

Encrypt-Augment-Recover: Computationally Function Private Predicate Encryption in the Public-Key Setting

Sikhar Patranabis and Debdeep Mukhopadhyay

Department of Computer Science and Engineering
Indian Institute of Technology Kharagpur
sikhar.patranabis@iitkgp.ac.in, debdeep@cse.iitkgp.ernet.in

Abstract. We solve the open problem of constructing *computationally function private* public-key predicate encryption schemes, supporting a rich class of predicates, without relying on strong assumptions such as indistinguishability obfuscation. Our framework for computational function privacy is indistinguishability-based. In other words, our notion requires that a secret-key corresponding to a predicate sampled from a distribution with min-entropy super logarithmic in the security parameter λ , is *computationally indistinguishable* from another secret-key corresponding to a uniformly and independently sampled predicate. Existing public-key constructions for predicate encryption either satisfy a *statistical* notion of function privacy, that was introduced for equality predicates by Boneh, Raghunathan and Segev in CRYPTO'13, and was generalized for subspace-membership predicates in ASIACRYPT'13, or use strong assumptions such as the existence of quasi-strong indistinguishability obfuscation (Iovino et al. in CANS'15). Computational function privacy for predicate encryption, without the need for strong additional assumptions, has only been concretely realized in the private-key setting, to the best of our knowledge.

In this paper, we develop a novel approach, denoted as *encrypt-augment-recover*, that takes an existing predicate encryption scheme and transforms it into a computationally function private one while retaining its original data privacy guarantees. Our approach leads to public-key constructions for identity-based encryption and inner-product encryption that are fully data private and computationally function private under a family of weaker variants of the DLIN assumption. Our constructions, in fact, satisfy an *enhanced* notion of function privacy, requiring that an adversary learns nothing more than the minimum necessary from a secret-key, even given corresponding ciphertexts with attributes that allow successful decryption.

Keywords: Predicate Encryption, Public-Key, Function Privacy, Computational Indistinguishability, Min-Entropy, Identity-Based Encryption, Inner-Product Encryption

1 Introduction

Predicate encryption systems [1–3] in the public-key setting allow a single public-key to be associated with multiple secret-keys, where each secret-key corresponds to a boolean predicate $f : \Sigma \rightarrow \{0, 1\}$ over a pre-defined set of attributes Σ . A plaintext message in a predicate encryption system is an attribute-payload message pair $(I, M) \in \Sigma \times \mathcal{M}$, with \mathcal{M} being the payload message space. A secret-key sk_f associated with a predicate f successfully decrypts a ciphertext C corresponding to a plaintext (I, M) and recovers the payload message M if and only if $f(I) = 1$. On the other hand, if $f(I) = 0$, attempting to decrypt C using sk_f returns the failure symbol \perp . A predicate encryption is said to be *attribute hiding* if the ciphertext C leaks no information about the underlying plaintext (I, M) to an adversary possessing benign secret-keys corresponding to predicates that do not trivially identify the attribute I .

Identity-Based Encryption. Identity-based encryption (IBE) [4–6] is the simplest sub-class of public-key predicate encryption. IBE supports a set of equality predicates of the form $f_{\text{id}} : \Sigma \rightarrow \{0, 1\}$ defined as $f_{\text{id}}(x) = 1$ if and only if $x = \text{id}$. The attribute space in this case is a set of identities \mathcal{ID} , and each identity $\text{id} \in \mathcal{ID}$ is associated with its own secret-key sk_{id} .

Inner-Product Encryption. Inner-product encryption (IPE) [2, 3, 7, 8] is the most expressive sub-class of predicate encryption, supporting a set of predicates $f_{\vec{v}} : \Sigma \rightarrow \{0, 1\}$ over a vector space of attributes $\Sigma = \mathbb{F}_q^n$ (q being a λ -bit prime). Of particular interest is a specific form of IPE called zero-IPE [3] where for $\vec{v}, \vec{x} \in \Sigma$, we have $f_{\vec{v}}(\vec{x}) = 1$ if and only if $\langle \vec{v}, \vec{x} \rangle = 0$ (here $\langle \vec{v}, \vec{x} \rangle$ denotes the standard inner-product of two vectors \vec{v} and \vec{x}). IPE is powerful enough to encompass IBE and many other predicate encryption systems [3].

Searchable Encryption and Function Privacy. Predicate encryption provides a generic framework for searchable encryption supporting a wide range of query predicates including conjunctive, disjunctive, range and subset queries [9, 1–3]. For instance, a predicate encryption system can be used to realize a mail gateway that follows some special instructions to route encrypted mails based on their header information (e.g. if the mail is from the boss and needs to be treated as urgent). The mail gateway is given the secret-key corresponding to the predicate *is-urgent*, the mail header serves as the attribute, while the routing instructions can be used as the payload message. Another application could be a payment gateway that flags encrypted payments if they correspond to amounts beyond some pre-defined threshold X . The payment gateway is given the secret-key corresponding to the predicate *greater-than- X* , the payment amount itself serves as the attribute, while the flag signal is encoded as the payload message. The attribute hiding property of the predicate encryption scheme ensures that neither gateway learns any information about the plaintext data from the entire operation.

A natural question now arises: should the gateways in the aforementioned examples be able to learn the underlying predicate from the secret-keys given to them? The answer in most scenarios is *no* - the secret-key sk_f should ideally reveal nothing about the predicate f beyond the absolute minimum. This notion of predicate hiding security is commonly referred to as *function privacy*, and predicate encryption scheme satisfying this notion of security are described as *function private*.

Function Privacy in the Public-Key Setting. As pointed out by Boneh, Raghunathan and Segev in [10, 11], formalizing a realistic notion of function privacy in the context of public-key predicate encryption is, in general, not straightforward. Consider, for example, an adversary against an IBE scheme who is given a secret-key sk_{id} corresponding to an identity id and has access to an encryption oracle. As long as the adversary has some apriori information that the identity id belongs to a small set \mathcal{S} , (e.g. id is sampled distribution with min-entropy at most polynomial in the security parameter λ), it can fully recover id from sk_{id} : it can simply resort to encrypting a random message M under each identity in \mathcal{S} , and decrypting using sk_{id} to check for a correct recovery. Consequently, [10, 11] consider a framework for function privacy under the minimal assumption that any predicate is sampled from a distribution with min-entropy at least super logarithmic in the security parameter λ .

Statistically Function Private Constructions. A class of existing public-key constructions for predicate encryption satisfy a *statistical* notion of function privacy, that was introduced for equality predicates in [10], and was subsequently generalized for subspace-membership predicates in [11]. In particular, the secret-keys in these constructions are *statistically indistinguishable* from random, as long the underlying predicates are sampled from sufficiently unpredictable distributions. In fact, designing public-key predicate encryption schemes whose function privacy can be based on well-known computational assumptions, was left as an open problem in [11].

Function Privacy from Indistinguishability Obfuscation. An alternative approach to achieving function privacy for more generalized circuits such as NC^1 circuits was proposed by Iovino et al. in [12]. Their constructions achieve function privacy assuming the existence of quasi-strong indistinguishability obfuscation (or quasi-siO in short), which are a weakening of the notion of siO originally proposed by Agrawal et al. in [13]. Their constructions can be used to generate functionally anonymous signatures, as well as black-box constructions for functional encryption schemes supporting randomized functionalities. In this paper, we explore the possibility of achieving predicate encryption schemes that can be proven

to be function private under well-known computational assumptions, as opposed to strong assumptions such as quasi-siO. While such schemes have been concretely realized in the private-key setting [14–17], there exist no such constructions in the public-key setting to the best of our knowledge.

1.1 Our Contributions

In this paper, we solve the open problem of constructing computationally function private public-key predicate encryption schemes, supporting a rich class of predicates, without relying on strong assumptions such as indistinguishability obfuscation. Our framework for computational function privacy, presented formally in Section 3, is indistinguishability-based. In other words, our notion requires that a secret-key corresponding to a predicate sampled from a distribution with min-entropy super logarithmic in the security parameter λ , is *computationally indistinguishable* from another secret-key corresponding to a uniformly and independently sampled predicate. Within this framework, we develop a novel approach, denoted as *encrypt-augment-recover*, that takes an existing predicate encryption scheme and transforms it into a computationally function private one. Our approach leads to the following constructions:

- In the random-oracle model, we present a family of identity-based encryption schemes from bilinear pairings based on the scheme of Boneh and Franklin [4]. Our schemes retain the adaptive data privacy of the underlying scheme (based on the same complexity assumption), and are computationally function private under progressively weaker variants of the well-known DLIN assumption. The detailed construction of these schemes, along with the proofs of data and function privacy, is presented in Section 4.
- In the standard model, we present a family of inner-product encryption schemes from bilinear pairings based on the scheme of Katz, Sahai and Waters [3]. Once again, our schemes retain the selectively attribute hiding property of the underlying scheme, and are computationally function private under progressively weaker variants of the DLIN assumption. The detailed construction of these schemes, along with the proofs of data and function privacy, is presented in Section 5.
- Our constructions, in fact, satisfy an *enhanced* notion of function privacy, requiring that an adversary learns nothing more than the minimum necessary from a secret-key, even given corresponding ciphertexts with attributes that allow successful decryption.

1.2 Overview of Our Approach: *Encrypt-Augment-Remove*

Our approach for achieving computationally function private predicate encryption schemes consists of three main steps - *encrypt*, *augment* and *recover*. We briefly describe the main ideas underlying each step, and exemplify them subsequently using a simple IBE scheme. Given a public-key predicate encryption scheme $\Pi = (\text{Setup}, \text{KeyGen}, \text{Enc}, \text{Dec})$, and a CPA-secure public-key encryption algorithm $\text{PKE} = (\text{KeyGen}, \text{Enc}, \text{Dec})$, we create a function private predicate encryption scheme $\Pi' = (\text{Setup}', \text{KeyGen}', \text{Enc}', \text{Dec}')$ as follows:

- The modified setup algorithm Setup' invokes PKE.KeyGen and obtains the key pair (PK, SK) . It also invokes $\Pi.\text{Setup}$ and obtains the public parameters pp , along with master secret-key msk . It outputs the modified public parameter $\text{pp}' = (\text{pp}, g(SK))$ and the modified master secret-key $\text{msk}' = (\text{msk}, PK)$, where g is a suitably chosen one way function.
- On input a predicate f and the augmented master secret-key $\text{msk}' = (\text{msk}, PK)$, the modified key-generation algorithm KeyGen' invokes $\Pi.\text{KeyGen}$ to obtain the original secret-key sk_f . It then outputs an *encrypted* secret-key sk'_f as $\text{PKE.Enc}(PK, \text{sk}_f)$.

This step automatically guarantees computational function privacy - any adversary that can distinguish the augmented secret-key from random must break the CPA security guarantee of the PKE

scheme. More specifically, it ensures *adaptive* function privacy - the inherently random nature of the augmented key generation algorithm ensures that the function privacy guarantees hold even when the adversary is allowed to specify predicate distributions in an adaptive manner after seeing the public parameters of the scheme. We assume that any adversarially-chosen distribution of predicates is sufficiently unpredictable, so as to rule out a trivial breach of function privacy as mentioned earlier. This minimal assumption is thus sufficient to transform the original predicate encryption scheme into a computationally function private one.

- An even greater challenge is to synchronize the encryption and decryption algorithms in the modified scheme. This is achieved as follows. On input the public parameter $\text{pp}' = (\text{pp}, g(SK))$, and a message M corresponding to an attribute I , the modified encryption algorithm Enc' first obtains $C = \Pi.\text{Enc}(\text{pp}, I, M)$. It then outputs the *augmented* ciphertext $C' = (C, \sigma(g(SK)))$, where σ is a randomized function sharing a source of randomness with $\Pi.\text{Enc}$.
- Finally, Dec' cleverly uses the additional ciphertext component $\sigma(g(SK))$ in C' to *remove* the effect of PKE from the encrypted secret-key sk'_f , and *recovers* the message M . Note that *removal* here is not same as decryption, since Dec' has access to only a one-way function of SK and not SK itself. It is, in fact, impossible to provide SK to Dec' in the clear without trivially compromising function privacy. The challenge is thus to ensure that Dec' can recover M without a complete decryption of sk'_f .

An Example of Our Approach. We present an example of a computationally function private IBE scheme in the random-oracle model achieved using our *encrypt-augment-decrypt* approach. A generalization of this scheme is presented in greater detail in Section 4, along with proofs for data and function privacy. Consider a public-key encryption scheme PKE with the key generation, encryption and decryption algorithms as described below:

- **KeyGen:** The key-generation algorithm samples $x_1, x_2, x_3 \xleftarrow{R} \mathbb{Z}_q^*$, where q is a λ -bit prime, and $g_1, g_2, g_3 \xleftarrow{R} \mathbb{G}$, where \mathbb{G} is a cyclic group of prime order q . It outputs the secret-key $SK = (x_1, x_2, x_3)$ and the public-key $PK = (g_1, g_2, g_3, (g_1^{x_1} \cdot g_3^{x_3}), (g_2^{x_2} \cdot g_3^{x_3}))$.
- **Enc:** The ciphertext C corresponding to a message $M \in \mathbb{G}$ is a tuple of the form:

$$C = (g_1^{y_1}, g_2^{y_2}, g_3^{y_1+y_2}, (g_1^{x_1} \cdot g_3^{x_3})^{y_1} \cdot (g_2^{x_2} \cdot g_3^{x_3})^{y_2} \cdot M)$$

where $y_1, y_2 \xleftarrow{R} \mathbb{Z}_q^*$.

- **Dec:** The decryption algorithm, on input the ciphertext $C = (c_1, c_2, c_3, c_4)$ and the secret-key (x_1, x_2, x_3) , recovers the message M as:

$$M = c_4 / (c_1^{x_1} \cdot c_2^{x_2} \cdot c_3^{x_3})$$

The above scheme is a simple variant of the Cramer-Shoup cryptosystem [18], and is CPA-secure under the DLIN assumption. We now present a computationally function private IBE scheme that is obtained by applying our *encrypt-augment-recover* approach to the anonymous IBE scheme based on bilinear maps proposed by Boneh and Franklin [4].

- **Setup:** The setup algorithm in the scheme of Boneh and Franklin samples $s \xleftarrow{R} \mathbb{Z}_q^*$, where q is a λ -bit prime. The public parameters are g and g^s , where g is the generator of a bilinear group \mathbb{G} of prime order q , while the master secret-key is s . Our scheme additionally samples $x_1, x_2, x_3 \xleftarrow{R} \mathbb{Z}_q^*$ and

$g_1, g_2, g_3 \xleftarrow{R} \mathbb{G}$. The augmented public parameter \mathbf{pp} and master secret-key \mathbf{msk} for our scheme are as follows:

$$\begin{aligned} \mathbf{pp} &= (g, g^s, g^{x_1}, g^{x_2}, g^{x_3}) \\ \mathbf{msk} &= (s, g_1, g_2, g_3, (g_1^{x_1} \cdot g_3^{x_3}), (g_2^{x_2} \cdot g_3^{x_3})) \end{aligned}$$

Observe that the additional components in \mathbf{pp} are one-way functions of x_1, x_2, x_3 - the secret-key SK of the PKE scheme. Additionally, the modified \mathbf{msk} contains the public-key PK of the PKE scheme. This is exactly in accordance with our proposed approach.

- **KeyGen:** The key-generation algorithm in the scheme of Boneh and Franklin computes a secret-key for an identity id as $\text{sk}_{\text{id}} = (H(\text{id}))^s$, where H is a random oracle mapping identities onto the group \mathbb{G} . In our scheme, we augment the key generation process as follows. We sample $y_1, y_2 \xleftarrow{R} \mathbb{Z}_q^*$, and output:

$$\text{sk}_{\text{id}} = (g_1^{y_1}, g_2^{y_2}, g_3^{y_1+y_2}, (g_1^{x_1} \cdot g_3^{x_3})^{y_1} \cdot (g_2^{x_2} \cdot g_3^{x_3})^{y_2} \cdot (H(\text{id}))^s)$$

Observe that $\text{sk}_{\text{id}} = \text{PKE.Enc}(PK, (H(\text{id}))^s)$, which is a direct exemplification of the *encrypt* step in our approach described above for achieving function privacy (the reader is referred to Section 4 for the detailed proof).

- **Enc:** An encryption of a message M for an identity id in the scheme of Boneh and Franklin is a tuple of the form $(g^r, M \cdot e(H(\text{id}), g^s)^r)$, where $r \xleftarrow{R} \mathbb{Z}_q^*$. In our scheme, we augment the encryption process to produce the ciphertext:

$$C = (g^r, (g^{x_1})^r, (g^{x_2})^r, (g^{x_3})^r, M \cdot e(H(\text{id}), g^s)^r)$$

Note that the augmented ciphertext in our scheme retains unaltered the original ciphertext. The main technical challenge is to prove that such an augmented ciphertext still provides the same data privacy guarantees as the original scheme of Boneh and Franklin (the reader is referred to Section 4 for the proof).

- **Dec:** Our decryption algorithm, on input of a ciphertext $C = (c_0, c_1, c_2, c_3, c_4)$, and a secret-key $\text{sk}_{\text{id}} = (d_0, d_1, d_2, d_3)$, recovers the encrypted message M as:

$$M = c_4 \cdot \frac{e(d_0, c_1) \cdot e(d_1, c_2) \cdot e(d_2, c_3)}{e(d_3, c_0)}$$

Observe that at the core of the above computation is the original decryption procedure in the scheme of Boneh and Franklin, with the additional components in the ciphertext and the secret-key canceling out each other to *recover the effect of the PKE* (the reader is referred to Section 4 for the detailed proof of correctness). It is important to note that this removal is different from directly decrypting sk_{id} , and in particular, does not require the knowledge of the secret-key of the PKE.

Comparison with a deterministic public-key encryption-based approach. An alternative approach for designing computationally function private identity-based encryption schemes, suggested in [10], is as follows: encrypt all identities using a DPKE scheme, and use any existing anonymous IBE scheme that treats the corresponding ciphertexts of the DPKE scheme as its identities. Since the security of any deterministic public-key encryption (DPKE) algorithm is also based on the minimal assumption that its plaintexts are sampled from a distribution with a certain amount of min-entropy, the above approach seems quite natural. However, this approach suffers from two inherent drawbacks:

1. In the setting of DPKE, the dependency of plaintexts on the public-key of the scheme is essentially limited. Intuitively, the reason is as follows: in a deterministic encryption setting, plaintext distributions can be chosen depending on the public-key such that the encryption algorithm acts as a subliminal

channel for leaking information, thus trivially violating all security guarantees [19]. However, this is too restrictive a security notion to be adopted in the context of IBE, where the key-generation process is allowed to be randomized. In particular, any realistic function privacy framework for IBE must allow adversaries to *adaptively* specify challenge identity distributions, after they have seen the public parameters of the scheme.

2. The DPKE-based approach is not directly generalizable for functionally richer predicates unless the underlying encryption scheme is somewhat function-preserving. For example, if the above approach were to be extended in the context of IPE, it would require the ciphertexts of the DPKE scheme to preserve the orthogonality of the underlying plaintext vectors. This could potentially weaken the function privacy guarantees even further.

Our approach overcomes these limitations by focusing on encrypting the secret-key sk_f of the original predicate encryption scheme rather than the underlying predicate f . The augmented key-generation process uses a PKE instead of a DPKE, and is hence allowed to be non-deterministic. This makes our notion of function privacy more realistic since it does not restrict adaptive choice of predicate distributions on part of the adversary. Finally, our approach is generally applicable to a large class of predicate encryption schemes, including IBE and IPE, without any additional constraints on the properties of the underlying PKE.

1.3 Other Related Work

Computational function privacy for predicate encryption has been studied in the private-key setting [20]. The inherent difficulty of achieving function privacy in the public-key setting does not apply to the private-key setting, where the encryptor and decryptor have a shared secret-key. In this setting, an adversary with access to a searching key cannot test the same on ciphertexts of its choice since it does not have access to the secret-key. Function privacy in the private-key setting is thus more natural to achieve. A general solution in this direction was proposed by Goldreich and Ostrovsky [21] in their construction of an oblivious RAM. More efficient constructions have been subsequently proposed for equality testing [22, 23, 14, 24, 25] and, more recently, for inner product testing [15–17].

