing. In *CRYPTO*, pages 213–229, 2001.
- [BFF⁺14] Gilles Barthe, Edvard Fagerholm, Dario Fiore, John C. Mitchell, Andre Scedrov, and Benedikt Schmidt. Automated analysis of cryptographic assumptions in generic group models. In *CRYPTO*, 2014.
- [BGG⁺14] 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.
- [BJK15] Allison Bishop, Abhishek Jain, and Lucas Kowalczyk. Function-hiding inner product encryption. In *Asiacrypt*, 2015.

- [BLP⁺13] Zvika Brakerski, Adeline Langlois, Chris Peikert, Oded Regev, and Damien Stehlé. Classical hardness of learning with errors. In *STOC*, STOC '13, 2013.
- [BSW11] Dan Boneh, Amit Sahai, and Brent Waters. Functional encryption: Definitions and challenges. In *TCC*, pages 253–273, 2011.
- [BV15] Nir Bitansky and Vinod Vaikuntanathan. Indistinguishability obfuscation from functional encryption. *FOCS*, 2015:163, 2015.
- [BWY11] Mihir Bellare, Brent Waters, and Scott Yilek. Identity-based encryption secure against selective opening attack. In *Theory of Cryptography Conference*, 2011.
- [CC09] Melissa Chase and Sherman S. M. Chow. Improving privacy and security in multi-authority attribute-based encryption. In *ACM CCS 2009*, 2009.
- [CDSG⁺18a] Jérémy Chotard, Edouard Dufour-Sans, Romain Gay, Duong Hieu Phan, and David Pointcheval. Decentralized multi-client functional encryption for inner product. In *Asiacrypt*, 2018.
- [CDSG⁺18b] Jérémy Chotard, Edouard Dufour-Sans, Romain Gay, Duong Hieu Phan, and David Pointcheval. Multi-client functional encryption with repetition for inner product. Cryptology ePrint Archive, Report 2018/1021, 2018.
- [CDSG⁺20] Jérémy Chotard, Edouard Dufour-Sans, Romain Gay, Duong Hieu Phan, and David Pointcheval. Dynamic decentralized functional encryption. In *Crypto*, 2020.
- [Cha07] Melissa Chase. Multi-authority attribute based encryption. In *TCC*, 2007.
- [CHKP10] David Cash, Dennis Hofheinz, Eike Kiltz, and Chris Peikert. Bonsai trees, or how to delegate a lattice basis. In *EUROCRYPT*, pages 523–552, 2010.
- [CZY19] Yuechen Chen, Linru Zhang, and Siu-Ming Yiu. Practical attribute based inner product functional encryption from simple assumptions. Cryptology ePrint Archive, 2019.
- [DOT18] Pratish Datta, Tatsuaki Okamoto, and Katsuyuki Takashima. Adaptively simulation-secure attribute-hiding predicate encryption. Cryptology ePrint Archive, Report 2018/1093, 2018.
- [GGG⁺14] Shafi Goldwasser, S. Dov Gordon, Vipul Goyal, Abhishek Jain, Jonathan Katz, Feng-Hao Liu, Amit Sahai, Elaine Shi, and Hong-Sheng Zhou. Multi-input functional encryption. In *Eurocrypt*, 2014.
- [GKL⁺13] S. Dov Gordon, Jonathan Katz, Feng-Hao Liu, Elaine Shi, and Hong-Sheng Zhou. Multi-input functional encryption. Cryptology ePrint Archive, Report 2013/774, 2013.
- [GPSW06] Vipul Goyal, Omkant Pandey, Amit Sahai, and Brent Waters. Attribute-based encryption for fine-grained access control of encrypted data. In *CCS*, 2006.
- [GPV08] Craig Gentry, Chris Peikert, and Vinod Vaikuntanathan. Trapdoors for hard lattices and new cryptographic constructions. In *STOC*, pages 197–206, 2008.
- [GV15] Sergey Gorbunov and Dhinakaran Vinayagamurthy. Riding on asymmetry: Efficient ABE for branching programs. In *ASIACRYPT*, 2015.
- [GVW12] Sergey Gorbunov, Vinod Vaikuntanathan, and Hoeteck Wee. Functional encryption with bounded collusions from multiparty computation. In *CRYPTO*, 2012.
- [GVW15] Sergey Gorbunov, Vinod Vaikuntanathan, and Hoeteck Wee. Predicate encryption for circuits from lwe. In *Crypto*, 2015.

- [KSW08] Jonathan Katz, Amit Sahai, and Brent Waters. Predicate encryption supporting disjunctions, polynomial equations, and inner products. In *EUROCRYPT*, pages 146–162, 2008.
- [KW19] Venkata Koppula and Brent Waters. Realizing chosen ciphertext security generically in attribute-based encryption and predicate encryption. In *CRYPTO 2019*, 2019.
- [LT19] Benoît Libert and Radu Tîţiu. Multi-client functional encryption for linear functions in the standard model from lwe. In *Asiacrypt*, 2019.
- [LW11] Allison Lewko and Brent Waters. Decentralizing attribute-based encryption. In *Proceedings of EuroCrypt*, 2011.
- [Mau05] Ueli Maurer. Abstract models of computation in cryptography. In *IMA Int. Conf.*, volume 3796 of *Lecture Notes in Computer Science*, pages 1–12. Springer, 2005.
- [MJ18] Yan Michalevsky and Marc Joye. Decentralized policy-hiding ABE with receiver privacy. In *ESORICS*, 2018.
- [MP12] Daniele Micciancio and Chris Peikert. Trapdoors for lattices: Simpler, tighter, faster, smaller. In *EUROCRYPT*, 2012.
- [Reg09] Oded Regev. On lattices, learning with errors, random linear codes, and cryptography. *J.ACM*, 56(6), 2009. extended abstract in STOC’05.
- [Sha84] Adi Shamir. Identity-based cryptosystems and signature schemes. In *Workshop on the theory and application of cryptographic techniques*, 1984.
- [Sho97] Victor Shoup. Lower bounds for discrete logarithms and related problems. In *Eurocrypt*, 1997.
- [SW05] Amit Sahai and Brent Waters. Fuzzy identity-based encryption. In *Eurocrypt*, 2005.
- [Wat09] Brent Waters. Dual system encryption: Realizing fully secure ibe and hibe under simple assumptions. In *Annual International Cryptology Conference*, 2009.

Appendix

A Feasibility for MPFE for General Circuits

In this section we discuss the feasibility of general MPFE. Since the framework of MPFE is very general and supports all existing constructions in the literature (that we are aware of), feasibility varies widely depending on the properties desired, such as whether the setup algorithm must be local, centralized or interactive, whether and encryption keys must be public or private and such others. Since the case of centralized setup has been most widely studied in the literature and it is evident that feasibility results for local and interactive setup also apply for a centralized setup, we focus on these below.

First, we provide some definitions needed for our constructions.

A.1 Definitions

A.1.1 Multi-Input Functional Encryption

An n -input FE scheme [GGG+14] MIFE for a message space $\{\mathcal{M}_\lambda\}_{\lambda \in \mathbb{N}}$ and a functionality $\{\mathcal{F}_\lambda\}_{\lambda \in \mathbb{N}}$, where for each $\lambda \in \mathbb{N}$, each $f \in \mathcal{F}_\lambda$ is a (description of a) function on $(\mathcal{M}_\lambda)^n$, is given by a set of algorithms with the following syntax:

- $\text{MIFE.Setup}(1^\lambda, 1^n)$: A PPT algorithm taking the security parameter λ and number of users n , and outputting the master secret key MSK and encryption keys $(\text{EK}_1, \dots, \text{EK}_n)$.
- $\text{MIFE.KeyGen}(\text{MSK}, f)$: A PT algorithm taking a master secret key MSK , a function $f \in \mathcal{F}_\lambda$ and outputting a corresponding decryption key SK_f .
- $\text{MIFE.Enc}(\text{EK}, \mathbf{x})$: A PPT algorithm taking an encryption key EK and a message $\mathbf{x} \in \mathcal{M}_\lambda$, and outputting a ciphertext c .
- $\text{MIFE.Dec}(\text{SK}_f, (c_1, \dots, c_n))$: A PT algorithm taking decryption key SK_f and vector of ciphertexts (c_1, \dots, c_n) , and outputting a string y .