The first meaningful notion of predicate privacy in the public-key setting was put forth by Boneh, Raghunathan and Segev for equality predicates [10], and subsequently for more general subspace-membership predicates [11]. Prior to their work, achieving predicate privacy in the public-key setting was considered impossible. Their notion of predicate privacy is based on the statistical indistinguishability of secret-keys from random, subject to the condition that the predicates are sampled from sufficiently unpredictable distributions. Prior to our work, no public-key constructions satisfying computational predicate privacy have been proposed to the best of our knowledge.

1.4 Paper Organization

The remainder of this paper is organized as follows. Section 2 presents background material on predicate encryption, and introduces several computational assumptions in bilinear groups. In Section 3, we formally define our framework for the computational function privacy of public-key predicate encryption. In Section 4, we present a family of adaptively data private and computationally function private IBE schemes in the random-oracle model. In Section 5, we present a family of selectively attribute hiding and computationally function private IPE schemes in the standard model. Finally, Section 6 concludes the paper and enumerates several open problems.

1.5 Notations Used

We write $x \xleftarrow{R} \mathcal{X}$ to represent that an element x is sampled uniformly at random from a set \mathcal{X} . The output a of a deterministic algorithm \mathcal{A} is denoted by $x \leftarrow \mathcal{A}$ and the output a' of a randomized algorithm \mathcal{A}' is denoted by $x' \xleftarrow{R} \mathcal{A}'$. We refer to $\lambda \in \mathbb{N}$ as the security parameter, and denote by $\exp(\lambda)$, $\text{poly}(\lambda)$ and

$\text{negl}(\lambda)$ any generic (unspecified) exponential function, polynomial function and negligible function in λ respectively. Note that a function $f : \mathbb{N} \rightarrow \mathbb{N}$ is said to be negligible in λ if for every positive polynomial p , $f(\lambda) < 1/p(\lambda)$ when λ is sufficiently large. Finally, for $a, b \in \mathbb{Z}$ such that $a \leq b$, we denote by $[a, b]$ the set of integers lying between a and b (both inclusive).

The min-entropy of a random variable Y is denoted as $\mathbf{H}_\infty(Y) = -\log(\max_y \Pr[Y = y])$; a random variable Y is said to be a k -source if $\mathbf{H}_\infty(Y) \geq k$. A (T, k) -block-source is a random variable $\mathbf{Y} = (Y_1, \dots, Y_T)$ where for each $i \in [1, T]$ and y_1, \dots, y_{i-1} , it holds that:

$$\mathbf{H}_\infty(Y_i | Y_1 = y_1, \dots, Y_{i-1} = y_{i-1}) \geq k$$

2 Preliminaries

2.1 Public-key Predicate Encryption

A public-key predicate encryption scheme for a class of predicates \mathcal{F} over an attribute space Σ and a payload-message space \mathcal{M} is a quadruple $\Pi = (\text{Setup}, \text{KeyGen}, \text{Enc}, \text{Dec})$ of probabilistic polynomial time algorithms. The **Setup** algorithm takes as input the security parameter λ , and generates the public parameter pp and the master secret-key msk for the system. The key-generation algorithm, **KeyGen** takes as input the master secret-key msk and a predicate $f \in \mathcal{F}$, and generates a secret-key sk_f corresponding to f . The **Enc** algorithm takes as input the public parameter pp , an attribute $I \in \Sigma$ and a payload-message $M \in \mathcal{M}$, and outputs the ciphertext $C = \text{Enc}(\text{pp}, I, M)$. The **Dec** algorithm takes as input the public parameter pp , a ciphertext C and a secret-key sk_f , and outputs either a payload-message $M \in \mathcal{M}$ or the symbol \perp .

Functional Correctness. A predicate encryption scheme $\Pi = (\text{Setup}, \text{KeyGen}, \text{Enc}, \text{Dec})$ is said to be functionally correct if for any security parameter λ , for any predicate $f \in \mathcal{F}$, for any attribute $I \in \Sigma$ and any payload-message $M \in \mathcal{M}$, the following hold with probability at least $1 - \text{negl}(\lambda)$:

1. If $f(I) = 1$, we have $\text{Dec}(\text{pp}, \text{Enc}(\text{pp}, I, M), \text{KeyGen}(\text{msk}, f)) = M$.
2. If $f(I) = 0$, we have $\text{Dec}(\text{pp}, \text{Enc}(\text{pp}, I, M), \text{KeyGen}(\text{msk}, f)) = \perp$.

where the probability is taken over the internal randomness of the algorithms **Setup**, **KeyGen**, **Enc**, and **Dec**.

Data Privacy. We briefly recall the notion of indistinguishability-based data privacy for a predicate encryption scheme under an *adaptive* chosen-attribute chosen-payload-message attack. Data privacy of a functional encryption scheme guarantees that any probabilistic polynomial-time adversary can gain no information about either the attribute I nor the payload-message M associated with a ciphertext C from the knowledge of the public parameters pp . We denote this notion of security by DP throughout the rest of the paper.

Definition 2.1 (Adaptively Data Private Predicate Encryption). A predicate encryption scheme $\Pi = (\text{Setup}, \text{KeyGen}, \text{Enc}, \text{Dec})$ is said to be *adaptively data private* if for any probabilistic polynomial-time adversary \mathcal{A} , the following holds:

$$\text{Adv}_{\Pi, \mathcal{A}}^{\text{DP}}(\lambda) \stackrel{\text{def}}{=} \left| \Pr \left[\text{Expt}_{\text{DP}, \Pi, \mathcal{A}}^{(0)}(\lambda) = 1 \right] - \Pr \left[\text{Expt}_{\text{DP}, \Pi, \mathcal{A}}^{(1)}(\lambda) = 1 \right] \right| \leq \text{negl}(\lambda)$$

where for each $\lambda \in \mathbb{N}$ and each $b \in \{0, 1\}$, the experiment $\text{Expt}_{\text{DP}, \Pi, \mathcal{A}}^{(b)}(\lambda)$ is defined as follows:

1. $(\text{pp}, \text{msk}) \xleftarrow{R} \text{Setup}(1^\lambda)$.

2. $((I_0^*, M_0^*), (I_1^*, M_1^*), \text{state}) \xleftarrow{R} \mathcal{A}^{\text{KeyGen}(\text{msk}, \cdot)}(\text{state})$, where $I_0^*, I_1^* \in \Sigma$ and $M_0^*, M_1^* \in \mathcal{M}$, subject to the restriction that for each predicate f_i with which \mathcal{A} queries $\text{KeyGen}(\text{msk}, \cdot)$, we have $f_i(I_0^*) = f_i(I_1^*)$.
3. $C^* \xleftarrow{R} \text{Enc}(\text{pp}, I_b^*, M_b^*)$.
4. $b' \xleftarrow{R} \mathcal{A}^{\text{KeyGen}(\text{msk}, \cdot)}(C^*, \text{state})$, once again subject to the restriction that for each predicate f_i with which \mathcal{A} queries $\text{KeyGen}(\text{msk}, \cdot)$, we have $f_i(I_0^*) = f_i(I_1^*)$.
5. Output b' .

We also consider a *selective* variant of the above security notion that requires the adversary to commit to the challenge pair of attributes before seeing the public parameters of the scheme. We denote this notion of security by **sDP** throughout the rest of the paper.

Definition 2.2 (Selectively Data Private Predicate Encryption). A predicate encryption scheme $\Pi = (\text{Setup}, \text{KeyGen}, \text{Enc}, \text{Dec})$ is said to be *selectively data private* if for any probabilistic polynomial-time adversary \mathcal{A} , the following holds:

$$\text{Adv}_{\Pi, \mathcal{A}}^{\text{sDP}}(\lambda) \stackrel{\text{def}}{=} \left| \Pr \left[\text{Expt}_{\text{sDP}, \Pi, \mathcal{A}}^{(0)}(\lambda) = 1 \right] - \Pr \left[\text{Expt}_{\text{sDP}, \Pi, \mathcal{A}}^{(1)}(\lambda) = 1 \right] \right| \leq \text{negl}(\lambda)$$

where for each $\lambda \in \mathbb{N}$ and each $b \in \{0, 1\}$, the experiment $\text{Expt}_{\text{sDP}, \Pi, \mathcal{A}}^{(b)}(\lambda)$ is defined as follows:

1. $(I_0^*, I_1^*, \text{state}) \xleftarrow{R} \mathcal{A}(1^\lambda)$, where $I_0^*, I_1^* \in \Sigma$.
2. $(\text{pp}, \text{msk}) \xleftarrow{R} \text{Setup}(1^\lambda)$.
3. $(M_0^*, M_1^*, \text{state}) \xleftarrow{R} \mathcal{A}^{\text{KeyGen}(\text{msk}, \cdot)}(\text{state})$, where $M_0^*, M_1^* \in \mathcal{M}$, subject to the restriction that for each predicate f_i with which \mathcal{A} queries $\text{KeyGen}(\text{msk}, \cdot)$, we have $f_i(I_0^*) = f_i(I_1^*)$.
4. $C^* \xleftarrow{R} \text{Enc}(\text{pp}, I_b^*, M_b^*)$.
5. $b' \xleftarrow{R} \mathcal{A}^{\text{KeyGen}(\text{msk}, \cdot)}(C^*, \text{state})$, once again subject to the restriction that for each predicate f_i with which \mathcal{A} queries $\text{KeyGen}(\text{msk}, \cdot)$, we have $f_i(I_0^*) = f_i(I_1^*)$.
6. Output b' .

Identity-Based Encryption. An identity-based encryption scheme Π^{IBE} over an identity space \mathcal{ID} and a message space \mathcal{M} is a public-key predicate encryption scheme supporting the set of equality predicates $f_{\text{id}} : \mathcal{ID} \rightarrow \{0, 1\}$ defined as $f_{\text{id}}(\text{id}') = 1$ if and only if $\text{id}' = \text{id}$. The secret-key associated with an identity $\text{id} \in \mathcal{ID}$ is denoted as sk_{id} . The notions of anonymity and message indistinguishability security popularly associated with IBE are equivalent to the notion of adaptive data privacy as described above.

Inner-Product Encryption. An inner-product encryption scheme Π^{IPE} over an attribute space $\Sigma = \mathbb{F}_q^n$ (q being a λ -bit prime) and a payload message space \mathcal{M} is a public-key predicate encryption scheme supporting the set of vector predicates $f_{\vec{v}} : \Sigma \rightarrow \{0, 1\}$. The secret-key associated with a vector $\vec{v} \in \Sigma$ is denoted as $\text{sk}_{\vec{v}}$. Zero-IPE is a specific sub-class of IPE where for $\vec{v}, \vec{x} \in \Sigma$, we have $f_{\vec{v}}(\vec{x}) = 1$ if and only if $\langle \vec{v}, \vec{x} \rangle = 0$.

2.2 Computational Assumptions in Bilinear Groups

The decisional bilinear Diffie-Hellman assumption (DBDH). Let $\text{GroupGen}(1^\lambda)$ be a probabilistic polynomial-time algorithm that takes as input a security parameter λ , and outputs the tuple $(\mathbb{G}, \mathbb{G}_T, q, g, e)$, where \mathbb{G} and \mathbb{G}_T are groups of order q (q being a λ -bit prime), g is a generator for \mathbb{G} and $e : \mathbb{G} \times \mathbb{G} \rightarrow \mathbb{G}_T$ is an efficiently computable non-degenerate bilinear map. The group \mathbb{G} is popularly referred to as a *bilinear group* [4]. The decisional bilinear Diffie-Hellman assumption is that the distribution ensembles:

$$\{(g, g^{a_1}, g^{a_2}, g^{a_3}, e(g, g)^{a_1 \cdot a_2 \cdot a_3})\}_{a_1, a_2, a_3 \xleftarrow{R} \mathbb{Z}_q^*} \text{ and } \{(g, g^{a_1}, g^{a_2}, g^{a_3}, Z)\}_{a_1, a_2, a_3 \xleftarrow{R} \mathbb{Z}_q^*, Z \xleftarrow{R} \mathbb{G}_T}$$

are computationally indistinguishable, where $(\mathbb{G}, \mathbb{G}_T, q, g, e) \leftarrow \text{GroupGen}(1^\lambda)$.

The decisional linear assumption (DLIN) [26]. Let \mathbb{G} be a group of prime order q and let g_1, g_2, g_3 be arbitrary generators for \mathbb{G} . The decisional linear assumption is that the distribution ensembles:

$$\left\{ (g_1, g_2, g_3, g_1^{a_1}, g_2^{a_2}, g_3^{a_1+a_2}) \right\}_{a_1, a_2 \xleftarrow{R} \mathbb{Z}_q^*} \quad \text{and} \quad \left\{ (g_1, g_2, g_3, g_1^{a_1}, g_2^{a_2}, g_3^{a_3}) \right\}_{a_1, a_2, a_3 \xleftarrow{R} \mathbb{Z}_q^*}$$

are computationally indistinguishable, where $g_1, g_2, g_3 \xleftarrow{R} \mathbb{G}$.

The DLIN assumption was introduced by Boneh, Boyen, and Shacham [26], and was intended to take the place of the more standard decisional Diffie Hellman (DDH) assumption in groups where the DDH assumption does not hold. In particular, for bilinear groups as defined above, the DLIN assumption holds even if the DDH assumption does not, at least in the generic group model.

The generalized decisional k -linear assumption (k -DLIN) [27]. Let \mathbb{G} be a group of prime order q and let g_1, \dots, g_k, g_{k+1} be arbitrary generators for \mathbb{G} . The generalized decisional k -linear assumption is that the distribution ensembles:

$$\left\{ \left(g_1, \dots, g_k, g_{k+1}, g_1^{a_1}, \dots, g_k^{a_k}, g_{k+1}^{\sum_{j=1}^k a_j} \right) \right\}_{a_1, \dots, a_k \xleftarrow{R} \mathbb{Z}_q^*} \quad \text{and} \\ \left\{ (g_1, \dots, g_k, g_{k+1}, g_1^{a_1}, \dots, g_k^{a_k}, g_{k+1}^{a_{k+1}}) \right\}_{a_1, \dots, a_k, a_{k+1} \xleftarrow{R} \mathbb{Z}_q^*}$$

are computationally indistinguishable, where $g_1, \dots, g_{k+1} \xleftarrow{R} \mathbb{G}$.

Quite evidently, this assumption is a generalization of the DLIN assumption stated above. Note that the k -DLIN assumption implies the $(k+1)$ -DLIN assumption for all $k \geq 1$, but the reverse is not necessarily true, implying that the k -DLIN assumption family is a family of progressively weaker assumptions [27].

3 Computational Function Privacy of Public-Key Predicate Encryption

We present our definitions for the computational function privacy of predicate encryption in the public-key setting. We consider adversaries that have access to the public parameters of the scheme, as well as a secret-key generation oracle. The adversary can also adaptively interact with a *real-or-random* function-privacy oracle RoR^{FP} . This oracle takes as input any adversarially-chosen distribution over the class of predicates \mathcal{F} , and outputs a secret-key either for a predicate sampled from the given distribution, or for an independently and uniformly sampled predicate. At the end of the interaction, the adversary should be able to distinguish between these *real* and *random* modes of operation of RoR^{FP} with only negligible probability.

Formal Definitions. We now formally present the computational function privacy definitions for public-key predicate encryption.

Definition 3.1 (Real-or-Random Function Privacy Oracle). The real-or-random function privacy oracle for predicate encryption RoR^{FP} takes as input triplets of the form $(\text{mode}, \text{msk}, \mathbf{F})$, where $\text{mode} \in \{\text{real}, \text{rand}\}$, msk is the master secret-key, and \mathbf{F} is a circuit representing a distribution over the class of predicates \mathcal{F} . If $\text{mode} = \text{real}$, the oracle samples $f \xleftarrow{R} \mathbf{F}$, while if $\text{mode} = \text{rand}$, it samples $f \xleftarrow{R} \mathcal{F}$. It then computes $\text{sk}_f \xleftarrow{R} \text{KeyGen}(\text{msk}, f)$ and responds with sk_f .

Definition 3.2 (Computational Function Privacy). A predicate encryption scheme $\Pi = (\text{Setup}, \text{KeyGen}, \text{Enc}, \text{Dec})$ is said to be *computationally function private* if for any probabilistic polynomial-time adversary \mathcal{A} , the following holds:

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

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

1. $(\text{pp}, \text{msk}) \xleftarrow{R} \text{Setup}(1^\lambda)$.
2. $b \xleftarrow{R} \mathcal{A}^{\text{RoR}^{\text{FP}}(\text{mode}, \text{msk}, \cdot), \text{KeyGen}(\text{msk}, \cdot)}(1^\lambda, \text{pp})$, subject to the restriction that each \mathbf{F}_i with which \mathcal{A} queries $\text{RoR}^{\text{FP}}(\text{mode}, \text{msk}, \cdot)$ represents a distribution with min-entropy $k = \omega(\log \lambda)$.
3. Output b .

Note that our definitions are generic, and may be suitably adopted for IBE, IPE and other classes of predicate encryption.

Min-Entropy Requirements. In our definitions for computational function privacy, the adversary is allowed to adaptively issue a polynomial number of queries to the RoR^{FP} oracle, as long as the queries correspond to distributions with min-entropy $k = \omega(\log \lambda)$. As discussed in Section 1.2, such a restriction is *necessary* for any definition of function privacy to be meaningful in the public-key setting. In the context of IBE, for example, the adversary is allowed to query the real-or-random oracle with $\mathbf{ID}^* \in \mathcal{ID}$ only if \mathbf{ID}^* represents a k -source such that $k = \omega(\log \lambda)$. In the context of IPE, on the other hand, the adversary can query the real-or-random oracle with $\mathbf{V}^* = (V_1^*, \dots, V_n^*) \in \mathbb{Z}_N^n$ only if \mathbf{V}^* is an (n, k) -block source such that $k = \omega(\log \lambda)$. Additionally, each component-wise distribution V_i^* for $i \in [1, n]$ should be completely uncorrelated with each of the other distributions in \mathbf{V}^* . This restriction is necessary to ensure that the adversary cannot carefully craft vectorial distributions with arbitrary inter-component correlations to trivially compromise function privacy (see [11] for a detailed explanation).

Multi-Shot v/s Single-Shot Adversaries. Definition 3.1 considers *multi-shot* adversaries that are allowed to query the RoR^{FP} oracle polynomially many times. However, it is polynomially equivalent to consider *single-shot* adversaries that can query the RoR^{FP} at most once. This is easily established by a hybrid argument, where the hybrids are constructed such that only one query is forwarded to the RoR^{FP} oracle, while the rest are answered by the key generation oracle.

3.1 Computational Function Privacy of Existing Predicate Encryption Schemes

Identity-Based Encryption. To the best of our knowledge, there exist no IBE schemes in literature that can be proven to be computationally function private under well-known cryptographic assumptions. The constructions in [4] and [28] have deterministic trapdoors, and are hence trivially not function private under Definition 3.1. For such schemes, an adversary could easily manufacture a circuit that uniformly samples id such that some function of the public parameters pp and the secret-key sk_{id} is already known to the adversary, thus leading to a straightforward attack breaking function privacy [10]. While such straightforward attacks cannot be demonstrated on the IBE constructions proposed in [5, 29–31], we are not aware if their function privacy can be based on standard computationally intractable problems. Finally, the IBE constructions presented in [10] are secure under a different notion of function privacy called statistical function privacy, that is based on the statistical closeness of adversarially-chosen and random distributions. Once again, to the best of our knowledge, these constructions are not computationally function private to the best of our knowledge. In Section 4, we present a family of IBE constructions that are computationally function private under well-known cryptographic assumptions.

Inner-Product Encryption. While computational functional privacy with respect to IPE in the private key setting is well-studied [15–17], there exist, to the best of our knowledge, no equivalent public-key counterparts in literature that can be proven to be function private under well-known cryptographic assumptions. The authors of [11] present a generic technique to achieve statistical function privacy in the context of inner-product encryption in the public-key setting; however, they leave the construction of computationally function private IPE schemes in the public-key setting as an open problem. It also seems that the function privacy of existing IPE constructions, such as in [2, 3], cannot directly be based on standard computational assumptions without suitable modifications. In Section 5, we present a family of IPE constructions that are computationally function private under well-known cryptographic assumptions.