Correctness We say that MIFE is *correct* if for all $\lambda, n \in \mathbb{N}$, $\mathbf{x}_1 \dots \mathbf{x}_n \in \mathcal{M}_\lambda$ and $f \in \mathcal{F}_\lambda$

$$\Pr \left[\begin{array}{l} ((\text{EK}_1, \dots, \text{EK}_n), \text{MSK}) \leftarrow \text{MIFE.Setup}(1^\lambda) \\ c_i \leftarrow \text{MIFE.Enc}(\text{EK}_i, \mathbf{x}_i) \quad \forall i \in [n] \\ \text{SK}_f \leftarrow \text{MIFE.KeyGen}(\text{MSK}, f) \\ y \leftarrow \text{MIFE.Dec}(\text{SK}_f, (c_1, \dots, c_n)) \end{array} \right] = 1.$$

We remark that our formulation of MIFE assumes that the senders (as in our application an encryptor is referred to as a sender or source) are ordered.

Indistinguishability-Based Security. For an n -input FE scheme MIFE as above and adversary $\mathcal{A} = (\mathcal{A}_0, \mathcal{A}_1, \mathcal{A}_2)$, consider the experiment in Figure A.1.

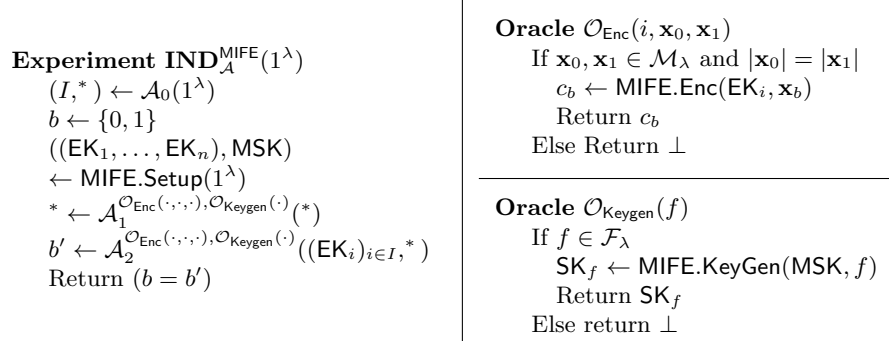


Figure A.1: Experiment for IND-security of standard MIFE.

We call \mathcal{A} *legitimate* if for all $\lambda \in \mathbb{N}$, in all transcripts $\text{IND}_{\mathcal{A}}^{\text{MIFE}}(1^\lambda)$ it holds that for every key generation query f there does not exist two sequences $(y_{1,0}, \dots, y_{n,0})$ and $(y_{1,1}, \dots, y_{n,1})$ such that

$$f(y_{1,0}, \dots, y_{n,0}) \neq f(y_{1,1}, \dots, y_{n,1})$$

and for every $j \in [n]$

- $j \in I$, *i.e.* j is corrupted (so there is no restriction on $y_{j,0}, y_{j,1}$ above), or
- there is an encryption query $(j, \mathbf{x}_0, \mathbf{x}_1)$ such that $y_{j,0} = \mathbf{x}_0$ and $y_{j,1} = \mathbf{x}_1$.

We assume adversaries are legitimate unless otherwise stated. We call \mathcal{A} *passive* if $I = \emptyset$. We call \mathcal{A} *selective* if \mathcal{A}_2 makes no queries. We say that MIFE is IND-secure if for any adversary \mathcal{A}

$$|\Pr \left[\text{IND}_{\mathcal{A}}^{\text{MIFE}}(\cdot) \text{ outputs } 1 \right] - 1/2| = \text{neg}(\lambda).$$

A.1.2 Functional Encryption

A functional encryption scheme, denoted as FE [BSW11], is a tuple of algorithms $\text{FE} = (\text{FE.Setup}, \text{FE.KeyGen}, \text{FE.Enc}, \text{FE.Dec})$ for a message space $\{\mathcal{M}_\lambda\}_{\lambda \in \mathbb{N}}$ and a functionality $\{\mathcal{F}_\lambda\}_{\lambda \in \mathbb{N}}$, where for each $\lambda \in \mathbb{N}$, each $f \in \mathcal{F}_\lambda$ is a (description of a) function on \mathcal{M}_λ . The syntax is the same as for a 1-input MIFE scheme where $\text{EK}_1 = \text{MSK}$ and $I = \emptyset$. The correctness requirement remains the same, as well the notion of indistinguishability based security (which we refer to as “message privacy”).

Function Privacy. We additionally define the notion of function privacy as follows. For an FE scheme FE as above and adversary $\mathcal{A} = (\mathcal{A}_0, \mathcal{A}_1, \mathcal{A}_2)$, consider the experiment in Figure A.2.