3.2 Enhanced Computational Function Privacy for Predicate Encryption

Boneh, Raghunathan and Segev put forth a stronger notion of function privacy in [10] with respect to IBE, referred to as *enhanced* function privacy. Informally, this notion requires that the function privacy guarantees of an IBE scheme hold even if the adversary can obtain, in addition to the secret-key sk_{id} corresponding to an identity id , an encryption of an arbitrary message M under id . We formally extend their notion of enhanced computational function privacy to address a more generalized class of predicates. We assume that the adversary, in addition to interacting with the key generation and real-or-random function privacy oracles, also interacts with a function-privacy encryption oracle Enc^{FP} . This oracle shares a state with the real-or-random function privacy oracle. For a multi-shot adversary, let f_j^* be the predicate sampled by the real-or-random function privacy oracle when responding to the j^{th} query made by the adversary to the real-or-random function privacy oracle. The Enc^{FP} oracle takes as input a payload message M and j , uniformly samples an attribute I_j^* such that $f_j^*(I_j^*) = 1$, and responds with $C_j^* = \text{Enc}(\text{pp}, (I_j^*, M))$. The equivalent definition for a single-shot adversary follows similarly.

Definition 3.3 (Enhanced Computational Function Privacy for Predicate Encryption). A predicate encryption scheme $\Pi = (\text{Setup}, \text{KeyGen}, \text{Enc}, \text{Dec})$ is said to be *computationally enhanced function private* if for any probabilistic polynomial-time adversary \mathcal{A} , the following holds:

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

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

1. $(\text{pp}, \text{msk}) \xleftarrow{R} \text{Setup}(1^\lambda)$.
2. $b \xleftarrow{R} \mathcal{A}^{\text{RoRFP}(\text{mode}, \text{msk}, \cdot), \text{Enc}^{\text{FP}}(\cdot), \text{KeyGen}(\text{msk}, \cdot)}(1^\lambda, \text{pp})$, subject to the restriction that each \mathbf{F}_i with which \mathcal{A} queries $\text{RoRFP}(\text{mode}, \text{msk}, \cdot)$ represents a distribution with min-entropy $k = \omega(\log \lambda)$.
3. Output b .

4 Computationally Function private Identity-Based Encryption

In this section, we apply our *encrypt-augment-recover* approach to the anonymous IBE scheme of Boneh and Franklin [4] to achieve a family of computationally function private IBE schemes $\{\Pi_k^{\text{IBE}}\}_{k \geq 1}$. At the core of Π_k^{IBE} is the following generalized version of the PKE scheme introduced in Section 1.2, which is CPA-secure under the $(k+1)$ -DLIN assumption:

- **KeyGen:** The key-generation algorithm samples $x_1, \dots, x_{k+2} \xleftarrow{R} \mathbb{Z}_q^*$, where q is a λ -bit prime, and $g_1, \dots, g_{k+2} \xleftarrow{R} \mathbb{G}$, where \mathbb{G} is a cyclic group of prime order q . It outputs the secret-key SK and the public-key PK as:

$$SK = (x_1, \dots, x_{k+2}), \quad PK = (g_1, \dots, g_{k+2}, \{(g_j^{x_j} \cdot g_{k+2}^{x_{k+2}})\}_{j \in [1, k+1]})$$

- **Enc:** The ciphertext C corresponding to a message $M \in \mathbb{G}$ is a tuple of the form:

$$C = \left(g_1^{y_1}, \dots, g_{k+1}^{y_{k+1}}, g_{k+2}^{\sum_{j=1}^{k+1} y_j}, \left(\prod_{j=1}^{k+1} (g_j^{x_j} \cdot g_{k+2}^{x_{k+2}})^{y_j} \right) \cdot M \right)$$

where $y_1, \dots, y_{k+1} \xleftarrow{R} \mathbb{Z}_q^*$.

- **Dec:** The decryption algorithm, on input the ciphertext $C = (c_1, \dots, c_{k+3})$ and the secret-key (x_1, \dots, x_{k+2}) , recovers the message M as:

$$M = c_{k+3} / \left(\prod_{j=1}^{k+1} c_j^{x_j} \right)$$

The IBE scheme Π_k^{IBE} , which is built by combining the Boneh-Franklin IBE scheme with the above PKE scheme, is adaptively data private under the DBDH assumption and function private under the $(k+1)$ -DLIN assumption. We present the construction for the same next.

The Construction. Let $\text{GroupGen}(1^\lambda)$ be a probabilistic polynomial-time algorithm that takes as input a security parameter λ , and outputs the tuple $(\mathbb{G}, \mathbb{G}_T, q, g, e)$, where \mathbb{G} and \mathbb{G}_T are groups of prime order $q = \mathcal{O}(2^\lambda)$, g is a generator for \mathbb{G} and $e : \mathbb{G} \times \mathbb{G} \rightarrow \mathbb{G}_T$ is an efficiently computable non-degenerate bilinear map. The IBE scheme $\Pi_k = (\text{Setup}, \text{KeyGen}, \text{Enc}, \text{Dec})$ is defined over the identity space $\mathcal{ID} = \{\mathcal{ID}_\lambda\}_{\lambda \in \mathbb{N}}$ and the message space $\mathcal{M} = \{\mathcal{M}_\lambda\}_{\lambda \in \mathbb{N}}$. We assume that \mathcal{M} is a small subset of \mathbb{G}_T , namely $|\mathcal{M}| < |\mathbb{G}_T|^{1/2}$. While the reason for this restriction will become clear eventually, we note that it is not very serious since the space of valid messages in reality is expected to be significantly smaller than $|\mathbb{G}_T|^{1/2}$. Finally, let $H : \mathcal{ID} \rightarrow \mathbb{G}$ be a publicly available collision-resistant hash function.

- **Setup:** The setup algorithm samples $(\mathbb{G}, \mathbb{G}_T, q, g, e) \xleftarrow{R} \text{GroupGen}(1^\lambda)$ on input the security parameter 1^λ . It also samples $s, x_1, \dots, x_{k+2} \xleftarrow{R} \mathbb{Z}_q^*$ and $g_1, \dots, g_{k+2} \xleftarrow{R} \mathbb{G}$. It outputs the public parameter pp and the master secret-key msk as:

$$\begin{aligned} \text{pp} &= (g, g^s, g^{x_1}, \dots, g^{x_{k+2}}) \\ \text{msk} &= (s, g_1, \dots, g_{k+2}, (g_1^{x_1} \cdot g_{k+2}^{x_{k+2}}), \dots, (g_{k+1}^{x_{k+1}} \cdot g_{k+2}^{x_{k+2}})) \end{aligned}$$

- **KeyGen:** On input the master secret-key msk and an identity $\text{id} \in \mathcal{ID}$, the key generation algorithm samples $y_1, \dots, y_{k+1} \xleftarrow{R} \mathbb{Z}_q^*$ and outputs the secret-key $\text{sk}_{\text{id}} = (d_0, \dots, d_{k+2})$ where:

$$d_{j-1} = g_j^{y_j} \text{ for } j \in [1, k+1], \quad d_{k+1} = g_{k+2}^{\sum_{j=1}^{k+1} y_j}, \quad d_{k+2} = \left(\prod_{j=1}^{k+1} (g_j^{x_j} \cdot g_{k+2}^{x_{k+2}})^{y_j} \right) \cdot (H(\text{id}))^s$$

Observe that $\text{sk}_{\text{id}} = \text{PKE.Enc}(PK, (H(\text{id}))^s)$.

- **Enc:** On input the public parameter pp , an identity $\text{id} \in \mathcal{ID}$ and a message $M \in \mathcal{M}$, the encryption algorithm samples $r \xleftarrow{R} \mathbb{Z}_q^*$ and outputs the ciphertext $C = (c_0, \dots, c_{k+3})$ where:

$$c_0 = g^r, \quad c_j = (g^{x_j})^r \text{ for } j \in [1, k+2], \quad c_{k+3} = M \cdot e(H(\text{id}), g^s)^r$$

- **Dec:** On input a ciphertext $C = (c_0, \dots, c_{k+3})$ and a secret-key $\text{sk}_{\text{id}} = (d_0, \dots, d_{k+2})$, the decryption algorithm computes:

$$M' = c_{k+3} \cdot \frac{\prod_{j=1}^{k+2} e(d_{j-1}, c_j)}{e(d_{k+2}, c_0)}$$

If $M' \in \mathcal{M}$, the decryption algorithm outputs M' , else it outputs \perp .

Correctness. First, consider a message $M \in \mathcal{M}$, a ciphertext $C = (c_0, \dots, c_{k+3})$ corresponding to M under an identity $\text{id} \in \mathcal{ID}$ and a secret-key $\text{sk}_{\text{id}} = (d_0, \dots, d_{k+2})$ corresponding to id . Then, we have:

$$\begin{aligned} M' &= M \cdot e(H(\text{id}), g^s)^r \cdot \frac{\prod_{j=1}^{k+1} e(g_j^{y_j}, (g^{x_j})^r) \cdot e(g^{\sum_{j=1}^{k+1} y_j}, (g^{x_{k+2}})^r)}{e\left(\left(\prod_{j=1}^{k+1} (g_j^{x_j} \cdot g_{k+2}^{x_{k+2}})^{y_j}\right) \cdot (H(\text{id}))^s, g^r\right)} \\ &= M \cdot e(H(\text{id}), g^s)^r \cdot \frac{\prod_{j=1}^{k+1} e(g_j^{y_j}, (g^{x_j})^r) \cdot e\left(g^{\sum_{j=1}^{k+1} y_j}, (g^{x_{k+2}})^r\right)}{\prod_{j=1}^{k+1} e(g_j^{x_j}, g^r)^{y_j} \cdot e(g_{k+2}^{x_{k+2}}, g^r)^{\sum_{j=1}^{k+1} y_j} \cdot e((H(\text{id}))^s, g^r)} \\ &= M \end{aligned}$$

Therefore as long as the ciphertext and the secret-key correspond to the same identity, the message is recovered correctly. Again, when the ciphertext and the secret-key correspond to two different identities, say id and id' respectively, the decryption algorithm computes:

$$M' = M \cdot e \left(H(\text{id}) \cdot (H(\text{id}'))^{-1}, g^s \right)^r$$

Now, the restriction that \mathcal{M} is a small subset of \mathbb{G}_T of size less than $|\mathbb{G}_T|^{1/2}$ ensures that the probability of M' still lying in \mathcal{M} , for $s, r \xleftarrow{R} \mathbb{Z}_q^*$ and a collision-resistant hash function H , is negligible in the security parameter λ . This completes the proof of correctness for our generalized IBE scheme Π_k^{IBE} .

4.1 Security of Our IBE Scheme

Data Privacy. We state the following theorem for the data privacy of Π_k^{IBE} :

Theorem 4.1 *Our IBE scheme Π_k^{IBE} is adaptively data private under the DBDH assumption.*

Proof Overview. We begin by providing a brief overview of the proof idea here. The proof considers a probabilistic polynomial-time that has access to the public parameters of our scheme, and can determine from a *challenge ciphertext* whether it is associated with either of two identity-message pairs (id_0^*, M_0^*) or (id_1^*, M_1^*) with non-negligible advantage. The adversary can request for secret-keys corresponding to identities of its choice, subject to the restriction that the requested identities differ from both id_0^* and id_1^* .

In order to provide adaptive data privacy under the DBDH assumption, the proof considers a simulator that receives the DBDH instance $(g, g^{a_1}, g^{a_2}, g^{a_3}, Z)$ as input and chooses $x_1, \dots, x_{k+2} \xleftarrow{R} \mathbb{Z}_q^*$ as input. The public parameters are set up as $(g, g^{a_1}, g^{x_1}, \dots, g^{x_{k+2}})$, where $x_1, \dots, x_{k+2} \xleftarrow{R} \mathbb{Z}_q^*$. The simulator simulates the hash function H as a random oracle. On receiving a hash query for some identity id_i , The simulator responds with either g^{α_i} or $(g^{a_2})^{\alpha_i}$ for $\alpha_i \xleftarrow{R} \mathbb{Z}_q^*$, and remembers the response for future queries.

The simulator also simulates the secret-key generation oracle. On receiving a secret-key query for id_i , the simulator looks up $H(\text{id}_i)$, and either aborts if $H(\text{id}_i) = (g^{a_2})^{\alpha_i}$, or chooses $y_1, \dots, y_{k+1} \xleftarrow{R} \mathbb{Z}_q^*$ and responds with:

$$\text{sk}_{\text{id}_i} = \left(g_1^{y_1}, \dots, g_2^{y_{k+1}}, g_{k+2}^{\sum_{j=1}^{k+1} y_j}, \left(\prod_{j=1}^{k+2} (g_j^{x_j} \cdot g_{k+2}^{x_{k+2}})^{y_j} \right) \cdot (g^{a_1})^{\alpha_i} \right)$$

It is easy to see when the simulator does not abort, the distribution of sk_{id_i} is exactly as in the real world.

Finally, in the challenge phase, the simulator chooses $b \xleftarrow{R} \{0, 1\}$ and looks up $H(\text{id}_b^*)$. It aborts if $H(\text{id}_b^*)$ is of the form g^{α^*} for some $\alpha^* \xleftarrow{R} \mathbb{Z}_q^*$. Otherwise, it must be the case that $H(\text{id}_b^*)$ is of the form $g^{a_2 \cdot \alpha^*}$. In this case, it embeds the DBDH challenge in the challenge ciphertext C^* as follows:

$$C^* = \left((g^{a_3})^{(\alpha_b^*)^{-1}}, (g^{a_3})^{x_1 \cdot (\alpha_b^*)^{-1}}, \dots, (g^{a_3})^{x_{k+2} \cdot (\alpha_b^*)^{-1}}, M_b^* \cdot Z \right)$$

This ciphertext is either well-formed or uniform and independent in the adversary's view depending on the DBDH instance received by the simulator as input. This allows us to relate the advantage of the simulator in solving the DBDH instance to the advantage of the adversary in breaking the adaptive data privacy of our scheme. Finally, the abort condition in the secret-key query phase and the challenge phase is mutually opposite, so that the overall probability of abortion can be suitably minimized. This is essentially a modification of the original proof idea in [4] to fit our function private scheme.

Detailed Proof. We now present the detailed proof. Let \mathcal{A} be any probabilistic polynomial-time adversary. The proof aims to show the following:

$$\mathbf{Adv}_{\Pi_k^{\text{IBE}}, \mathcal{A}}^{\text{DP}}(\lambda) = \left| \Pr \left[\text{Expt}_{\text{DP}, \Pi_k^{\text{IBE}}, \mathcal{A}}^{(0)}(\lambda) = 1 \right] - \Pr \left[\text{Expt}_{\text{DP}, \Pi_k^{\text{IBE}}, \mathcal{A}}^{(1)}(\lambda) = 1 \right] \right| \leq \text{negl}(\lambda)$$

We define a second experiment $\text{Expt}_{\text{rand}, \text{DP}, \Pi_k^{\text{IBE}}, \mathcal{A}}^{(b)}(\lambda)$ that is identical to $\text{Expt}_{\text{DP}, \Pi_k^{\text{IBE}}, \mathcal{A}}^{(b)}(\lambda)$ except in Step 3, where the challenge ciphertext C^* is generated uniformly and independent of the challenge pair (id_b^*, M_b^*) . Then, the following is obvious to see:

$$\left| \Pr \left[\text{Expt}_{\text{rand}, \text{DP}, \Pi_k^{\text{IBE}}, \mathcal{A}}^{(0)}(\lambda) = 1 \right] - \Pr \left[\text{Expt}_{\text{rand}, \text{DP}, \Pi_k^{\text{IBE}}, \mathcal{A}}^{(1)}(\lambda) = 1 \right] \right| = 0$$

We now state and prove the following claim:

Claim 4.1 *For any probabilistic polynomial-time adversary \mathcal{A} and for $b \in \{0, 1\}$, the following holds:*

$$\left| \Pr \left[\text{Expt}_{\text{DP}, \Pi_k^{\text{IBE}}, \mathcal{A}}^{(b)}(\lambda) = 1 \right] - \Pr \left[\text{Expt}_{\text{rand}, \text{DP}, \Pi_k^{\text{IBE}}, \mathcal{A}}^{(b)}(\lambda) = 1 \right] \right| \leq \text{negl}(\lambda)$$

To prove this claim, we assume the contrary. Let \mathcal{A} be a probabilistic polynomial-time adversary such that:

$$\left| \Pr \left[\text{Expt}_{\text{DP}, \Pi_k^{\text{IBE}}, \mathcal{A}}^{(b)}(\lambda) = 1 \right] - \Pr \left[\text{Expt}_{\text{rand}, \text{DP}, \Pi_k^{\text{IBE}}, \mathcal{A}}^{(b)}(\lambda) = 1 \right] \right| = \epsilon > \text{negl}(\lambda)$$

Also, let $Q_K = \text{poly}(\lambda)$ be the maximum number of secret-key queries made by \mathcal{A} in either of the experiments $\text{Expt}_{\text{DP}, \Pi_k^{\text{IBE}}, \mathcal{A}}^{(b)}(\lambda)$ and $\text{Expt}_{\text{rand}, \text{DP}, \Pi_k^{\text{IBE}}, \mathcal{A}}^{(b)}(\lambda)$. We construct an algorithm \mathcal{B} that solves an instance of the DBDH problem with non-negligible advantage $\epsilon' \geq \epsilon / (\exp(1) \cdot (Q_K + 1))$. \mathcal{B} is given $(g, g^{a_1}, g^{a_2}, g^{a_3}, Z)$ and interacts with \mathcal{A} as follows:

- **Setup:** \mathcal{B} samples $x_1, \dots, x_k \xleftarrow{R} \mathbb{Z}_q^*$ and $g_1, \dots, g_{k+2} \xleftarrow{R} \mathbb{G}$. It sets $\text{pp} = (g, g^{a_1}, g^{x_1}, \dots, g^{x_{k+2}})$ and provides pp to \mathcal{A} . Observe that the distribution of pp is exactly as in the real world.
- **H – Queries:** \mathcal{B} maintains a list of tuples of the form $\langle \text{id}_j, H(\text{id}_j), \alpha_j, \beta_j \rangle$ such that $H(\text{id}_j) \in \mathbb{G}$, $\alpha_j \in \mathbb{Z}_q^*$ and $\beta_j \in \{0, 1\}$. When \mathcal{A} issues a hash query for $\text{id}_i \in \mathcal{ID}$, \mathcal{B} responds as follows:
 1. \mathcal{B} searches the list for a matching tuple of the form $\langle \text{id}_i, H(\text{id}_i), \alpha_i, \beta_i \rangle$. If such a tuple is found, it returns $H(\text{id}_i)$ to \mathcal{A} .
 2. Otherwise, \mathcal{B} samples $\alpha_i \xleftarrow{R} \mathbb{Z}_q^*$ and $\beta_i \leftarrow \{0, 1\}$ such that $\Pr[\beta_i = 0] = 1 - 1/(Q_K + 1)$.
 - If $\beta_i = 0$, it sets $H(\text{id}_i) = g^{\alpha_i}$
 - If $\beta_i = 1$, it sets $H(\text{id}_i) = (g^{a_2})^{\alpha_i}$, where g^{a_2} is a part of its input instance.
 3. \mathcal{B} now adds the tuple $\langle \text{id}_i, H(\text{id}_i) \rangle$ to its existing list, and returns $H(\text{id}_i)$ to \mathcal{A} .
- **Secret-Key Queries:** When \mathcal{A} issues a secret-key query for $\text{id}_i \in \mathcal{ID}$, \mathcal{B} responds as follows:
 1. \mathcal{B} runs the aforementioned procedure to look up $H(\text{id}_i)$ and obtains the tuple $\langle \text{id}_i, H(\text{id}_i), \alpha_i, \beta_i \rangle$.
 2. If $\beta_i = 1$, \mathcal{B} aborts by outputting a uniformly sampled bit.
 3. If $\beta_i = 0$, we must have $H(\text{id}_i) = g^{\alpha_i}$; hence, \mathcal{B} chooses $y_1, \dots, y_{k+1} \xleftarrow{R} \mathbb{Z}_q^*$ and responds with:

$$\text{sk}_{\text{id}_i} = \left(g_1^{y_1}, \dots, g_2^{y_{k+1}}, g_{k+2}^{\sum_{j=1}^{k+1} y_j}, \left(\prod_{j=1}^{k+2} (g_j^{x_j} \cdot g_{k+2}^{x_{k+2}})^{y_j} \right) \cdot (g^{a_1})^{\alpha_i} \right)$$

Once again, observe that the distribution of sk_{id_i} is exactly as in the real world.

- **Challenge:** \mathcal{A} outputs the challenge pair $((\text{id}_0^*, M_0^*), (\text{id}_1^*, M_1^*))$. \mathcal{B} samples $b \xleftarrow{R} \{0, 1\}$ and does the following:

1. \mathcal{B} runs the aforementioned procedure to look up $H(\text{id}_b^*)$ and obtains the tuple $(\text{id}_b^*, H(\text{id}_b^*), \alpha_b^*, \beta_b^*)$.
2. If $\beta_b^* = 0$, \mathcal{B} aborts by outputting a uniformly sampled bit.
3. If $\beta_b^* = 1$, we must have $H(\text{id}_b^*) = (g^{a_2})^{\alpha_b^*}$. \mathcal{B} accordingly responds with the challenge ciphertext C^* as:

$$C^* = \left((g^{a_3})^{(\alpha_b^*)^{-1}}, (g^{a_3})^{x_1 \cdot (\alpha_b^*)^{-1}}, \dots, (g^{a_3})^{x_{k+2} \cdot (\alpha_b^*)^{-1}}, M_b^* \cdot Z \right)$$

where g^{a_3} and Z are part of its input instance.

- **Guess:** At the end of the game, \mathcal{A} outputs a bit b' . \mathcal{B} outputs the same bit b' .

It is easy to see that when $Z = e(g, g)^{a_1 \cdot a_2 \cdot a_3}$, the ciphertext C^* is well-formed and identically distributed to the challenge in the experiment $\text{Expt}_{\text{DP}, \Pi_k^{\text{IBE}}, \mathcal{A}}^{(b)}(\lambda)$. To see this, set $s = a_1$ and $r = a_3 \cdot (\alpha_b^*)^{-1}$ and observe that:

$$\begin{aligned} (g^{a_3})^{(\alpha_b^*)^{-1}} &= g^r, \quad (g^{a_3})^{x_j \cdot (\alpha_b^*)^{-1}} = (g^{x_j})^r \text{ for } j \in [1, k+2] \\ Z &= e(g, g)^{a_1 \cdot a_2 \cdot a_3} = e\left((g^{a_2})^{\alpha_b^*}, g^{a_1}\right)^{a_3 \cdot (\alpha_b^*)^{-1}} = e(H(\text{id}_b^*), g^s)^r \end{aligned}$$

On the other hand, if Z is uniform random in \mathbb{G}_T , then C^* is independent of b and hence identically distributed to the challenge in the experiment $\text{Expt}_{\text{rand}, \text{DP}, \Pi_k^{\text{IBE}}, \mathcal{A}}^{(b)}(\lambda)$.

It remains to bound the probability that \mathcal{B} aborts either during a secret-key query or during the challenge phase, denoted as $\Pr[\text{Abort}]$. It is easy to see that the probability that \mathcal{B} aborts during a given secret-key query is $1/(Q_K + 1)$. Assuming that \mathcal{B} makes at most Q_K secret-key queries, the probability that it never aborts for any of these queries is thus lower bounded as $(1 - 1/(Q_K + 1))^{Q_K}$. The probability that \mathcal{B} does not abort during the challenge phase is lower bounded as $1/(Q_K + 1)$. Thus the overall probability that \mathcal{B} does not abort during the entire game (equivalently $\Pr[\overline{\text{Abort}}]$) is lower bounded as $(1 - 1/(Q_K + 1))^{Q_K} / (Q_K + 1)$.

Finally, the advantage ϵ' of \mathcal{B} in solving the DBDH instance (where the probability is taken over all possible choices of $a_1, a_2, a_3 \xleftarrow{R} \mathbb{Z}_q^*$) and all possible choice of $Z \xleftarrow{R} \mathbb{G}_T$) may be quantified as:

$$\begin{aligned} \epsilon' &= \left| \Pr[\mathcal{B}(g, g^{a_1}, g^{a_2}, g^{a_3}, e(g, g)^{a_1 \cdot a_2 \cdot a_3}) = 1] - \Pr[\mathcal{B}(g, g^{a_1}, g^{a_2}, g^{a_3}, Z) = 1] \right| \\ &= \left| \Pr[\mathcal{B}(g, g^{a_1}, g^{a_2}, g^{a_3}, e(g, g)^{a_1 \cdot a_2 \cdot a_3}) = 1 | \overline{\text{Abort}}] \cdot \Pr[\overline{\text{Abort}}] - \right. \\ &\quad \left. \Pr[\mathcal{B}(g, g^{a_1}, g^{a_2}, g^{a_3}, Z) = 1 | \overline{\text{Abort}}] \cdot \Pr[\overline{\text{Abort}}] \right| \\ &= \left| \Pr[\text{Expt}_{\text{DP}, \Pi_k^{\text{IBE}}, \mathcal{A}}^{(b)}(\lambda) = 1] \cdot \Pr[\overline{\text{Abort}}] - \Pr[\text{Expt}_{\text{rand}, \text{DP}, \Pi_k^{\text{IBE}}, \mathcal{A}}^{(b)}(\lambda) = 1] \cdot \Pr[\overline{\text{Abort}}] \right| \\ &\geq \epsilon \cdot (1 - 1/(Q_K + 1))^{Q_K} / (Q_K + 1) \\ &\geq \epsilon / (\exp(1) \cdot (Q_K + 1)) \end{aligned}$$

The above derivation uses the fact that the abort condition for \mathcal{B} is independent of the view of the adversary \mathcal{A} . This completes the proof of Claim 4.1. The proof of data privacy for Π_k^{IBE} now follows from the following observation:

$$\begin{aligned}
\mathbf{Adv}_{\Pi_k^{\text{IBE}}, \mathcal{A}}^{\text{DP}}(\lambda) &= \left| \Pr \left[\text{Expt}_{\text{DP}, \Pi_k^{\text{IBE}}, \mathcal{A}}^{(0)}(\lambda) = 1 \right] - \Pr \left[\text{Expt}_{\text{DP}, \Pi_k^{\text{IBE}}, \mathcal{A}}^{(1)}(\lambda) = 1 \right] \right| \\
&\leq \left| \Pr \left[\text{Expt}_{\text{DP}, \Pi_k^{\text{IBE}}, \mathcal{A}}^{(0)}(\lambda) = 1 \right] - \Pr \left[\text{Expt}_{\text{rand}, \text{DP}, \Pi_k^{\text{IBE}}, \mathcal{A}}^{(0)}(\lambda) = 1 \right] \right| + \\
&\quad \left| \Pr \left[\text{Expt}_{\text{DP}, \Pi_k^{\text{IBE}}, \mathcal{A}}^{(1)}(\lambda) = 1 \right] - \Pr \left[\text{Expt}_{\text{rand}, \text{DP}, \Pi_k^{\text{IBE}}, \mathcal{A}}^{(1)}(\lambda) = 1 \right] \right| + \\
&\quad \left| \Pr \left[\text{Expt}_{\text{rand}, \text{DP}, \Pi_k^{\text{IBE}}, \mathcal{A}}^{(0)}(\lambda) = 1 \right] - \Pr \left[\text{Expt}_{\text{rand}, \text{DP}, \Pi_k^{\text{IBE}}, \mathcal{A}}^{(1)}(\lambda) = 1 \right] \right| \\
&\leq 2 \left| \Pr \left[\text{Expt}_{\text{DP}, \Pi_k^{\text{IBE}}, \mathcal{A}}^{(b)}(\lambda) = 1 \right] - \Pr \left[\text{Expt}_{\text{rand}, \text{DP}, \Pi_k^{\text{IBE}}, \mathcal{A}}^{(b)}(\lambda) = 1 \right] \right| \quad (\text{from Equation 1}) \\
&= 2\epsilon \leq \text{negl}(\lambda)
\end{aligned}$$

The above proof reinforces the fact that applying our *encrypt-augment-cancel* approach *does not affect* the adaptive data privacy of the original IBE scheme of Boneh and Franklin.

Computational Function Privacy. We state the following theorem for the computational function privacy of Π_k^{IBE} :

Theorem 4.2 *Our IBE scheme Π_k^{IBE} is computationally function private under the $(k+1)$ -DLIN assumption for identities sampled uniformly from k -sources with $k = \omega(\log \lambda)$.*

Proof. We begin by stating the following claim:

Claim 4.2 *For any probabilistic polynomial-time adversary \mathcal{A} and for mode $\xleftarrow{R} \{\text{real}, \text{rand}\}$, the following holds:*

$$\left| \Pr \left[\text{Expt}_{\text{FP}, \Pi_k^{\text{IBE}}, \mathcal{A}}^{\text{mode}}(\lambda) = 1 \right] - \Pr \left[\text{Expt}_{\text{FP}, \Pi_k^{\text{IBE}}, \mathcal{A}}^{\text{rand}}(\lambda) = 1 \right] \right| \leq \text{negl}(\lambda)$$

To prove this claim, we assume the contrary. Let \mathcal{A} be a probabilistic polynomial-time adversary such that for mode $\xleftarrow{R} \{\text{real}, \text{rand}\}$, we have:

$$\left| \Pr \left[\text{Expt}_{\text{FP}, \Pi_k^{\text{IBE}}, \mathcal{A}}^{\text{mode}}(\lambda) = 1 \right] - \Pr \left[\text{Expt}_{\text{FP}, \Pi_k^{\text{IBE}}, \mathcal{A}}^{\text{rand}}(\lambda) = 1 \right] \right| = \epsilon > \text{negl}(\lambda)$$

We assume that the adversary \mathcal{A} issues a single query to the real-or-random oracle. As discussed in Section 3, such a single-shot adversary is polynomially equivalent to its multi-shot variant considered in Definition 3.1. We construct an algorithm \mathcal{B} that solves an instance of the $(k+1)$ -DLIN problem with non-negligible advantage $\epsilon' = \epsilon$. \mathcal{B} is given $(g_1, \dots, g_{k+2}, g_1^{a_1}, \dots, g_{k+2}^{a_{k+2}})$ and interacts with \mathcal{A} as follows:

- **Setup:** \mathcal{B} samples $s, x_1, \dots, x_{k+2} \xleftarrow{R} \mathbb{Z}_q^*$ and $g \xleftarrow{R} \mathbb{G}$. It sets $\mathbf{pp} = (g, g^s, g^{x_1}, \dots, g^{x_{k+2}})$ and provides \mathbf{pp} to \mathcal{A} . Observe that the distribution of \mathbf{pp} is exactly as in the real world.
- **H – Queries:** \mathcal{B} maintains a list of tuples of the form $\langle \text{id}_j, H(\text{id}_j) \rangle$ such that $H(\text{id}_j) \in \mathbb{G}$. When \mathcal{A} issues a hash query for $\text{id}_i \in \mathcal{ID}$, \mathcal{B} responds as follows:
 1. \mathcal{B} searches the list for a matching tuple of the form $\langle \text{id}_i, H(\text{id}_i) \rangle$. If such a tuple is found, it returns $H(\text{id}_i)$ to \mathcal{A} .
 2. Otherwise, \mathcal{B} samples $H(\text{id}_i) \xleftarrow{R} \mathbb{G}$ and adds the tuple $\langle \text{id}_i, H(\text{id}_i) \rangle$ to its existing list. It returns $H(\text{id}_i)$ to \mathcal{A} .

- **Secret-Key Queries:** When \mathcal{A} issues a secret-key query for $\text{id}_i \in \mathcal{ID}$, \mathcal{B} responds as follows:
 1. \mathcal{B} runs the aforementioned procedure to look up $H(\text{id}_i)$.
 2. \mathcal{B} samples $y_1, \dots, y_{k+1} \xleftarrow{R} \mathbb{Z}_q^*$ and responds with:

$$\text{sk}_{\text{id}_i} = \left(g_1^{y_1}, \dots, g_{k+1}^{y_{k+1}}, g_{k+2}^{\sum_{j=1}^{k+1} y_j}, \left(\prod_{j=1}^{k+1} (g_j^{x_j} \cdot g_{k+2}^{x_{k+2}})^{y_j} \right) \cdot (H(\text{id}_i))^s \right)$$

where g_1, \dots, g_{k+2} are part of its input instance. Once again, observe that the distribution of sk_{id_i} is exactly as in the real world.

- **Real-or-Random Query:** Suppose \mathcal{A} queries the real-or-random oracle with \mathbf{ID}^* - a circuit representing a k -source over the identity space \mathcal{ID} such that $k = \omega(\log \lambda)$. \mathcal{B} samples $\text{mode} \xleftarrow{R} \{\text{real}, \text{rand}\}$ and does the following:
 1. If $\text{mode} = \text{real}$, \mathcal{B} samples $\text{id}^* \xleftarrow{R} \mathbf{ID}^*$, while if $\text{mode} = \text{rand}$, it samples $\text{id}^* \xleftarrow{R} \mathcal{ID}$. It then runs the aforementioned procedure to look up $H(\text{id}^*)$.
 2. \mathcal{B} responds with the secret-key sk_{id^*} as:

$$\text{sk}_{\text{id}^*} = \left(g_1^{a_1}, \dots, g_{k+2}^{a_{k+2}}, \left(\prod_{j=1}^{k+2} (g_j^{a_j})^{x_j} \right) \cdot (H(\text{id}^*))^s \right)$$

where $g_1^{a_1}, \dots, g_{k+2}^{a_{k+2}}$ are part of its input instance.

- **Guess:** At the end of the game, \mathcal{A} outputs a bit b' . \mathcal{B} outputs the same bit b' .

It is easy to see that when $a_{k+2} = \sum_{j=1}^{k+1} a_j$, the secret-key sk_{id^*} is well-formed and identically distributed to the response of the real-or-random oracle in the experiment $\text{Expt}_{\text{FP}, \Pi_k^{\text{IBE}}, \mathcal{A}}^{\text{mode}}(\lambda)$. On the other hand, when a_{k+2} is uniformly random in \mathbb{Z}_q^* , the secret-key sk_{id^*} is uniformly random, and hence identically distributed to the response of the real-or-random oracle in the experiment $\text{Expt}_{\text{FP}, \Pi_k^{\text{IBE}}, \mathcal{A}}^{\text{rand}}(\lambda)$. Now, the advantage ϵ' of \mathcal{B} in solving the $(k+1)$ -DLIN instance (where the probability is taken over all possible choices of $a_1, \dots, a_{k+2} \xleftarrow{R} \mathbb{Z}_q^*$ and all possible choices of $g_1, \dots, g_{k+2} \xleftarrow{R} \mathbb{G}$) may be quantified as:

$$\begin{aligned} \epsilon' &= \left| \Pr \left[\mathcal{B} \left(g_1, \dots, g_{k+2}, g_1^{a_1}, \dots, g_{k+2}^{\sum_{j=1}^{k+1} a_j} \right) = 1 \right] - \Pr \left[\mathcal{B} \left(g_1, \dots, g_{k+2}, g_1^{a_1}, \dots, g_{k+2}^{a_{k+2}} \right) = 1 \right] \right| \\ &= \left| \Pr \left[\text{Expt}_{\text{FP}, \Pi_k^{\text{IBE}}, \mathcal{A}}^{\text{mode}}(\lambda) = 1 \right] - \Pr \left[\text{Expt}_{\text{FP}, \Pi_k^{\text{IBE}}, \mathcal{A}}^{\text{rand}}(\lambda) = 1 \right] \right| \\ &= \epsilon \end{aligned}$$

This completes the proof of Claim 4.2. The proof of function privacy for Π_k^{IBE} now follows from the following observation:

$$\begin{aligned} \mathbf{Adv}_{\Pi_k^{\text{IBE}}, \mathcal{A}}^{\text{FP}}(\lambda) &= \left| \Pr \left[\text{Expt}_{\text{FP}, \Pi_k^{\text{IBE}}, \mathcal{A}}^{\text{real}}(\lambda) = 1 \right] - \Pr \left[\text{Expt}_{\text{FP}, \Pi_k^{\text{IBE}}, \mathcal{A}}^{\text{rand}}(\lambda) = 1 \right] \right| \\ &\leq 2 \left| \Pr \left[\text{Expt}_{\text{FP}, \Pi_k^{\text{IBE}}, \mathcal{A}}^{\text{mode}}(\lambda) = 1 \right] - \Pr \left[\text{Expt}_{\text{FP}, \Pi_k^{\text{IBE}}, \mathcal{A}}^{\text{rand}}(\lambda) = 1 \right] \right| \\ &= 2\epsilon \leq \text{negl}(\lambda) \end{aligned}$$

Enhanced Computational Function Privacy. We state the following theorem for the enhanced computational function privacy of Π_k^{IBE} :

Theorem 4.3 *Our IBE scheme Π_k^{IBE} is computationally enhanced function private under the $(k+1)$ -DLIN assumption for identities sampled uniformly from k -sources with $k = \omega(\log \lambda)$.*

Proof Overview. The proof is very similar to that of Theorem 4.2; we begin by providing a proof overview here. An important component of the proof is to show that the challenger \mathcal{B} can additionally simulate the function privacy encryption oracle, and that the real and random modes of operation of the real-or-random function privacy oracle are indistinguishable even in the presence of the encryption oracle. The first part is easy to demonstrate. Let $\text{sk}_{\text{id}^*} = (d_0^*, \dots, d_{k+2}^*)$ be the response of the real-or-random function privacy oracle to the adversary \mathcal{A} . On receiving a function-privacy encryption oracle query for a payload message M , \mathcal{B} responds as follows: it samples $r \xleftarrow{R} \mathbb{Z}_q^*$ and outputs the ciphertext $C^* = (c_0^*, \dots, c_{k+3}^*)$ where:

$$c_0^* = g^r, c_j^* = (g^{x_j})^r \text{ for } j \in [1, k+2]$$

and

$$c_{k+3}^* = M \cdot \frac{e(d_{k+2}^*, c_0^*)}{\prod_{j=1}^{k+2} e(d_{j-1}^*, c_j^*)}$$

For the second part, observe that when the input $(k+1)$ -DLIN challenge for \mathcal{B} is valid, the ciphertext C^* is well-formed and identically distributed to the response of the function privacy encryption oracle in the experiment $\text{Expt}_{\text{EFP}, \Pi_k^{\text{IBE}}, \mathcal{A}}^{\text{mode}}(\lambda)$. On the other hand, when the input $(k+1)$ -DLIN challenge for \mathcal{B} is invalid, the ciphertext C^* is uniformly random, and hence identically distributed to the response of the function privacy encryption oracle in the experiment $\text{Expt}_{\text{EFP}, \Pi_k^{\text{IBE}}, \mathcal{A}}^{\text{mode}}(\lambda)$. In either case, decrypting C^* using sk_{id^*} correctly retrieves M . Hence, for the real and random modes of operation of the function privacy oracle, the response of the encryption oracle is also indistinguishable under the $(k+1)$ -DLIN assumption. This completes the proof overview for the enhanced computational function privacy of Π_k^{IBE} .

Detailed Proof. We now present the detailed proof. We begin by stating the following claim:

Claim 4.3 *For any probabilistic polynomial-time adversary \mathcal{A} and for mode $\xleftarrow{R} \{\text{real}, \text{rand}\}$, the following holds:*

$$\left| \Pr \left[\text{Expt}_{\text{EFP}, \Pi_k^{\text{IBE}}, \mathcal{A}}^{\text{mode}}(\lambda) = 1 \right] - \Pr \left[\text{Expt}_{\text{EFP}, \Pi_k^{\text{IBE}}, \mathcal{A}}^{\text{rand}}(\lambda) = 1 \right] \right| \leq \text{negl}(\lambda)$$

To prove this claim, we assume the contrary. Let \mathcal{A} be a probabilistic polynomial-time adversary such that for mode $\xleftarrow{R} \{\text{real}, \text{rand}\}$, we have:

$$\left| \Pr \left[\text{Expt}_{\text{EFP}, \Pi_k^{\text{IBE}}, \mathcal{A}}^{\text{mode}}(\lambda) = 1 \right] - \Pr \left[\text{Expt}_{\text{EFP}, \Pi_k^{\text{IBE}}, \mathcal{A}}^{\text{rand}}(\lambda) = 1 \right] \right| = \epsilon > \text{negl}(\lambda)$$

We assume that the adversary \mathcal{A} issues a single query to the real-or-random oracle. As discussed in Section 3, such a single-shot adversary is polynomially equivalent to its multi-shot variant considered in Definition 3.1. We construct an algorithm \mathcal{B} that solves an instance of the $(k+1)$ -DLIN problem with non-negligible advantage $\epsilon' = \epsilon$. \mathcal{B} is given $(g_1, \dots, g_{k+2}, g_1^{a_1}, \dots, g_{k+2}^{a_{k+2}})$ and interacts with \mathcal{A} as follows:

- **Setup:** \mathcal{B} samples $s, x_1, \dots, x_{k+2} \xleftarrow{R} \mathbb{Z}_q^*$ and a generator $g \xleftarrow{R} \mathbb{G}$. It sets $\text{pp} = (g, g^s, g^{x_1}, \dots, g^{x_{k+2}})$ and provides pp to \mathcal{A} . Observe that the distribution of pp is exactly as in the real world.
- **H – Queries:** \mathcal{B} maintains a list of tuples of the form $\langle \text{id}_j, H(\text{id}_j) \rangle$ such that $H(\text{id}_j) \in \mathbb{G}$. When \mathcal{A} issues a hash query for $\text{id}_i \in \mathcal{ID}$, \mathcal{B} responds as follows:
 1. \mathcal{B} searches the list for a matching tuple of the form $\langle \text{id}_i, H(\text{id}_i) \rangle$. If such a tuple is found, it returns $H(\text{id}_i)$ to \mathcal{A} .
 2. Otherwise, \mathcal{B} samples $H(\text{id}_i) \xleftarrow{R} \mathbb{G}$ and adds the tuple $\langle \text{id}_i, H(\text{id}_i) \rangle$ to its existing list. It returns $H(\text{id}_i)$ to \mathcal{A} .
- **Secret-Key Queries:** When \mathcal{A} issues a secret-key query for $\text{id}_i \in \mathcal{ID}$, \mathcal{B} responds as follows:

1. \mathcal{B} runs the aforementioned procedure to look up $H(\text{id}_i)$.
2. \mathcal{B} samples $y_1, \dots, y_{k+1} \xleftarrow{R} \mathbb{Z}_q^*$ and responds with:

$$\text{sk}_{\text{id}_i} = \left(g_1^{y_1}, \dots, g_{k+1}^{y_{k+1}}, g_{k+2}^{\sum_{j=1}^{k+1} y_j}, \left(\prod_{j=1}^{k+1} (g_j^{x_j} \cdot g_{k+2}^{x_{k+2}})^{y_j} \right) \cdot (H(\text{id}_i))^s \right)$$

where g_1, \dots, g_{k+2} are part of its input instance. Once again, observe that the distribution of sk_{id_i} is exactly as in the real world.