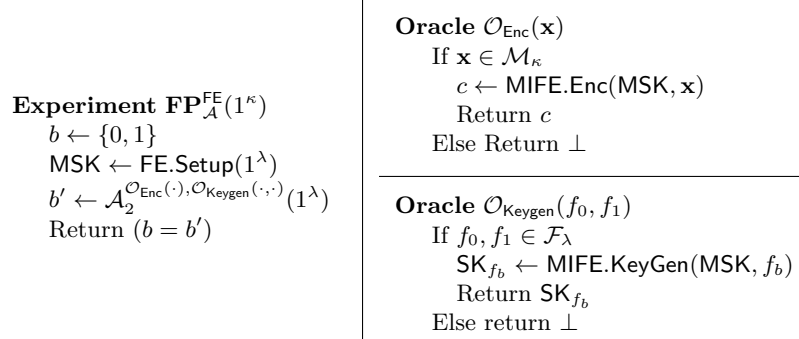


Figure A.2: Experiment for FP-security of FE.

We call \mathcal{A} *legitimate* if for all $\lambda \in \mathbb{N}$, in all transcripts $\mathbf{FP}_{\mathcal{A}}^{\text{FE}}(1^\kappa)$ it holds that for every key generation query f_0, f_1 there does not exist an encryption query $\mathbf{x} \in \mathcal{M}_\kappa$ such that

$$f_0(\mathbf{x}) \neq f_1(\mathbf{x}) .$$

We assume adversaries are legitimate unless otherwise stated. We say that FE is FP-secure if for any adversary \mathcal{A}

$$|\Pr \left[\mathbf{FP}_{\mathcal{A}}^{\text{FE}}(\cdot) \text{ outputs } 1 \right] - 1/2| = \text{negl}(\cdot) .$$

A.1.3 Two-Round MPC

A *2-round MPC protocol* MPC for message-space $\{\mathcal{M}_\lambda\}_{\lambda \in \mathbb{N}}$ and functionality $\{\mathcal{F}_\lambda\}_{\lambda \in \mathbb{N}}$ where for each $\lambda \in \mathbb{N}$ each $f \in \mathcal{F}_\lambda$ is a function on $(\mathcal{M}_\lambda)^n$ for some n , consists of three algorithms with the following syntax:

- **RunRoundOne** $(1^\lambda, 1^n, f, i, x)$: A PPT algorithm taking the security parameter λ , number of users n , a (description of a) function $f \in \mathcal{F}_\lambda$ of arity n , an index $i \in [n]$, an input $x \in \mathcal{M}_\lambda$, and outputting a first protocol message $\rho^{(1)}$ and secret \mathfrak{s} .
- **RunRoundTwo** $(\mathfrak{s}, (\rho_1^{(1)}, \dots, \rho_n^{(1)}))$: A PPT algorithm taking a secret \mathfrak{s} and the first protocol message for all n parties $\rho_1^{(1)}, \dots, \rho_n^{(1)}$, and outputting a second protocol message $\rho^{(2)}$.
- **ComputeResult**: A PT algorithm taking as input the n second-round protocol messages $\rho_1^{(2)}, \dots, \rho_n^{(2)}$ for each party and outputting a value y .

Correctness. We say that MPC is *correct* if for all $\lambda, n, \in \mathbb{N}$, $\mathbf{x}_1 \dots \mathbf{x}_n \in \mathcal{M}_\lambda$ and $f \in \mathcal{F}_\lambda$

$$\Pr \left[\begin{array}{l} (\rho_i^{(1)}, \mathfrak{s}_i) \leftarrow \text{RunRoundOne}(1^\lambda, 1^n, f, i, \mathbf{x}_i) \quad \forall i \in [n] \\ y = f(\mathbf{x}) : \rho_i^{(2)} \leftarrow \text{RunRoundTwo}(\mathfrak{s}_i, (\rho_1^{(1)}, \dots, \rho_n^{(1)})) \quad \forall i \in [n] \\ y \leftarrow \text{ComputeResult}(\rho_1^{(2)}, \dots, \rho_n^{(2)}) \end{array} \right] = 1.$$

Remark A.1. The above definition of two-round MPC is without setup (*i.e.*, a CRS). We also consider the case that there is an additional algorithm CRSGen taking 1^λ and outputting a common reference string CRS that is input to the remaining algorithms. We call this two-round MPC *in the CRS model*.

We call MPC *input-rerunnable* (resp. *function-rerunnable*) if it is *input-delayed* (resp. *function-delayed*) and if RunRoundTwo can be executed multiple times with different input choices (resp. function choices) while still preserving the security properties of the MPC protocol.

Security. Let MPC be a 2-round MPC protocol as above. Let Coins be the coin-space for the protocol. For an adversary $\mathcal{A} = (\mathcal{A}_0, \mathcal{A}_1)$ and simulator SS , consider the experiments in Figure A.3. We say that \mathcal{A} is *passive* (aka. semi-honest) of $I = \emptyset$. We say that MPC is SIM-secure if for any PPT adversary \mathcal{A} there is a stateful PPT simulator $SS = (\text{CRSGen}, \text{Extract}, \text{Sim})$ such that $\text{REAL}_\mathcal{A}^{\text{MPC}}(\cdot)$ and $\text{IDEAL}_{\mathcal{A}, SS}^{\text{MPC}}(\cdot)$ are computationally indistinguishable. Note that simulation of the first-round protocol messages for the

Experiment $\text{REAL}_\mathcal{A}^{\text{MPC}}(1^\kappa)$	Experiment $\text{IDEAL}_{\mathcal{A}, SS}^{\text{MPC}}(1^\kappa)$
<p>(Optional) $\text{crs} \leftarrow \text{CRSGen}(1^\lambda)$ $(1^n, I, f, (x_i)_{i \notin I}) \leftarrow \mathcal{A}_0(1^\kappa)$ <i>//</i> $f \in \mathcal{F}_\lambda$ of arity n, $x_i \in \mathcal{M}_\lambda$ $((\rho_i^{(1)})_{i \in I}, \text{st}) \leftarrow \mathcal{A}_1(n, I, f)$ For $i \notin I$ do: $r_i \leftarrow \text{Coins}(1^\kappa)$ $(\rho_i^{(1)}, \mathfrak{s}_i)$ $\leftarrow \text{RunRoundOne}(1^\kappa, 1^n, f, i, x_i; r_i)$ For $i \notin I$ do: $\rho_i^{(2)}$ $\leftarrow \text{RunRoundTwo}(\mathfrak{s}_i, (\rho_1^{(1)}, \dots, \rho_n^{(1)}); r_i)$ $\alpha \leftarrow \mathcal{A}_2(\text{st}, (\rho_i^{(1)}, \rho_i^{(2)})_{i \notin I})$ Return α</p>	<p>(Optional) $\text{crs} \leftarrow \widetilde{\text{CRSGen}}(1^\lambda)$ $(1^n, I, f, (x_i)_{i \notin I}) \leftarrow \mathcal{A}_0(1^\kappa)$ <i>//</i> $f \in \mathcal{F}_\lambda$ of arity n, $x_i \in \mathcal{M}_\lambda$ $((\rho_i^{(1)})_{i \in I}, \text{st}) \leftarrow \mathcal{A}_1(n, I, f)$ $x_i \leftarrow \text{Extract}(\rho_i^{(1)}) \quad \forall i \in I$ $(\rho_i^{(1)}, \rho_i^{(2)})_{i \notin I}$ $\leftarrow \widetilde{\text{Sim}}((x_i)_{i \in I}, f(x_1, \dots, x_n))$ $\alpha \leftarrow \mathcal{A}_2(\text{st}, (\rho_i^{(1)}, \rho_i^{(2)})_{i \notin I})$ Return α</p>

Figure A.3: Experiments for SIM-security of two-round MPC.

honest parties are independent of the inputs, so for convenience and ease of presentation we will partition the algorithm $\widetilde{\text{Sim}}$ into two algorithms: $\widetilde{\text{Sim}}_1$ and $\widetilde{\text{Sim}}_2$, defined as follows:

- $\widetilde{\text{Sim}}_1() \mapsto (\rho_i^{(1)})_{i \notin I}$: Outputs the first-round protocol messages for the honest parties.
- $\widetilde{\text{Sim}}_2((x_i)_{i \in I}, y) \mapsto (\rho_i^{(2)})_{i \notin I}$: On input the inputs of the corrupted parties along with the target output value of the protocol, $\widetilde{\text{Sim}}_2$ outputs the second-round protocol messages for the honest parties.