- **Real-or-Random Query:** Suppose \mathcal{A} queries the real-or-random oracle with ID^* - a circuit representing a k -source over the identity space \mathcal{ID} such that $k = \omega(\log \lambda)$. \mathcal{B} samples $\text{mode} \xleftarrow{R} \{\text{real}, \text{rand}\}$ and does the following:

1. If $\text{mode} = \text{real}$, \mathcal{B} samples $\text{id}^* \xleftarrow{R} \text{ID}^*$, while if $\text{mode} = \text{rand}$, it samples $\text{id}^* \xleftarrow{R} \mathcal{ID}$. It then runs the aforementioned procedure to look up $H(\text{id}^*)$.
2. \mathcal{B} responds with the secret-key sk_{id^*} as:

$$\text{sk}_{\text{id}^*} = \left(g_1^{a_1}, \dots, g_{k+2}^{a_{k+2}}, \left(\prod_{j=1}^{k+2} (g_j^{a_j})^{x_j} \right) \cdot (H(\text{id}^*))^s \right)$$

where $g_1^{a_1}, \dots, g_{k+2}^{a_{k+2}}$ are part of its input instance.

- **Function Privacy Encryption Oracle Query:** Suppose \mathcal{A} queries the function privacy encryption oracle with a message M . Let $\text{sk}_{\text{id}^*} = (d_0^*, \dots, d_{k+2}^*)$ be the response of the real-or-random function privacy oracle to the adversary \mathcal{A} . On receiving a function-privacy encryption oracle query for a payload message M , \mathcal{B} responds as follows: it samples $r \xleftarrow{R} \mathbb{Z}_q^*$ and outputs the ciphertext $C^* = (c_0^*, \dots, c_{k+3}^*)$ where:

$$c_0^* = g^r, \quad c_j^* = (g^{x_j})^r \quad \text{for } j \in [1, k+2]$$

and

$$c_{k+3}^* = M \cdot \frac{e(d_{k+2}^*, c_0^*)}{\prod_{j=1}^{k+2} e(d_{j-1}^*, c_j^*)}$$

- **Guess:** At the end of the game, \mathcal{A} outputs a bit b' . \mathcal{B} outputs the same bit b' .

It is now easy to see the following:

- When $a_{k+2} = \sum_{j=1}^{k+1} a_j$, the secret-key sk_{id^*} is well-formed and identically distributed to the response of the real-or-random oracle in the experiment $\text{Expt}_{\text{EFP}, \Pi_k^{\text{IBE}}, \mathcal{A}}^{\text{mode}}(\lambda)$. Additionally, the ciphertext C^* is also well-formed and identically distributed to the response of the function privacy encryption oracle in the experiment $\text{Expt}_{\text{EFP}, \Pi_k^{\text{IBE}}, \mathcal{A}}^{\text{mode}}(\lambda)$.
- On the other hand, when a_{k+2} is uniformly random in \mathbb{Z}_q^* , the secret-key sk_{id^*} is uniformly random, and hence identically distributed to the response of the real-or-random oracle in the experiment $\text{Expt}_{\text{EFP}, \Pi_k^{\text{IBE}}, \mathcal{A}}^{\text{rand}}(\lambda)$. Correspondingly, the ciphertext C^* is also uniformly random, and hence identically distributed to the response of the function privacy encryption oracle in the experiment $\text{Expt}_{\text{EFP}, \Pi_k^{\text{IBE}}, \mathcal{A}}^{\text{mode}}(\lambda)$.
- In either case, decrypting C^* using sk_{id^*} correctly retrieves M . This essentially implies that they correspond to the same underlying identity, which is either id^* sampled by \mathcal{B} when $a_{k+2} = \sum_{j=1}^{k+1} a_j$, or some other uniformly random identity in \mathcal{ID} when a_{k+2} is uniformly random in \mathbb{Z}_q^* .

Now, the advantage ϵ' of \mathcal{B} in solving the $(k+1)$ -DLIN instance (where the probability is taken over all possible choices of $a_1, \dots, a_{k+2} \xleftarrow{R} \mathbb{Z}_q^*$ and all possible choices of $g_1, \dots, g_{k+2} \xleftarrow{R} \mathbb{G}$) may be quantified as:

$$\begin{aligned} \epsilon' &= \left| \Pr \left[\mathcal{B} \left(g_1, \dots, g_{k+2}, g_1^{a_1}, \dots, g_{k+2}^{\sum_{j=1}^{k+1} a_j} \right) = 1 \right] - \Pr \left[\mathcal{B} \left(g_1, \dots, g_{k+2}, g_1^{a_1}, \dots, g_{k+2}^{a_{k+2}} \right) = 1 \right] \right| \\ &= \left| \Pr \left[\text{Expt}_{\text{EFP}, \Pi_k^{\text{IBE}}, \mathcal{A}}^{\text{mode}}(\lambda) = 1 \right] - \Pr \left[\text{Expt}_{\text{EFP}, \Pi_k^{\text{IBE}}, \mathcal{A}}^{\text{rand}}(\lambda) = 1 \right] \right| \\ &= \epsilon \end{aligned}$$

This completes the proof of Claim 4.3. The proof of enhanced function privacy for Π_k^{IBE} now follows from the following observation:

$$\begin{aligned} \text{Adv}_{\Pi_k^{\text{IBE}}, \mathcal{A}}^{\text{EFP}}(\lambda) &= \left| \Pr \left[\text{Expt}_{\text{EFP}, \Pi_k^{\text{IBE}}, \mathcal{A}}^{\text{real}}(\lambda) = 1 \right] - \Pr \left[\text{Expt}_{\text{EFP}, \Pi_k^{\text{IBE}}, \mathcal{A}}^{\text{rand}}(\lambda) = 1 \right] \right| \\ &\leq 2 \left| \Pr \left[\text{Expt}_{\text{EFP}, \Pi_k^{\text{IBE}}, \mathcal{A}}^{\text{mode}}(\lambda) = 1 \right] - \Pr \left[\text{Expt}_{\text{EFP}, \Pi_k^{\text{IBE}}, \mathcal{A}}^{\text{rand}}(\lambda) = 1 \right] \right| \\ &= 2\epsilon \leq \text{negl}(\lambda) \end{aligned}$$

5 Computationally Function private Inner-Product Encryption

In this section, we present a family of selectively data private zero-IPE schemes $\{\Pi_k^{\text{IPE}}\}_{k \geq 1}$ that are also computationally function private under the generalized family of k -DLIN assumptions in the standard model. Our schemes are defined over the set of attributes $\Sigma = \mathbb{Z}_N^n$ (N being a product of three primes q_1, q_2 and q_3), and the class of vectorial predicates $\mathcal{F} = \{f_{\vec{v}} \mid \vec{v} \in \mathbb{Z}_N^n\}$, such that for $I = (I_1, \dots, I_n) \in \mathbb{Z}_N^n$, we have $f_{\vec{v}}(I) = 1$ if and only if $\langle \vec{v}, I \rangle = 0 \pmod N$. Once again, our constructions are obtained by applying our *encrypt-augment-recover* approach to the zero-IPE scheme of Katz, Sahai and Waters [3].

Construction Overview. Let \mathbb{G} be a bilinear group of order $N = q_1 q_2 q_3$ (each of q_1, q_2 and q_3 being λ -bit primes), and let $\mathbb{G}_1, \mathbb{G}_2$ and \mathbb{G}_3 denote the subgroups of \mathbb{G} of order q_1, q_2 and q_3 , respectively. Also, let $\hat{e} : \mathbb{G} \times \mathbb{G} \rightarrow \mathbb{G}_T$ be an efficiently computable non-degenerate bilinear map, where \mathbb{G}_T is also a group of order N . Note that if g is the generator for \mathbb{G} , then the element $g_1 = g^{q_2 \cdot q_3}$ is a generator for \mathbb{G}_1 , the element $g_2 = g^{q_1 \cdot q_3}$ is a generator for \mathbb{G}_2 , and the element $g_3 = g^{q_1 \cdot q_2}$ is a generator for \mathbb{G}_3 . Furthermore, for any elements $h_1 \in \mathbb{G}_1, h_2 \in \mathbb{G}_2$ and $h_3 \in \mathbb{G}_3$, we have $\hat{e}(h_1, h_2) = \hat{e}(h_2, h_3) = \hat{e}(h_1, h_3) = 1$. Also, let $\text{GroupGen}'(1^\lambda)$ be a probabilistic polynomial-time algorithm that takes as input a security parameter λ , and outputs the tuple $(\mathbb{G}, \mathbb{G}_T, q_1, q_2, q_3, g_1, g_2, g_3, \hat{e})$. Finally, the payload message space \mathcal{M} is assumed to be a small subset of \mathbb{G}_T , namely $|\mathcal{M}| < |\mathbb{G}_T|^{1/2}$. Our function private zero-IPE scheme uses the three subgroups for three distinct roles:

- The subgroup \mathbb{G}_2 is used to encode the vectors \vec{v} and I in the secret-key and the ciphertexts, respectively, and to compute the inner product $\langle \vec{v}, I \rangle$ in the exponent of a bilinear map computation.
- The subgroup \mathbb{G}_1 serves a dual purpose in our scheme. On the one hand, it has the effect of *masking* the inner product computation in \mathbb{G}_2 , and preventing the adversary from improperly manipulating the computation in any way to reveal information about the underlying attributes. In particular, it is pivotal in ensuring the non-malleability of the secret-keys and ciphertexts generated by the scheme. On the other hand, it is in the \mathbb{G}_1 subgroup that we incorporate our *encrypt-augment-recover* methodology to achieve computational function privacy.
- The subgroup \mathbb{G}_3 serves as an additional layer of masking for the other subgroups. In particular, random elements sampled from \mathbb{G}_3 are multiplied with various components in both the secret-keys as well as the ciphertexts to hide possible information leakages from the subgroups \mathbb{G}_1 and \mathbb{G}_2 .

Encrypt-Augment-Recover. We apply our *encrypt-augment-recover* approach to the zero-IPE scheme of Katz, Waters and Sahai to achieve computational function privacy. At the core of our approach is the public-key encryption algorithm $\text{PKE} = (\text{KeyGen}, \text{Enc}, \text{Dec})$ that is CPA-secure under the $(k+1)$ -DLIN assumption (the reader is referred to Section 4 for recalling the PKE scheme). The PKE essentially operates in the subgroup \mathbb{G}_1 of prime order q_1 , and its outputs are suitably masked before being incorporated in our scheme. We modify the algorithms of the original scheme as follows:

- The modified setup algorithm runs $(SK, PK) \xleftarrow{R} \text{PKE.KeyGen}$. It incorporates PK in the master secret-key of the original scheme, and modifies the public parameter to include a one way function of SK .
- The original zero-IPE scheme of Katz, Sahai and Waters comprises of secret-keys of the form $(d_0, \{d_{1,i}, d_{2,i}\}_{i \in [1,n]})$. The modified key-generation algorithm in our scheme generates a secret-key of the form $\text{sk}_{\vec{v}} = (d_0, \{d_{1,i}^j, d_{2,i}^j\}_{i \in [1,n], j \in [0, k+2]})$ such that $(\{d_{1,i}^j\}_{j \in [0, k+2]}) = \text{PKE.Enc}(PK, d_{1,i})$ and $(\{d_{2,i}^j\}_{j \in [0, k+2]}) = \text{PKE.Enc}(PK, d_{2,i})$ for $i \in [1, n]$ (along with suitable masking as necessary). Observe that this naturally ensures that each component of the modified secret-key is independent and identically distributed. In the proof of function privacy, we argue the indistinguishability of a well-formed secret-key component from a uniformly random one by relating it to the hardness of solving a $(k+1)$ -DLIN instance in \mathbb{G}_1 .
- The modified encryption algorithm generates an *augmented* ciphertext that retains the ciphertext of the original scheme unaltered as one of its components. The additional ciphertext components are used by the modified decryption algorithm subsequently to remove the effect of PKE and recover the payload message M . The additional components are also in the group \mathbb{G}_1 , and are suitably masked using uniformly random elements from \mathbb{G}_3 . The masking ensures that the data privacy guarantees of the original scheme are not weakened.

Construction Details. We now present the construction for Π_k^{IPE} in details.

- **Setup:** The setup algorithm samples $(\mathbb{G}, \mathbb{G}_T, q_1, q_2, q_3, g_1, g_2, g_3, \hat{e}) \xleftarrow{R} \text{GroupGen}'(1^\lambda)$. It also samples $\{x_{1,j}, x_{2,j} \xleftarrow{R} \mathbb{Z}_{q_1}^*\}_{j \in [1, k+2]}$, $\{g_{1,j}, g_{2,j} \xleftarrow{R} \mathbb{G}_1\}_{j \in [1, k+2]}$, $\{h_{1,i}, h_{2,i} \xleftarrow{R} \mathbb{G}_1\}_{i \in [1, n]}$ and $\{R_{1,i}^j, R_{2,i}^j \xleftarrow{R} \mathbb{G}_3\}_{i \in [1, n], j \in [0, k+2]}$. It additionally samples $h \xleftarrow{R} \mathbb{G}_1$, $\gamma \xleftarrow{R} \mathbb{Z}_{q_1}^*$ and $R_3 \xleftarrow{R} \mathbb{G}_3$, and sets:

$$\begin{aligned} Q &= g_2 \cdot R_3 \\ S_{1,i}^0 &= h_{1,i} \cdot R_{1,i}^0, \quad S_{2,i}^0 = h_{2,i} \cdot R_{2,i}^0 \text{ for } i \in [1, n] \\ S_{1,i}^j &= h_{1,i}^{x_{1,j}} \cdot R_{1,i}^j, \quad S_{2,i}^j = h_{2,i}^{x_{2,j}} \cdot R_{2,i}^j \text{ for } i \in [1, n], j \in [1, k+2] \end{aligned}$$

It outputs the public parameter pp and the master secret-key msk as:

$$\begin{aligned} \text{pp} &= (g_1, g_3, Q, \{S_{1,i}^j, S_{2,i}^j\}_{i \in [1, n], j \in [0, k+2]}, \hat{e}(g_1, h)^\gamma) \\ \text{msk} &= (q_1, q_2, q_3, g_2, \{g_{1,j}, g_{2,j}\}_{j \in [1, k+2]}, \{g_{1,j}^{x_{1,j}} \cdot g_{1, k+2}^{x_{1, k+2}}, g_{2,j}^{x_{2,j}} \cdot g_{2, k+2}^{x_{2, k+2}}\}_{j \in [1, k+2]}, \{h_{1,i}, h_{2,i}\}_{i \in [1, n]}, h^\gamma) \end{aligned}$$

- **KeyGen:** On input the master secret-key msk and a vector $\vec{v} = (v_1, \dots, v_n)$, the key generation algorithm samples $\{z_{1,i}, z_{2,i} \xleftarrow{R} \mathbb{Z}_{q_1}^*\}_{i \in [1, n]}$, $\{y_{1,i}^j, y_{2,i}^j \xleftarrow{R} \mathbb{Z}_{q_1}^*\}_{i \in [1, n], j \in [1, k+1]}$, $Q_4 \xleftarrow{R} \mathbb{G}_2$, $R_5 \xleftarrow{R} \mathbb{G}_3$ and $f_1, f_2 \xleftarrow{R} \mathbb{Z}_{q_2}^*$. It then sets $d_0 = Q_4 \cdot R_5 / (h^\gamma \cdot \prod_{i=1}^n h_{1,i}^{z_{1,i}} \cdot h_{2,i}^{z_{2,i}})$. It also sets:

$$\begin{aligned} d_{1,i}^0 &= g_1^{z_{1,i}} \cdot g_2^{f_1 \cdot v_i} / \left(\prod_{j=1}^{k+1} (g_{1,j}^{x_{1,j}} \cdot g_{1, k+2}^{x_{1, k+2}})^{y_{1,i}^j} \right) \text{ for } i \in [1, n] \\ d_{2,i}^0 &= g_1^{z_{2,i}} \cdot g_2^{f_2 \cdot v_i} / \left(\prod_{j=1}^{k+1} (g_{2,j}^{x_{2,j}} \cdot g_{2, k+2}^{x_{2, k+2}})^{y_{2,i}^j} \right) \text{ for } i \in [1, n] \end{aligned}$$

Finally, it sets the following additional components:

$$\begin{aligned} d_{1,i}^j &= g_{1,j}^{y_{1,i}^j}, \quad d_{2,i}^j = g_{2,j}^{y_{2,i}^j} \text{ for } i \in [1, n], j \in [1, k+1] \\ d_{1,i}^{k+2} &= g_{1,k+2}^{\sum_{j=1}^{k+1} y_{1,i}^j}, \quad d_{2,i}^{k+2} = g_{2,k+2}^{\sum_{j=1}^{k+1} y_{2,i}^j} \text{ for } i \in [1, n] \end{aligned}$$

and outputs the secret-key $\text{sk}_{\vec{\nu}}$ as:

$$\text{sk}_{\vec{\nu}} = \left(d_0, \{d_{1,i}^j, d_{2,i}^j\}_{i \in [1, n], j \in [0, k+2]} \right)$$

- **Enc:** On input the public parameter pp , an attribute $I = (I_1, \dots, I_n) \in \mathbb{Z}_N^n$ and a payload message $M \in \mathcal{M}$, the encryption algorithm samples $r, \alpha, \beta \xleftarrow{R} \mathbb{Z}_N^*$ and $\{R_{6,i}^j, R_{7,i}^j \xleftarrow{R} \mathbb{G}_3\}_{i \in [1, n], j \in [0, k+2]}$. It then sets $c_0 = g_1^r$. It also sets :

$$\begin{aligned} c_{1,i}^0 &= (S_{1,i}^0)^r \cdot Q^{\alpha \cdot I_i} \cdot R_{6,i}^0, \quad c_{2,i}^0 = (S_{2,i}^0)^r \cdot Q^{\beta \cdot I_i} \cdot R_{7,i}^0 \text{ for } i \in [1, n] \\ c_{1,i}^j &= (S_{1,i}^j)^r \cdot R_{6,i}^j, \quad c_{2,i}^j = (S_{2,i}^j)^r \cdot R_{7,i}^j \text{ for } i \in [1, n], j \in [1, k+2] \end{aligned}$$

Finally, it sets $c_3 = M \cdot (\hat{e}(g_1, h)^\gamma)^r$ and outputs the ciphertext C as:

$$C = \left(c_0, \{c_{1,i}^j, c_{2,i}^j\}_{i \in [1, n], j \in [0, k+2]}, c_3 \right)$$

- **Dec:** On input a ciphertext $C = \left(c_0, \{c_{1,i}^j, c_{2,i}^j\}_{i \in [1, n], j \in [0, k+2]} \right)$ and a secret-key $\text{sk}_{\vec{\nu}} = \left(d_0, \{d_{1,i}^j, d_{2,i}^j\}_{i \in [1, n], j \in [0, k+2]} \right)$, the decryption algorithm computes:

$$M' = c_3 \cdot \hat{e}(d_0, c_0) \cdot \left(\prod_{i=1}^n \prod_{j=0}^{k+2} \hat{e}(d_{1,i}^j, c_{1,i}^j) \cdot \hat{e}(d_{2,i}^j, c_{2,i}^j) \right)$$

If $M' \in \mathcal{M}$, the decryption algorithm outputs M' , else it outputs \perp .

Correctness. To see that correctness holds for our zero-IPE scheme, let C and $\text{sk}_{\vec{v}}$ be as described in Section 5. Then we have:

$$\begin{aligned}
M' &= c_3 \cdot \hat{e}(d_0, c_0) \cdot \left(\prod_{i=1}^n \prod_{j=0}^{k+2} \hat{e}(d_{1,i}^j, c_{1,i}^j) \cdot \hat{e}(d_{2,i}^j, c_{2,i}^j) \right) \\
&= M \cdot (\hat{e}(g_1, h)^\gamma)^r \cdot \hat{e} \left(Q_4 \cdot R_5 / \left(h^\gamma \cdot \prod_{i=1}^n h_{1,i}^{z_{1,i}} \cdot h_{2,i}^{z_{2,i}} \right), g_1^r \right) \\
&\quad \cdot \prod_{i=1}^n \left(\hat{e} \left(g_1^{z_{1,i}} \cdot g_2^{f_1 \cdot v_i} / \left(\prod_{j=1}^{k+1} (g_{1,j}^{x_{1,j}} \cdot g_{1,k+2}^{x_{1,k+2}})^{y_{1,i}^j} \right), (S_{1,i}^0)^r \cdot Q^{\alpha \cdot I_i} \cdot R_{6,i}^0 \right) \right) \\
&\quad \cdot \prod_{i=1}^n \left(\hat{e} \left(g_1^{z_{2,i}} \cdot g_2^{f_2 \cdot v_i} / \left(\prod_{j=1}^{k+1} (g_{2,j}^{x_{2,j}} \cdot g_{2,k+2}^{x_{2,k+2}})^{y_{2,i}^j} \right), (S_{2,i}^0)^r \cdot Q^{\beta \cdot I_i} \cdot R_{7,i}^0 \right) \right) \\
&\quad \cdot \left(\prod_{i=1}^n \prod_{j=1}^{k+1} \hat{e} \left(g_{1,j}^{y_{1,i}^j}, (S_{1,i}^j)^r \cdot R_{6,i}^j \right) \cdot \hat{e} \left(g_{2,j}^{y_{2,i}^j}, (S_{2,i}^j)^r \cdot R_{7,i}^j \right) \right) \\
&\quad \cdot \left(\prod_{i=1}^n \hat{e} \left(g_{1,k+2}^{\sum_{j=1}^{k+1} y_{1,i}^j}, (S_{1,i}^{k+2})^r \cdot R_{6,i}^{k+2} \right) \cdot \hat{e} \left(g_{2,k+2}^{\sum_{j=1}^{k+1} y_{2,i}^j}, (S_{2,i}^{k+2})^r \cdot R_{7,i}^{k+2} \right) \right) \\
&= M \cdot (\hat{e}(g_1, h)^\gamma)^r \cdot \left(\frac{\prod_{i=1}^n \hat{e}(g_1^{z_{1,i}}, h_{1,i}^r) \cdot \hat{e}(g_1^{z_{2,i}}, h_{2,i}^r)}{\hat{e}(h^\gamma, g_1^r) \cdot \hat{e}(\prod_{i=1}^n h_{1,i}^{z_{1,i}} \cdot h_{2,i}^{z_{2,i}}, g_1^r)} \right) \cdot \left(\prod_{i=1}^n \hat{e}(g_2^{f_1 \cdot v_i}, g_2^{\alpha \cdot I_i}) \cdot \hat{e}(g_2^{f_2 \cdot v_i}, g_2^{\beta \cdot I_i}) \right) \\
&\quad \cdot \prod_{i=1}^n \left(\frac{\prod_{j=1}^{k+1} \left(\hat{e}(g_{1,j}^{y_{1,i}^j}, (h_{1,i}^{x_{1,j}})^r) \cdot \hat{e}(g_{2,j}^{y_{2,i}^j}, (h_{2,i}^{x_{2,j}})^r) \right)}{\hat{e} \left(\prod_{j=1}^{k+1} (g_{1,j}^{x_{1,j}} \cdot g_{1,k+2}^{x_{1,k+2}})^{y_{1,i}^j}, h_{1,i}^r \right) \cdot \hat{e} \left(\prod_{j=1}^{k+1} (g_{2,j}^{x_{2,j}} \cdot g_{2,k+2}^{x_{2,k+2}})^{y_{2,i}^j}, h_{2,i}^r \right)} \right) \\
&= M \cdot \prod_{i=1}^n \hat{e}(g_2, g_2)^{(\alpha f_1 + \beta f_2) \cdot v_i \cdot I_i} \\
&= M \cdot \hat{e}(g_2, g_2)^{(\alpha f_1 + \beta f_2 \bmod q_2) \cdot \langle \vec{v}, I \rangle}
\end{aligned}$$

where α, β are uniformly random in \mathbb{Z}_N^* and f_1, f_2 are uniformly random in $\mathbb{Z}_{q_2}^*$. If $\langle \vec{v}, I \rangle = 0 \bmod N$, then we have $M' = M$. If $\langle \vec{v}, I \rangle \neq 0 \bmod N$, there are two cases: if $\langle \vec{v}, I \rangle \neq 0 \bmod q_2$, then with all but negligible probability (over random choice of α, β, f_1, f_2), M' does not lie in \mathbb{G}_T (since \mathcal{M} is a small subset of \mathbb{G}_T). Otherwise, we have $\langle \vec{v}, I \rangle = 0 \bmod q_2$, in which case M' will always be equal to M ; however, this would reveal a non-trivial factor of N , and so this too occurs with negligible probability. In fact, the data privacy property of this zero-IPE construction relies on a set of assumptions in bilinear groups that imply the hardness of finding a non-trivial factor of N .

The Need for Two Sub-Systems. Note that our zero-IPE scheme uses two parallel sub-systems (in the key generation and encryption algorithms) that are apparently redundant since they perform the same functions. Indeed, our scheme inherits this feature from the original zero-IPE scheme of Katz, Sahai and Waters [3]. Eliminating one of the sub-systems from our scheme would retain functional correctness as well as computational function privacy, while also improving performance and efficiency. However, the proof methodology in [3] for data privacy relies on the existence of the parallel sub-systems in an essential way. Since our aim is to retain the same data privacy guarantees as in the original scheme, we stick to the use of two parallel sub-systems in our augmented zero-IPE scheme.

5.1 Security of Our IPE Scheme

Data Privacy. We state the following theorem for the data privacy of Π_k^{IPE} :

Theorem 5.1 *Our zero-IPE scheme Π_k^{IPE} retains the selective data privacy guarantees of the original zero-IPE scheme of Katz, Sahai and Waters [3].*

Proof Overview. We provide a brief overview of the proof technique for our scheme, which essentially follows the proof technique presented in [3]. We consider a probabilistic polynomial-time adversary that tries to determine whether the *challenge ciphertext* is associated with either of the two attributes I_0 or I_1 . The proof proceeds via a sequence of hybrid games in which an entire attribute used in the challenge ciphertext is changed in one step, instead of changing them component by component for reasons mentioned in the proof of the original scheme in [3]. This is facilitated by the presence of the two parallel sub-systems, which allows the hybrid games to use *ill-formed* ciphertexts that are encrypted with respect to two *different* attributes I and I' and in the two sub-systems. Let such a ciphertext be denoted informally as (I, I') . The proof establishes indistinguishability between the well-formed ciphertexts (I_0, I_0) and (I_1, I_1) via a sequence of intermediate hybrid games using the ill-formed ciphertexts $(I_0, \vec{0})$, (I_0, I_1) and $(\vec{0}, I_1)$. The zero vector is used since it is orthogonal to any other vector. The simulator in our proof works in one sub-system independent of what happens in the other one. In each hybrid game, the simulator embeds a subgroup-decision like assumption in the challenge ciphertext, and the structure of the challenge determines whether a sub-system embeds a given vector or a zero vector. This is essentially an adoption of the proof technique originally presented in [3] to our function private scheme.

An additional requirement in our proof is that the simulator should be able to embed the $(k + 1)$ -DLIN instances when responding to the key generation queries from the adversary. As demonstrated in the proof of Theorem 4.1, this is straightforward to achieve: since the simulator in the data privacy game is allowed to set up the $(k + 1)$ -DLIN instances entirely on its own, it can easily augment the secret-key generation process in the proof of the original scheme by appropriately embedding these instances where necessary. Moreover, since the $(k + 1)$ -DLIN instances are sampled uniformly at random, the resulting distribution of secret-keys is exactly as in the real world from the point of view of the adversary. Similarly, in the challenge phase, the simulator generates the additional components in the augmented ciphertext uniformly at random, without altering the nature of the ciphertext distribution from the adversary's point of view.

Computational Function Privacy. We state the following theorem for the computational function privacy of Π_k :

Theorem 5.2 *Our zero-IPE scheme Π_k^{IPE} is computationally function private under the $(k + 1)$ -DLIN assumption for predicate vectors sampled uniformly from (n, k) -block sources with $k = \omega(\log \lambda)$.*

Proof. The proof aims to show that any probabilistic poly-time adversary \mathcal{A} cannot distinguish between the real and random modes of operation of the function privacy oracle, provided that the oracle is queried with circuits that sample sufficiently unpredictable distributions over the space of predicates. In particular, such distributions should be (n, k) -block sources over \mathbb{Z}_N^n , such that each component of a vector \vec{v} sampled from an adversarially chosen distribution has a min-entropy of $k = \omega(\log \lambda)$, and is uncorrelated with all other components. We define a series of hybrid experiments $\text{Expt}_{\text{FP}, \Pi_k^{IPE}, \mathcal{A}}^{\text{mode}, m}(\lambda)$ for $\text{mode} \in \{\text{real}, \text{rand}\}$ and $m \in [0, n]$ as follows:

- $\text{Expt}_{\text{FP}, \Pi_k^{IPE}, \mathcal{A}}^{\text{mode}, 0}(\lambda)$ is exactly identical to $\text{Expt}_{\text{FP}, \Pi_k^{IPE}, \mathcal{A}}^{\text{mode}}(\lambda)$.
- $\text{Expt}_{\text{FP}, \Pi_k^{IPE}, \mathcal{A}}^{\text{mode}, m}(\lambda)$ for $m \in [1, n]$ is identical to $\text{Expt}_{\text{FP}, \Pi_k^{IPE}, \mathcal{A}}^{\text{mode}}(\lambda)$ except that the secret-key $\text{sk}_{\vec{v}^*} = (d_0^*, \{d_{1,i}^{*j}, d_{2,i}^{*j}\}_{i \in [1, n], j \in [0, k+2]})$ generated by the real-or-random oracle is such that the set of components $\{d_{1,i}^{*j}, d_{2,i}^{*j}\}_{i \in [1, m], j \in [0, k+2]}$ are uniformly random and independent of the underlying vector \vec{v}^* .

Quite evidently, the following holds:

$$\left| \Pr \left[\text{Expt}_{\text{FP}, \Pi_k^{\text{IPE}}, \mathcal{A}}^{\text{real}, n}(\lambda) = 1 \right] - \Pr \left[\text{Expt}_{\text{FP}, \Pi_k^{\text{IPE}}, \mathcal{A}}^{\text{rand}, n}(\lambda) = 1 \right] \right| = 0 \quad (1)$$

We now state and prove the following claim:

Claim 5.1 *For any probabilistic polynomial-time adversary \mathcal{A} , for $\text{mode} \in \{\text{real}, \text{rand}\}$ and for $m \in [0, n-1]$, the following holds:*

$$\left| \Pr \left[\text{Expt}_{\text{FP}, \Pi_k^{\text{IPE}}, \mathcal{A}}^{\text{mode}, m}(\lambda) = 1 \right] - \Pr \left[\text{Expt}_{\text{FP}, \Pi_k^{\text{IPE}}, \mathcal{A}}^{\text{mode}, m+1}(\lambda) = 1 \right] \right| \leq \text{negl}(\lambda)$$

To prove this claim, we assume the contrary. Let \mathcal{A} be a probabilistic polynomial-time adversary such that:

$$\left| \Pr \left[\text{Expt}_{\text{FP}, \Pi_k^{\text{IPE}}, \mathcal{A}}^{\text{mode}, m}(\lambda) = 1 \right] - \Pr \left[\text{Expt}_{\text{FP}, \Pi_k^{\text{IPE}}, \mathcal{A}}^{\text{mode}, m+1}(\lambda) = 1 \right] \right| = \epsilon > \text{negl}(\lambda)$$

for some $m \in [0, n-1]$. We construct an algorithm \mathcal{B} such that:

$$\left| \Pr \left[\mathcal{B} \left(\left(g_{1,1}, \dots, g_{1,k+2}, g_{1,1}^{a_1}, \dots, g_{1,k+2}^{a_1} \right), \left(g_{2,1}, \dots, g_{2,k+2}, g_{2,1}^{a'_1}, \dots, g_{2,k+2}^{a'_1} \right) \right) = 1 \right] - \Pr \left[\mathcal{B} \left(\left(g_{1,1}, \dots, g_{1,k+2}, g_{1,1}^{a_1}, \dots, g_{1,k+2}^{a_{k+2}} \right), \left(g_{2,1}, \dots, g_{2,k+2}, g_{2,1}^{a'_1}, \dots, g_{2,k+2}^{a'_{k+2}} \right) \right) = 1 \right] \right| = \epsilon$$

where the probability is over random choice of $\{a_j, a'_j \xleftarrow{R} \mathbb{Z}_{q_1}^*\}_{j \in [1, k+2]}$, and over random choice of $\{g_{1,j}, g_{2,j} \xleftarrow{R} \mathbb{G}_1\}_{j \in [1, k+2]}$ (\mathbb{G}_1 being a group of prime order q_1). Observe that \mathcal{B} can in turn be trivially used to construct another algorithm that has advantage at least ϵ in solving a given instance of the $(k+1)$ -DLIN problem. \mathcal{B} interacts with \mathcal{A} as follows:

- **Setup:** \mathcal{B} uniformly samples two other λ -bit primes q_2, q_3 apart from q_1 which is the order of the group \mathbb{G}_1 in its input instance. It also sets $N = q_1 q_2 q_3$, and then sets up the groups \mathbb{G}, \mathbb{G}_2 and \mathbb{G}_3 of order N, q_2 and q_3 respectively, along with $\hat{e}: \mathbb{G} \times \mathbb{G} \rightarrow \mathbb{G}_T$. It then samples $\{x_{1,j}, x_{2,j} \xleftarrow{R} \mathbb{Z}_{q_1}^*\}_{j \in [1, k+2]}$, $\{h_{1,i}, h_{2,i} \xleftarrow{R} \mathbb{G}_1\}_{i \in [1, n]}$ and $\{R_{1,i}^j, R_{2,i}^j \xleftarrow{R} \mathbb{G}_3\}_{i \in [1, n], j \in [0, k+2]}$. It additionally samples $h \xleftarrow{R} \mathbb{G}_1, \gamma \xleftarrow{R} \mathbb{Z}_{q_1}^*$ and $R_3 \xleftarrow{R} \mathbb{G}_3$, and sets:

$$\begin{aligned} Q &= g_2 \cdot R_3 \\ S_{1,i}^0 &= h_{1,i} \cdot R_{1,i}^0, \quad S_{2,i}^0 = h_{2,i} \cdot R_{2,i}^0 \text{ for } i \in [1, n] \\ S_{1,i}^j &= h_{1,i}^{x_{1,j}} \cdot R_{1,i}^j, \quad S_{2,i}^j = h_{2,i}^{x_{2,j}} \cdot R_{2,i}^j \text{ for } i \in [1, n], j \in [1, k+2] \end{aligned}$$

Finally, it sets the public parameter pp as:

$$\text{pp} = \left(g_1, g_3, Q, \{S_{1,i}^j, S_{2,i}^j\}_{i \in [1, n], j \in [0, k+2]}, \hat{e}(g_1, h)^\gamma \right)$$

and provides the same to \mathcal{A} . Observe that pp is distributed exactly as in the real world.

- **Secret-Key Queries:** When \mathcal{A} issues a secret-key query for $\vec{v} \in \mathbb{Z}_N^n$, \mathcal{B} samples $\{z_{1,i}, z_{2,i} \xleftarrow{R} \mathbb{Z}_{q_1}^*\}_{i \in [1, n]}$, $\{y_{1,i}^j, y_{2,i}^j \xleftarrow{R} \mathbb{Z}_{q_1}^*\}_{i \in [1, n], j \in [1, k+1]}$, $Q_4 \xleftarrow{R} \mathbb{G}_2, R_5 \xleftarrow{R} \mathbb{G}_3$ and $f_1, f_2 \xleftarrow{R} \mathbb{Z}_{q_2}^*$. It then sets:

$$\begin{aligned} d_0 &= Q_4 \cdot R_5 / \left(\prod_{i=1}^n h_{1,i}^{z_{1,i}} \cdot h_{2,i}^{z_{2,i}} \right) \\ d_{1,i}^0 &= g_1^{z_{1,i}} \cdot g_2^{f_1 \cdot v_i} / \left(\prod_{j=1}^{k+1} \left(g_{1,j}^{x_{1,j}} \cdot g_{1,k+2}^{x_{1,k+2}} \right)^{y_{1,i}^j} \right) \text{ for } i \in [1, n] \\ d_{2,i}^0 &= g_1^{z_{2,i}} \cdot g_2^{f_2 \cdot v_i} / \left(\prod_{j=1}^{k+1} \left(g_{2,j}^{x_{2,j}} \cdot g_{2,k+2}^{x_{2,k+2}} \right)^{y_{2,i}^j} \right) \text{ for } i \in [1, n] \end{aligned}$$

Finally, it sets the following additional components:

$$\begin{aligned} d_{1,i}^j &= g_{1,j}^{y_{1,i}^j} \quad , \quad d_{2,i}^j = g_{2,j}^{y_{2,i}^j} \quad \text{for } i \in [1, n], j \in [1, k+1] \\ d_{1,i}^{k+2} &= g_{1,k+2}^{\sum_{j=1}^{k+1} y_{1,i}^j} \quad , \quad d_{2,i}^{k+2} = g_{2,k+2}^{\sum_{j=1}^{k+1} y_{2,i}^j} \quad \text{for } i \in [1, n] \end{aligned}$$

and outputs the secret-key $\text{sk}_{\vec{v}}$ as:

$$\text{sk}_{\vec{v}} = \left(d_0, \{d_{1,i}^j, d_{2,i}^j\}_{i \in [1, n], j \in [0, k+2]} \right)$$

- **Real-or-Random Query:** Suppose \mathcal{A} queries the real-or-random oracle with an (n, k) -block source $\mathbf{V}^* = (V_1^*, \dots, V_n^*)$ over \mathbb{Z}_N^n such that $k = \omega(\log \lambda)$. \mathcal{B} samples mode $\xleftarrow{R} \{\text{real}, \text{rand}\}$ and does the following:

1. For each $i \in [1, n]$, \mathcal{B} samples $v_i^* \xleftarrow{R} V_i^*$ if mode = real, or $v_i^* \xleftarrow{R} \mathbb{Z}_N$ if mode = rand. The vector $\vec{v}^* = (v_1^*, \dots, v_n^*)$ is the challenge vector that \mathcal{B} uses to respond to the query from \mathcal{A} .
2. As in the response to the secret-key queries, \mathcal{B} samples $\{z_{1,i}, z_{2,i} \xleftarrow{R} \mathbb{Z}_{q_1}^*\}_{i \in [1, n]}$, $\{y_{1,i}^j, y_{2,i}^j \xleftarrow{R} \mathbb{Z}_{q_1}^*\}_{i \in [1, n] \setminus \{m+1\}, j \in [1, k+1]}$, $Q_4 \xleftarrow{R} \mathbb{G}_2$, $R_5 \xleftarrow{R} \mathbb{G}_3$ and $f_1, f_2 \xleftarrow{R} \mathbb{Z}_{q_2}^*$. It also samples $\{y_{1,i}^{k+2}, y_{2,i}^{k+2} \xleftarrow{R} \mathbb{Z}_{q_1}^*\}_{i \in [1, m]}$. Observe that \mathcal{B} does not sample $y_{1,m+1}^j$ or $y_{2,m+1}^j$ for $j \in [1, k+2]$. This is because \mathcal{B} embeds the its input pair of $(k+1)$ -DLIN instances in its response by formally setting $y_{1,m+1}^j = a_j$ and $y_{2,m+1}^j = a'_j$ for each $j \in [1, k+2]$ as described next.
3. \mathcal{B} now sets the various components of the secret-key as follows:

$$\begin{aligned} d_0^* &= Q_4 \cdot R_5 / \left(\prod_{i=1}^n h_{1,i}^{z_{1,i}} \cdot h_{2,i}^{z_{2,i}} \right) \\ d_{1,i}^{*0} &= g_1^{z_{1,i}} \cdot g_2^{f_1 \cdot v_i} / \left(\prod_{j=1}^{k+1} \left(g_{1,j}^{x_{1,j}} \cdot g_{1,k+2}^{x_{1,k+2}} \right)^{y_{1,i}^j} \right) \quad \text{for } i \in [1, n] \setminus \{m+1\} \\ d_{2,i}^{*0} &= g_1^{z_{2,i}} \cdot g_2^{f_2 \cdot v_i} / \left(\prod_{j=1}^{k+1} \left(g_{2,j}^{x_{2,j}} \cdot g_{2,k+2}^{x_{2,k+2}} \right)^{y_{2,i}^j} \right) \quad \text{for } i \in [1, n] \setminus \{m+1\} \end{aligned}$$

It also sets the following additional components:

$$\begin{aligned} d_{1,i}^{*j} &= g_{1,j}^{y_{1,i}^j} \quad , \quad d_{2,i}^{*j} = g_{2,j}^{y_{2,i}^j} \quad \text{for } i \in [1, n] \setminus \{m+1\}, j \in [1, k+1] \\ d_{1,i}^{*k+2} &= g_{1,k+2}^{y_{1,i}^{k+2}} \quad , \quad d_{2,i}^{*k+2} = g_{2,k+2}^{y_{2,i}^{k+2}} \quad \text{for } i \in [1, m] \\ d_{1,i}^{*k+2} &= g_{1,k+2}^{\sum_{j=1}^{k+1} y_{1,i}^j} \quad , \quad d_{2,i}^{*k+2} = g_{2,k+2}^{\sum_{j=1}^{k+1} y_{2,i}^j} \quad \text{for } i \in [m+2, n] \end{aligned}$$

Observe that the first m components of the secret-key are crafted to be uniformly random, while the last $(n - m - 1)$ components are well-formed.