Input/Function-Rerunnability. For simplicity, the definition in Figure A.3 does not capture input/function-rerunnability. It is straightforward to see how the definition can be extended. For example in the case of function-rerunnability (the situation is analogous for input-rerunnability where inputs and functions are swapped), the changes to the definition are (1) \mathcal{A}_0 outputs a set of functions $\{f_i\}$ instead of a single function, (2) in the real experiment RunRoundTwo is executed for each function, (3) in the ideal experiment $\widetilde{\text{Sim}}_2$ is called for each function, and (4) the complete set of second-round protocol messages for all functions is given to \mathcal{A}_2 .

A.2 Constructions

Local Setup. In multi-party FE schemes, local setup is highly desirable, since it overcomes the key escrow problem which is one of the main drawbacks of present day FE constructions. In this setting, each user i runs the setup algorithm independently, and obtains her own public key PK_i and master secret key MSK_i . No communication or co-ordination is required between the users. In the symmetric key setting, the user encrypts her input using her master key MSK_i , while in the public key setting, anyone may encrypt using the public key(s) PK_i .

Constructions of multi-party FE with local setup have been notoriously hard to build. For general circuits and in the symmetric key setting, the primitive of adhoc multi-input functional encryption (aMIFE) recently proposed by Agrawal et al. [ACF⁺20] comes close to a general feasibility result in our model. Adhoc MIFE enables multiple parties to run local, independent setup algorithms, and generate their own private master keys as well as corresponding public keys. To encrypt data, each party (say i) uses its private master key to compute a ciphertext for any message of its choice (say x_i), and to issue function specific keys, the party uses the same master key to compute a partial function key corresponding to some function (say f). These ciphertexts and partial keys are sent to the decryptor who can aggregate them and compute $f(x_1, \dots, x_n)$ as long as all the partial decryption keys correspond to the same function f . The authors provide a construction of aMIFE for circuits by using the standard notion of MIFE, a single input FE and a special purpose two round MPC.

There are some important differences in the formulation of aMIFE and our MPFE:

1. In aMIFE, users each provide partial keys which may be combined when these keys refer to the *same* function. In contrast, our notion allows different partial keys to refer to different functions which may be combined in diverse ways as codified by the function Agg_y .
2. The encryptors and key authorities are the same in aMIFE, and have the same master key MSK. This is not general enough to capture even regular single user FE, where the key authority and encryptor are different entities that may not share the master key, for instance in the public key setting. To capture these notions, we defined our framework to allow for users and key authorities to have separate encryption and master keys, where the encryption keys may be public.

Below, we provide an MPFE construction using aMIFE as a black-box – this transformation works in the symmetric key setting. The more general construction supporting the public key setting requires an interactive setup, and is discussed next. The construction for symmetric key MPFE with local setup using aMIFE is as follows:

1. Setup in MPFE is run with `mode = Local`. Every party in the MPFE protocol, whether a key generator or encryptor runs the aMIFE setup and obtains its own public and private key. If the party is a key generator, it sets this as its MSK, if it is an encryptor, it sets it as EK.
2. Next, each party in MPFE, whether encryptor or key generator, computes an aMIFE *ciphertext* for its input x_i or y_j . Thus, we obtain n_x partial ciphertexts for inputs $x_i, i \in [n_x]$ and n_y partial keys for inputs y_j , where $j \in [n_y]$.
3. Each party, whether encryptor or key generator, also computes an aMIFE partial key for the function $\mathcal{U}(\text{Agg}_x(\cdot), \text{Agg}_y(\cdot))$.
4. The decryptor/aggregator receives the partial keys and ciphertexts and performs aMIFE decryption to compute $\mathcal{U}(\text{Agg}_x(x_1, \dots, x_{n_x}), \text{Agg}_y(y_1, \dots, y_{n_y}))$.

Correctness and security follow by those of aMIFE. The decryptor obtains aMIFE ciphertexts for x_i where $i \in [n_x]$ and y_j where $j \in [n_y]$ along with $n_x + n_y$ aMIFE partial keys for the function $\mathcal{U}(\text{Agg}_x(\cdot), \text{Agg}_y(\cdot))$. This enables the decryptor to recover $\mathcal{U}(\text{Agg}_x(x_1, \dots, x_{n_x}), \text{Agg}_y(y_1, \dots, y_{n_y}))$ as desired. Security follows directly from the security of aMIFE.

Interactive Setup. A standard MIFE scheme can be modified to support distributed keys by replacing the setup algorithm with an interactive protocol. This makes use of a function delayed, rerunnable two round MPC protocol similar to the construction of aMIFE provided by [ACF+20]. Note that the limitation of interactive setup is mitigated by the fact that it must only be run once.