4. Finally, \mathcal{B} embeds its input pair of $(k+1)$ -DLIN instances in the $(m+1)^{\text{th}}$ component of the secret-key by setting:

$$\begin{aligned} d_{1,m+1}^{*0} &= g_1^{z_{1,m+1}} \cdot g_1^{f_1 \cdot v_{m+1}} / \left(\prod_{j=1}^{k+2} \left(g_{1,j}^{a_j} \right)^{x_{1,j}} \right) \\ d_{2,m+1}^{*0} &= g_1^{z_{2,m+1}} \cdot g_2^{f_2 \cdot v_{m+1}} / \left(\prod_{j=1}^{k+2} \left(g_{2,j}^{a'_j} \right)^{x_{2,j}} \right) \\ d_{1,m+1}^{*j} &= g_{1,j}^{a_j} \quad , \quad d_{2,m+1}^{*j} = g_{2,j}^{a'_j} \quad \text{for } j \in [1, k+2] \end{aligned}$$

\mathcal{B} responds to \mathcal{A} with the secret-key $\text{sk}_{\vec{v}^*}$ as:

$$\text{sk}_{\vec{v}^*} = \left(d_0^*, \{d_{1,i}^{*j}, d_{2,i}^{*j}\}_{i \in [1,n], j \in [0,k+2]} \right)$$

- **Guess:** At the end of the game, \mathcal{A} outputs a bit b' . \mathcal{B} outputs the same bit b' .

It is easy to see that when $a_{k+2} = \sum_{j=1}^{k+1} a_j$ and $a'_{k+2} = \sum_{j=1}^{k+1} a'_j$, the secret-key $\text{sk}_{\vec{v}^*}$ is identically distributed to the response of the real-or-random oracle in the experiment $\text{Expt}_{\text{FP}, \Pi_k^{\text{IPE}}, \mathcal{A}}^{\text{mode}, m}(\lambda)$. On the other hand, when either or both of a_{k+2} and a'_{k+2} are uniformly random in \mathbb{Z}_q^* , the secret-key $\text{sk}_{\vec{v}^*}$ is identically distributed to the response of the real-or-random oracle in the experiment $\text{Expt}_{\text{FP}, \Pi_k^{\text{IPE}}, \mathcal{A}}^{\text{mode}, m+1}(\lambda)$. It follows readily that \mathcal{B} has the same advantage ϵ as \mathcal{A} in solving its input instance pair. This completes the proof of Claim 5.1.

We now make the following observation:

$$\begin{aligned} & \left| \Pr \left[\text{Expt}_{\text{FP}, \Pi_k^{\text{IPE}}, \mathcal{A}}^{\text{mode}}(\lambda) = 1 \right] - \Pr \left[\text{Expt}_{\text{FP}, \Pi_k^{\text{IPE}}, \mathcal{A}}^{\text{mode}, n}(\lambda) = 1 \right] \right| \\ & \leq \sum_{m=0}^{n-1} \left| \Pr \left[\text{Expt}_{\text{FP}, \Pi_k^{\text{IPE}}, \mathcal{A}}^{\text{mode}, m}(\lambda) = 1 \right] - \Pr \left[\text{Expt}_{\text{FP}, \Pi_k^{\text{IPE}}, \mathcal{A}}^{\text{mode}, m+1}(\lambda) = 1 \right] \right| \\ & \leq \text{negl}(\lambda) \text{ (from Claim 5.1) for } n = \text{poly}(\lambda) \end{aligned}$$

Consequently, for $\text{mode} \stackrel{R}{\leftarrow} \{\text{real}, \text{rand}\}$, we have:

$$\begin{aligned} \text{Adv}_{\Pi_k^{\text{IPE}}, \mathcal{A}}^{\text{FP}}(\lambda) &= \left| \Pr \left[\text{Expt}_{\text{FP}, \Pi_k^{\text{IPE}}, \mathcal{A}}^{\text{real}}(\lambda) = 1 \right] - \Pr \left[\text{Expt}_{\text{FP}, \Pi_k^{\text{IPE}}, \mathcal{A}}^{\text{rand}}(\lambda) = 1 \right] \right| \\ &\leq \left| \Pr \left[\text{Expt}_{\text{FP}, \Pi_k^{\text{IPE}}, \mathcal{A}}^{\text{real}}(\lambda) = 1 \right] - \Pr \left[\text{Expt}_{\text{FP}, \Pi_k^{\text{IPE}}, \mathcal{A}}^{\text{real}, n}(\lambda) = 1 \right] \right| \\ &\quad + \left| \Pr \left[\text{Expt}_{\text{FP}, \Pi_k^{\text{IPE}}, \mathcal{A}}^{\text{rand}}(\lambda) = 1 \right] - \Pr \left[\text{Expt}_{\text{FP}, \Pi_k^{\text{IPE}}, \mathcal{A}}^{\text{rand}, n}(\lambda) = 1 \right] \right| \\ &\quad + \left| \Pr \left[\text{Expt}_{\text{FP}, \Pi_k^{\text{IPE}}, \mathcal{A}}^{\text{real}, n}(\lambda) = 1 \right] - \Pr \left[\text{Expt}_{\text{FP}, \Pi_k^{\text{IPE}}, \mathcal{A}}^{\text{rand}, n}(\lambda) = 1 \right] \right| \\ &\leq 2 \left| \Pr \left[\text{Expt}_{\text{FP}, \Pi_k^{\text{IPE}}, \mathcal{A}}^{\text{mode}}(\lambda) = 1 \right] - \Pr \left[\text{Expt}_{\text{FP}, \Pi_k^{\text{IPE}}, \mathcal{A}}^{\text{mode}, n}(\lambda) = 1 \right] \right| \text{ (from Equation 1)} \\ &= 2\epsilon' \leq \text{negl}(\lambda) \end{aligned}$$

This completes the proof of function privacy for Π_k^{IPE} . Note that the proof does not essentially rely on the presence of two parallel sub-systems in our scheme. In fact, eliminating one of the subsystems gives a simpler proof. However, the two sub-system version allows us to adopt the proof of data privacy in [3] for our scheme.

Enhanced Computational Function Privacy. We state the following theorem for the enhanced computational function privacy of Π_k :

Theorem 5.3 *Our zero-IPE scheme Π_k^{IPE} is computationally enhanced function private under the $(k+1)$ -DLIN assumption for predicate vectors sampled uniformly from (n, k) -block sources with $k = \omega(\log \lambda)$.*

Proof Overview. Once again, the proof is very similar to that of Theorem 5.2; we begin by providing a proof overview here. The proof basically shows that the simulator \mathcal{B} can additionally simulate the function privacy encryption oracle, and that the real and random modes of operation of the function privacy oracle are indistinguishable even in the presence of the encryption oracle. We again consider a series of hybrid experiments that gradually randomize the secret-key $\mathbf{sk}_{\vec{v}^*}$ generated by the real-or-random function privacy oracle component by component. The challenge is to simulate the function privacy encryption oracle such that its behavior does not provide the adversary with any additional distinguisher between two consecutive experiments.

Let $\mathbf{sk}_{\vec{v}^*} = \left(d_0^*, \{d_{1,i}^{*j}, d_{2,i}^{*j}\}_{i \in [1,n], j \in [0,k+2]} \right)$ be the response of the real-or-random function privacy oracle to the adversary \mathcal{A} . On receiving a function-privacy encryption oracle query for a payload message M , \mathcal{B} responds as follows: it samples $I^* = (I_1^*, \dots, I_n^*)$ such that $\langle \vec{v}^*, I^* \rangle = 0$. It then samples $r, \alpha, \beta \xleftarrow{R} \mathbb{Z}_N^*$ and $\{R_{6,i}^j, R_{7,i}^j\} \xleftarrow{R} \mathbb{G}_3\}_{i \in [1,n], j \in [0,k+2]}$ and sets $c_0^* = g_1^r$. It also sets:

$$\begin{aligned} c_{1,i}^{*0} &= (S_{1,i}^0)^r \cdot Q^{\alpha \cdot I_i^*} \cdot R_{6,i}^0, & c_{2,i}^{*0} &= (S_{2,i}^0)^r \cdot Q^{\beta \cdot I_i^*} \cdot R_{7,i}^0 \text{ for } i \in [1, n] \\ c_{1,i}^{*j} &= (S_{1,i}^j)^r \cdot R_{6,i}^j, & c_{2,i}^{*j} &= (S_{2,i}^j)^r \cdot R_{7,i}^j \text{ for } i \in [1, n], j \in [1, k+2] \end{aligned}$$

Finally, it sets:

$$c_3^* = M / \hat{e}(d_0^*, c_0^*) \cdot \left(\prod_{i=1}^n \prod_{j=0}^{k+2} \hat{e}(d_{1,i}^{*j}, c_{1,i}^{*j}) \cdot \hat{e}(d_{2,i}^{*j}, c_{2,i}^{*j}) \right)$$

and outputs the ciphertext $C^* = \left(c_0^*, \{c_{1,i}^{*j}, c_{2,i}^{*j}\}_{i \in [1,n], j \in [0,k+2]}, c_3^* \right)$.

Clearly, the ciphertext C^* is dependent solely on $\mathbf{sk}_{\vec{v}^*}$ and the input message M , and is hence independent of the vector \vec{v}^* . In addition, it produces M upon decryption using $\mathbf{sk}_{\vec{v}^*}$ irrespective of \mathcal{B} 's input challenge. Consequently, the distribution of C^* in the two experiments is indistinguishable from \mathcal{A} 's point of view. This in turn implies that the adversary's advantage in distinguishing between two hybrid experiments is the same as \mathcal{B} 's advantage in solving its input $(k+1)$ -DLIN challenge pair.

Detailed Proof. We now present the detailed proof. The proof aims to show that any probabilistic poly-time adversary \mathcal{A} cannot distinguish between the real and random modes of operation of the function privacy oracle, provided that the oracle is queried with circuits that sample sufficiently unpredictable distributions over the space of predicates. In particular, such distributions should be (n, k) -block sources over \mathbb{Z}_N^n , such that each component of a vector \vec{v} sampled from an adversarially chosen distribution has a min-entropy of $k = \omega(\log \lambda)$, and is uncorrelated with all other components. We define a series of hybrid experiments $\text{Expt}_{\text{EFP}, \Pi_k^{\text{IPE}}, \mathcal{A}}^{\text{mode}, m}(\lambda)$ for $\text{mode} \in \{\text{real}, \text{rand}\}$ and $m \in [0, n]$ as follows:

- $\text{Expt}_{\text{EFP}, \Pi_k^{\text{IPE}}, \mathcal{A}}^{\text{mode}, 0}(\lambda)$ is exactly identical to $\text{Expt}_{\text{EFP}, \Pi_k^{\text{IPE}}, \mathcal{A}}^{\text{mode}}(\lambda)$.
- $\text{Expt}_{\text{EFP}, \Pi_k^{\text{IPE}}, \mathcal{A}}^{\text{mode}, m}(\lambda)$ for $m \in [1, n]$ is identical to $\text{Expt}_{\text{EFP}, \Pi_k^{\text{IPE}}, \mathcal{A}}^{\text{mode}}(\lambda)$ except that the secret-key $\mathbf{sk}_{\vec{v}^*} = \left(d_0^*, \{d_{1,i}^{*j}, d_{2,i}^{*j}\}_{i \in [1,n], j \in [0,k+2]} \right)$ generated by the real-or-random oracle is such that the set of components $\{d_{1,i}^{*j}, d_{2,i}^{*j}\}_{i \in [1,m], j \in [0,k+2]}$ are uniformly random and independent of the underlying vector \vec{v}^* . In addition, the ciphertext C^* for a message M generated by the function privacy encryption oracle is uniformly random and independent of \vec{v}^* ; it, however, produces M upon decryption using the $\mathbf{sk}_{\vec{v}^*}$ generated by the real-or-random oracle.

Quite evidently, the following holds:

$$\left| \Pr \left[\text{Expt}_{\text{EFP}, \Pi_k^{\text{IPE}}, \mathcal{A}}^{\text{real}, n}(\lambda) = 1 \right] - \Pr \left[\text{Expt}_{\text{EFP}, \Pi_k^{\text{IPE}}, \mathcal{A}}^{\text{rand}, n}(\lambda) = 1 \right] \right| = 0 \quad (2)$$

We now state and prove the following claim:

Claim 5.2 For any probabilistic polynomial-time adversary \mathcal{A} , for $\text{mode} \in \{\text{real}, \text{rand}\}$ and for $m \in [0, n-1]$, the following holds:

$$\left| \Pr \left[\text{Expt}_{\text{EFP}, \Pi_k^{\text{IPE}}, \mathcal{A}}^{\text{mode}, m}(\lambda) = 1 \right] - \Pr \left[\text{Expt}_{\text{EFP}, \Pi_k^{\text{IPE}}, \mathcal{A}}^{\text{mode}, m+1}(\lambda) = 1 \right] \right| \leq \text{negl}(\lambda)$$

To prove this claim, we assume the contrary. Let \mathcal{A} be a probabilistic polynomial-time adversary such that:

$$\left| \Pr \left[\text{Expt}_{\text{EFP}, \Pi_k^{\text{IPE}}, \mathcal{A}}^{\text{mode}, m}(\lambda) = 1 \right] - \Pr \left[\text{Expt}_{\text{EFP}, \Pi_k^{\text{IPE}}, \mathcal{A}}^{\text{mode}, m+1}(\lambda) = 1 \right] \right| = \epsilon > \text{negl}(\lambda)$$

for some $m \in [0, n-1]$. We construct an algorithm \mathcal{B} such that:

$$\left| \Pr \left[\mathcal{B} \left(\left(g_{1,1}, \dots, g_{1,k+2}, g_{1,1}^{a_1}, \dots, g_{1,k+2}^{\sum_{j=1}^{k+1} a_j} \right), \left(g_{2,1}, \dots, g_{2,k+2}, g_{2,1}^{a'_1}, \dots, g_{2,k+2}^{\sum_{j=1}^{k+1} a'_j} \right) \right) = 1 \right] - \Pr \left[\mathcal{B} \left(\left(g_{1,1}, \dots, g_{1,k+2}, g_{1,1}^{a_1}, \dots, g_{1,k+2}^{a_{k+2}} \right), \left(g_{2,1}, \dots, g_{2,k+2}, g_{2,1}^{a'_1}, \dots, g_{2,k+2}^{a'_{k+2}} \right) \right) = 1 \right] \right| = \epsilon$$

where the probability is over random choice of $\{a_j, a'_j \xleftarrow{R} \mathbb{Z}_{q_1}^*\}_{j \in [1, k+2]}$, and over random choice of $\{g_{1,j}, g_{2,j} \xleftarrow{R} \mathbb{G}_1\}_{j \in [1, k+2]}$ (\mathbb{G}_1 being a group of prime order q_1). Observe that \mathcal{B} can in turn be trivially used to construct another algorithm that has advantage at least ϵ in solving a given instance of the $(k+1)$ -DLIN problem. \mathcal{B} interacts with \mathcal{A} as follows:

- **Setup:** \mathcal{B} uniformly samples two other λ -bit primes q_2, q_3 apart from q_1 which is the order of the group \mathbb{G}_1 in its input instance. It also sets $N = q_1 q_2 q_3$, and then sets up the groups \mathbb{G}, \mathbb{G}_2 and \mathbb{G}_3 of order N, q_2 and q_3 respectively, along with the bilinear map $\hat{e}: \mathbb{G} \times \mathbb{G} \rightarrow \mathbb{G}_T$. It then samples $\{x_{1,j}, x_{2,j} \xleftarrow{R} \mathbb{Z}_{q_1}^*\}_{j \in [1, k+2]}$, $\{h_{1,i}, h_{2,i} \xleftarrow{R} \mathbb{G}_1\}_{i \in [1, n]}$ and $\{R_{1,i}^j, R_{2,i}^j \xleftarrow{R} \mathbb{G}_3\}_{i \in [1, n], j \in [0, k+2]}$. It additionally samples $h \xleftarrow{R} \mathbb{G}_1, \gamma \xleftarrow{R} \mathbb{Z}_{q_1}^*$ and $R_3 \xleftarrow{R} \mathbb{G}_3$, and sets:

$$\begin{aligned} Q &= g_2 \cdot R_3 \\ S_{1,i}^0 &= h_{1,i} \cdot R_{1,i}^0, S_{2,i}^0 = h_{2,i} \cdot R_{2,i}^0 \text{ for } i \in [1, n] \\ S_{1,i}^j &= h_{1,i}^{x_{1,j}} \cdot R_{1,i}^j, S_{2,i}^j = h_{2,i}^{x_{2,j}} \cdot R_{2,i}^j \text{ for } i \in [1, n], j \in [1, k+2] \end{aligned}$$

Finally, it sets the public parameter pp as:

$$\text{pp} = \left(g_1, g_3, Q, \{S_{1,i}^j, S_{2,i}^j\}_{i \in [1, n], j \in [0, k+2]}, \hat{e}(g_1, h)^\gamma \right)$$

and provides the same to \mathcal{A} . Observe that pp is distributed exactly as in the real world.