We provide a general feasibility result for the case of interactive setup below.

1. The MPFE setup protocol does the following:
 - (a) Invoke $(\text{MIFE.MSK}, \text{MIFE.EK}_1, \dots, \text{MIFE.EK}_{n_x+n_y}) \leftarrow \text{MIFE.Setup}(1^\lambda, 1^{n_x+n_y})$.
 - (b) Using an N out of N secret sharing scheme, compute shares MIFE.MSK_i of the MIFE.MSK and output MIFE.MSK_i to each key authority i for $i \in [n_y]$.
 - (c) Using MIFE.MSK_i as input, run the first round of function-delayed, rerunnable two round MPC protocol as $\text{MPC.RunRoundOne}(1^\lambda, 1^{n_y}, i, \text{MIFE.MSK}_i)$ to obtain first protocol message $\rho_i^{(1)}$ and secret \mathfrak{s}_i for each party $i \in [n_y]$.
 - (d) For each key authority $i \in [n_y]$, output $\text{PK}_i = \rho_i^{(1)}$ as the public key and $\text{MSK}_i = (\text{MIFE.EK}_{n_x+i}, \text{MIFE.MSK}_i, \mathfrak{s}_i)$ as the master secret key.
 - (e) For each encryptor $j \in [n_x]$, output the encryption key $\text{EK}_j = \text{MIFE.EK}_j$ and $\text{PK}_j = \text{MIFE.PK}$.
2. To encrypt, encryptor j for $j \in [n_x]$ computes a ciphertext $\text{MIFE.CT}_j = \text{MIFE.Encrypt}(\text{EK}_j, x_j)$ for any message x_j of its choice.
3. To generate a partial key for a given function f_i , the i^{th} key authority does the following:
 - (a) Compute a ciphertext $\text{MIFE.CT}_{n_x+i} = \text{MIFE.Encrypt}(\text{EK}_{n_x+i}, f_i)$.
 - (b) Define a function \mathcal{U}' , which upon inputs $\text{MIFE.MSK}_1, \dots, \text{MIFE.MSK}_{n_y}$ does the following: i) it computes MIFE.MSK using share reconstruction. ii) it computes $\text{MIFE.SK} = \text{MIFE.KeyGen}(\text{MIFE.MSK}, \mathcal{U}(\text{Agg}_x(\cdot), \text{Agg}_y(\cdot)))$.
 - (c) Compute $\text{MIFE.SK}_{i,\mathcal{U}'} = \text{MPC.RunRoundTwo}(\mathcal{U}', \mathfrak{s}_i, (\rho_1^{(1)}, \dots, \rho_n^{(1)}))$ for $i \in [n_y]$ and outputs it¹⁹.
4. To decrypt, the decryptor computes MPC.ComputeResult to obtain $\mathcal{U}'(\text{MIFE.MSK}_1, \dots, \text{MIFE.MSK}_{n_y}) = \text{MIFE.SK}_{\mathcal{U}'}$. It computes $\text{MIFE.Decrypt}(\text{MIFE.CT}_1, \dots, \text{MIFE.CT}_{n_x+n_y}, \text{MIFE.SK}_{\mathcal{U}'})$ and outputs it.

This construction is simpler than that of aMIFE since we permit interactive setup and allow the arity of the function being computed to be fixed. Note that if the MIFE is public key, then so is the above MPFE, since the symmetric/public key property of the above MPFE is inherited from the underlying MIFE. By correctness of MPC, the decryptor recovers MIFE.SK for the function $\mathcal{U}(\text{Agg}_x(\cdot), \text{Agg}_y(\cdot))$. The decryptor also obtains ciphertexts $\text{MIFE.CT}(x_1), \dots, \text{MIFE.CT}(x_{n_x})$ and $\text{MIFE.CT}(f_1), \dots, \text{MIFE.CT}(f_{n_y})$. By correctness of MIFE, it therefore obtains $\mathcal{U}(\text{Agg}_x(x_1, \dots, x_{n_x}), \text{Agg}_y(y_1, \dots, y_{n_y}))$ as desired. Security follows as a special case of aMIFE security, by leveraging security of MPC and MIFE.

¹⁹We note that all authorities need to agree on the choice of Agg_x and Agg_y for correctness of this step to hold. Also note that the choice of Agg_x and Agg_y may be made during the partial key generation step, and is not required at setup.