- **Secret-Key Queries:** When \mathcal{A} issues a secret-key query for $\vec{v} \in \mathbb{Z}_N^n$, \mathcal{B} samples $\{z_{1,i}, z_{2,i} \xleftarrow{R} \mathbb{Z}_{q_1}^*\}_{i \in [1, n]}$, $\{y_{1,i}^j, y_{2,i}^j \xleftarrow{R} \mathbb{Z}_{q_1}^*\}_{i \in [1, n], j \in [1, k+1]}$, $Q_4 \xleftarrow{R} \mathbb{G}_2, R_5 \xleftarrow{R} \mathbb{G}_3$ and $f_1, f_2 \xleftarrow{R} \mathbb{Z}_{q_2}^*$. It then sets:

$$\begin{aligned} d_0 &= Q_4 \cdot R_5 / \left(\prod_{i=1}^n h_{1,i}^{z_{1,i}} \cdot h_{2,i}^{z_{2,i}} \right) \\ d_{1,i}^0 &= g_1^{z_{1,i}} \cdot g_2^{f_1 \cdot v_i} / \left(\prod_{j=1}^{k+1} \left(g_{1,j}^{x_{1,j}} \cdot g_{1,k+2}^{x_{1,k+2}} \right)^{y_{1,i}^j} \right) \text{ for } i \in [1, n] \\ d_{2,i}^0 &= g_1^{z_{2,i}} \cdot g_2^{f_2 \cdot v_i} / \left(\prod_{j=1}^{k+1} \left(g_{2,j}^{x_{2,j}} \cdot g_{2,k+2}^{x_{2,k+2}} \right)^{y_{2,i}^j} \right) \text{ for } i \in [1, n] \end{aligned}$$

Finally, it sets the following additional components:

$$\begin{aligned} d_{1,i}^j &= g_{1,j}^{y_{1,i}^j} \quad , \quad d_{2,i}^j = g_{2,j}^{y_{2,i}^j} \quad \text{for } i \in [1, n], j \in [1, k+1] \\ d_{1,i}^{k+2} &= g_{1,k+2}^{\sum_{j=1}^{k+1} y_{1,i}^j} \quad , \quad d_{2,i}^{k+2} = g_{2,k+2}^{\sum_{j=1}^{k+1} y_{2,i}^j} \quad \text{for } i \in [1, n] \end{aligned}$$

and outputs the secret-key $\text{sk}_{\vec{v}}$ as:

$$\text{sk}_{\vec{v}} = \left(d_0, \{d_{1,i}^j, d_{2,i}^j\}_{i \in [1, n], j \in [0, k+2]} \right)$$

- **Real-or-Random Query:** Suppose \mathcal{A} queries the real-or-random oracle with an (n, k) -block source $\mathbf{V}^* = (V_1^*, \dots, V_n^*)$ over \mathbb{Z}_N^n such that $k = \omega(\log \lambda)$. \mathcal{B} samples mode $\xleftarrow{R} \{\text{real}, \text{rand}\}$ and does the following:

1. For each $i \in [1, n]$, \mathcal{B} samples $v_i^* \xleftarrow{R} V_i^*$ if mode = real, or $v_i^* \xleftarrow{R} \mathbb{Z}_N$ if mode = rand. The vector $\vec{v}^* = (v_1^*, \dots, v_n^*)$ is the challenge vector that \mathcal{B} uses to respond to the query from \mathcal{A} .
2. As in the response to the secret-key queries, \mathcal{B} samples $\{z_{1,i}, z_{2,i} \xleftarrow{R} \mathbb{Z}_{q_1}^*\}_{i \in [1, n]}$, $\{y_{1,i}^j, y_{2,i}^j \xleftarrow{R} \mathbb{Z}_{q_1}^*\}_{i \in [1, n] \setminus \{m+1\}, j \in [1, k+1]}$, $Q_4 \xleftarrow{R} \mathbb{G}_2$, $R_5 \xleftarrow{R} \mathbb{G}_3$ and $f_1, f_2 \xleftarrow{R} \mathbb{Z}_{q_2}^*$. It also samples $\{y_{1,i}^{k+2}, y_{2,i}^{k+2} \xleftarrow{R} \mathbb{Z}_{q_1}^*\}_{i \in [1, m]}$. Observe that \mathcal{B} does not sample $y_{1,m+1}^j$ or $y_{2,m+1}^j$ for $j \in [1, k+2]$. This is because \mathcal{B} embeds the its input pair of $(k+1)$ -DLIN instances in its response by formally setting $y_{1,m+1}^j = a_j$ and $y_{2,m+1}^j = a'_j$ for each $j \in [1, k+2]$ as described next.
3. \mathcal{B} now sets the various components of the secret-key as follows:

$$\begin{aligned} d_0^* &= Q_4 \cdot R_5 / \left(\prod_{i=1}^n h_{1,i}^{z_{1,i}} \cdot h_{2,i}^{z_{2,i}} \right) \\ d_{1,i}^{*0} &= g_1^{z_{1,i}} \cdot g_2^{f_1 \cdot v_i} / \left(\prod_{j=1}^{k+1} \left(g_{1,j}^{x_{1,j}} \cdot g_{1,k+2}^{x_{1,k+2}} \right)^{y_{1,i}^j} \right) \quad \text{for } i \in [1, n] \setminus \{m+1\} \\ d_{2,i}^{*0} &= g_1^{z_{2,i}} \cdot g_2^{f_2 \cdot v_i} / \left(\prod_{j=1}^{k+1} \left(g_{2,j}^{x_{2,j}} \cdot g_{2,k+2}^{x_{2,k+2}} \right)^{y_{2,i}^j} \right) \quad \text{for } i \in [1, n] \setminus \{m+1\} \end{aligned}$$

It also sets the following additional components:

$$\begin{aligned} d_{1,i}^{*j} &= g_{1,j}^{y_{1,i}^j} \quad , \quad d_{2,i}^{*j} = g_{2,j}^{y_{2,i}^j} \quad \text{for } i \in [1, n] \setminus \{m+1\}, j \in [1, k+1] \\ d_{1,i}^{*k+2} &= g_{1,k+2}^{y_{1,i}^{k+2}} \quad , \quad d_{2,i}^{*k+2} = g_{2,k+2}^{y_{2,i}^{k+2}} \quad \text{for } i \in [1, m] \\ d_{1,i}^{*k+2} &= g_{1,k+2}^{\sum_{j=1}^{k+1} y_{1,i}^j} \quad , \quad d_{2,i}^{*k+2} = g_{2,k+2}^{\sum_{j=1}^{k+1} y_{2,i}^j} \quad \text{for } i \in [m+2, n] \end{aligned}$$

Observe that the first m components of the secret-key are crafted to be uniformly random, while the last $(n - m - 1)$ components are well-formed.

4. Finally, \mathcal{B} embeds its input pair of $(k+1)$ -DLIN instances in the $(m+1)^{\text{th}}$ component of the secret-key by setting:

$$\begin{aligned} d_{1,m+1}^{*0} &= g_1^{z_{1,m+1}} \cdot g_1^{f_1 \cdot v_{m+1}} / \left(\prod_{j=1}^{k+2} \left(g_{1,j}^{a_j} \right)^{x_{1,j}} \right) \\ d_{2,m+1}^{*0} &= g_1^{z_{2,m+1}} \cdot g_2^{f_2 \cdot v_{m+1}} / \left(\prod_{j=1}^{k+2} \left(g_{2,j}^{a'_j} \right)^{x_{2,j}} \right) \\ d_{1,m+1}^{*j} &= g_{1,j}^{a_j} \quad , \quad d_{2,m+1}^{*j} = g_{2,j}^{a'_j} \quad \text{for } j \in [1, k+2] \end{aligned}$$

\mathcal{B} responds to \mathcal{A} with the secret-key $\text{sk}_{\vec{v}^*}$ as:

$$\text{sk}_{\vec{v}^*} = \left(d_0^*, \{d_{1,i}^{*j}, d_{2,i}^{*j}\}_{i \in [1,n], j \in [0,k+2]} \right)$$

- **Function Privacy Encryption Oracle Query:** Let $\text{sk}_{\vec{v}^*} = \left(d_0^*, \{d_{1,i}^{*j}, d_{2,i}^{*j}\}_{i \in [1,n], j \in [0,k+2]} \right)$ be the response of the real-or-random function privacy oracle to the adversary \mathcal{A} . Suppose \mathcal{A} queries the function privacy encryption oracle with a message M . \mathcal{B} responds as follows: it samples $I^* = (I_1^*, \dots, I_n^*)$ such that $\langle \vec{v}^*, I^* \rangle = 0$. It then samples $r, \alpha, \beta \xleftarrow{R} \mathbb{Z}_N^*$ and $\{R_{6,i}^j, R_{7,i}^j \xleftarrow{R} \mathbb{G}_3\}_{i \in [1,n], j \in [0,k+2]}$ and sets $c_0^* = g_1^r$. It also sets:

$$\begin{aligned} c_{1,i}^{*0} &= (S_{1,i}^0)^r \cdot Q^{\alpha \cdot I_i^*} \cdot R_{6,i}^0, & c_{2,i}^{*0} &= (S_{2,i}^0)^r \cdot Q^{\beta \cdot I_i^*} \cdot R_{7,i}^0 \text{ for } i \in [1, n] \\ c_{1,i}^{*j} &= (S_{1,i}^j)^r \cdot R_{6,i}^j, & c_{2,i}^{*j} &= (S_{2,i}^j)^r \cdot R_{7,i}^j \text{ for } i \in [1, n], j \in [1, k+2] \end{aligned}$$

Finally, it sets:

$$c_3^* = M / \hat{e}(d_0^*, c_0^*) \cdot \left(\prod_{i=1}^n \prod_{j=0}^{k+2} \hat{e}(d_{1,i}^{*j}, c_{1,i}^{*j}) \cdot \hat{e}(d_{2,i}^{*j}, c_{2,i}^{*j}) \right)$$

and outputs the ciphertext $C^* = (c_0^*, \{c_{1,i}^{*j}, c_{2,i}^{*j}\}_{i \in [1,n], j \in [0,k+2]}, c_3^*)$. Clearly, for $m = 0$, the ciphertext C^* is well-formed, while for $m \in [1, n]$ independent of \vec{v}^* and is crafted by \mathcal{B} solely based on $\text{sk}_{\vec{v}^*}$ to produce the correct output M upon decryption.

- **Guess:** At the end of the game, \mathcal{A} outputs a bit b' . \mathcal{B} outputs the same bit b' .

It is easy to see that when $a_{k+2} = \sum_{j=1}^{k+1} a_j$ and $a'_{k+2} = \sum_{j=1}^{k+1} a'_j$, the secret-key $\text{sk}_{\vec{v}^*}$ is identically distributed to the response of the real-or-random oracle in the experiment $\text{Expt}_{\text{EFP}, \Pi_k^{\text{IPE}}, \mathcal{A}}^{\text{mode}, m}(\lambda)$. On the other hand, when either or both of a_{k+2} and a'_{k+2} are uniformly random in \mathbb{Z}_q^* , the secret-key $\text{sk}_{\vec{v}^*}$ is identically distributed to the response of the real-or-random oracle in the experiment $\text{Expt}_{\text{EFP}, \Pi_k^{\text{IPE}}, \mathcal{A}}^{\text{mode}, m+1}(\lambda)$. In either case, the ciphertext C^* is dependent solely on $\text{sk}_{\vec{v}^*}$ and the input message M , and is hence independent of the vector \vec{v}^* . In addition, it produces M upon decryption using $\text{sk}_{\vec{v}^*}$ irrespective of whether $a_{k+2} = \sum_{j=1}^{k+1} a_j$ and $a'_{k+2} = \sum_{j=1}^{k+1} a'_j$. Consequently, the distribution of C^* in the two experiments is indistinguishable from \mathcal{A} 's point of view. It follows readily that \mathcal{B} has the same advantage ϵ as \mathcal{A} in solving its input instance pair. This completes the proof of Claim 5.2.

We now make the following observation:

$$\begin{aligned} & \left| \Pr \left[\text{Expt}_{\text{EFP}, \Pi_k^{\text{IPE}}, \mathcal{A}}^{\text{mode}}(\lambda) = 1 \right] - \Pr \left[\text{Expt}_{\text{EFP}, \Pi_k^{\text{IPE}}, \mathcal{A}}^{\text{mode}, n}(\lambda) = 1 \right] \right| \\ & \leq \sum_{m=0}^{n-1} \left| \Pr \left[\text{Expt}_{\text{EFP}, \Pi_k^{\text{IPE}}, \mathcal{A}}^{\text{mode}, m}(\lambda) = 1 \right] - \Pr \left[\text{Expt}_{\text{EFP}, \Pi_k^{\text{IPE}}, \mathcal{A}}^{\text{mode}, m+1}(\lambda) = 1 \right] \right| \\ & \leq \text{negl}(\lambda) \text{ (from Claim 5.2) for } n = \text{poly}(\lambda) \end{aligned}$$

Consequently, for mode $\xleftarrow{R} \{\text{real}, \text{rand}\}$, we have:

$$\begin{aligned}
\text{Adv}_{\Pi_k^{\text{IPE}}, \mathcal{A}}^{\text{EFP}}(\lambda) &= \left| \Pr \left[\text{Expt}_{\text{EFP}, \Pi_k^{\text{IPE}}, \mathcal{A}}^{\text{real}}(\lambda) = 1 \right] - \Pr \left[\text{Expt}_{\text{EFP}, \Pi_k^{\text{IPE}}, \mathcal{A}}^{\text{rand}}(\lambda) = 1 \right] \right| \\
&\leq \left| \Pr \left[\text{Expt}_{\text{EFP}, \Pi_k^{\text{IPE}}, \mathcal{A}}^{\text{real}}(\lambda) = 1 \right] - \Pr \left[\text{Expt}_{\text{EFP}, \Pi_k^{\text{IPE}}, \mathcal{A}}^{\text{real}, n}(\lambda) = 1 \right] \right| \\
&\quad + \left| \Pr \left[\text{Expt}_{\text{EFP}, \Pi_k^{\text{IPE}}, \mathcal{A}}^{\text{rand}}(\lambda) = 1 \right] - \Pr \left[\text{Expt}_{\text{EFP}, \Pi_k^{\text{IPE}}, \mathcal{A}}^{\text{rand}, n}(\lambda) = 1 \right] \right| \\
&\quad + \left| \Pr \left[\text{Expt}_{\text{EFP}, \Pi_k^{\text{IPE}}, \mathcal{A}}^{\text{real}, n}(\lambda) = 1 \right] - \Pr \left[\text{Expt}_{\text{EFP}, \Pi_k^{\text{IPE}}, \mathcal{A}}^{\text{rand}, n}(\lambda) = 1 \right] \right| \\
&\leq 2 \left| \Pr \left[\text{Expt}_{\text{EFP}, \Pi_k^{\text{IPE}}, \mathcal{A}}^{\text{mode}}(\lambda) = 1 \right] - \Pr \left[\text{Expt}_{\text{EFP}, \Pi_k^{\text{IPE}}, \mathcal{A}}^{\text{mode}, n}(\lambda) = 1 \right] \right| \quad (\text{from Equation 2}) \\
&= 2\epsilon' \leq \text{negl}(\lambda)
\end{aligned}$$

This completes the proof of function privacy for Π_k^{IPE} . Note that the proof does not essentially rely on the presence of two parallel sub-systems in our scheme. In fact, eliminating one of the subsystems gives a simpler proof. However, the two sub-system version allows us to adopt the proof of data privacy in [3] for our scheme.

6 Extensions and Open Problems

Our work solves the open problem of constructing *computationally function private* public-key predicate encryption schemes, supporting a rich class of predicates without relying on strong assumptions such as indistinguishability obfuscation. Our approach, denoted as *encrypt-augment-recover*, offers a methodology for converting an existing public-key predicate encryption scheme into a computationally function private one, without compromising the data privacy guarantees of the original scheme. Our approach yields the first fully data private and computationally function private constructions for IBE and public-key IPE, and, in general, for searchable encryption schemes supporting a wide range of predicates. In this section, we discuss some interesting extensions and open problems that arise from our work.

Extension to Private-Key Predicate Encryption. Our *encrypt-augment-recover* methodology is equally applicable for achieving computationally function private predicate encryption schemes in the private-key setting, even when the underlying predicates are not necessarily sampled from distributions with at least super-logarithmic min-entropy. In particular, the core function privacy arguments for our constructions presented in this paper do not essentially rely on the unpredictability of the predicate distributions; this assumption is additionally made to rule out trivial attacks in the public-key setting. In the private-key setting, where the adversary does not have access to an encryption oracle, our approach anticipates an expansion to the existing body of work in designing function private predicate encryption schemes.

Extension to Multi-Input Predicate Encryption. In this work, we have focused on designing function private predicate encryption systems in the single-input setting. Multi-input predicate encryption (MIPE) introduced by Goldwasser et al. [32] is a generalization of functional encryption to the setting of multi-input predicates. An MIPE scheme has several encryption slots and each decryption key sk_f for a multi-input predicate f jointly decrypts the ciphertexts $\text{Enc}(I_1), \dots, \text{Enc}(I_n)$ for all slots to obtain $f(I_1, \dots, I_n)$ without revealing anything more about the encrypted attributes. In particular, this provides a framework to evaluate *bounded-norm* multi-input IPE: each predicate is specified by a collection of vectors $\vec{v}_1, \dots, \vec{v}_n$, and takes as input a collection of vectors $\vec{x}_1, \dots, \vec{x}_n$ to output $f_{\vec{v}_1, \dots, \vec{v}_n}(\vec{x}_1, \dots, \vec{x}_n) = \sum_{i=1}^n \{\vec{v}_i, \vec{x}_i\}$.

We point out that our technique can be easily generalized to obtain function private IPE schemes in the multi-input setting as follows: we first use our technique to obtain a function private IPE construction

in the single-input setting, and then run n independent copies of this construction. The i^{th} copy is used to encrypt \vec{x}_i in the i^{th} slot, while the new secret-key is the ensemble of the n secret-keys corresponding to $\vec{v}_1, \dots, \vec{v}_n$. The decryption algorithm computes each inner product individually, and returns their sum. Although this means that the adversary also learns each individual inner product, this is an inherent leakage in the public-key setting and does not weaken the security guarantees. The data privacy guarantees of the underlying scheme ensure no further leakage, while the function privacy guarantees of the underlying scheme continue to hold as long as each \vec{v}_i is sampled from block sources with sufficient min-entropy, and is independent of the other $n - 1$ vectors.

Generalization of Our Approach. In this work, we have demonstrated concrete constructions for function private IBE and IPE in the public-key setting based on existing constructions in the literature. An interesting open problem is to explore whether our approach can be generalized so as to be applicable to *any* public-key predicate encryption scheme, particularly those based on lattices [2], without compromising on their data privacy guarantees. It would also be interesting to identify the properties (if any) of existing predicate encryption schemes that make them amenable to modification using our approach.

Hidden Vector Encryption and Polynomial Evaluation. Boneh and Waters [1] proposed hidden vector encryption (HVE), a pre-cursor to IPE, that supports search using conjunctive, range and comparison-based query predicates. In HVE, attributes correspond to vectors over an alphabet Σ , while secret-keys correspond to predicate vectors over the augmented alphabet $\Sigma_\star = \Sigma \cup \{\star\}$ containing the wild card character \star . Decryption succeeds if the attribute matches the predicate vector in every coordinate that is not \star . We note that although IPE can be used to realize HVE [3], our computational function privacy definitions do not naturally extend to HVE. In particular, the presence of the wild card character \star in the predicate vectors of HVE trivially violates our min-entropy requirements, making it difficult to *hide* their presence in the secret-key.

A weaker notion of function privacy for HVE, that our framework can realistically incorporate, is to ask that the secret-key reveals nothing more about the the predicate vector than the locations of the wild card character \star . We believe that HVE schemes satisfying this weaker notion of computational function privacy can be achieved using our *encrypt-augment-recover* approach. It is still an open problem, however, to formalize a stronger function privacy definition for HVE, and to realize function private constructions satisfying this definition. This would also provide insight into the limits of function privacy for searchable encryption schemes supporting comparison and range queries. Finally, it is also open to formalize security definitions and realize constructions for function private encryption schemes that support arbitrary polynomial evaluation predicates [3].

References

1. Dan Boneh and Brent Waters. Conjunctive, Subset, and Range Queries on Encrypted Data. In *Theory of Cryptography, 4th Theory of Cryptography Conference, TCC 2007, Amsterdam, The Netherlands, February 21-24, 2007, Proceedings*, pages 535–554, 2007.
2. Shweta Agrawal, David Mandell Freeman, and Vinod Vaikuntanathan. Functional Encryption for Inner Product Predicates from Learning with Errors. In *Advances in Cryptology - ASIACRYPT 2011 - 17th International Conference on the Theory and Application of Cryptology and Information Security, Seoul, South Korea, December 4-8, 2011. Proceedings*, pages 21–40, 2011.
3. Jonathan Katz, Amit Sahai, and Brent Waters. Predicate Encryption Supporting Disjunctions, Polynomial Equations, and Inner Products. *J. Cryptology*, 26(2):191–224, 2013.
4. Dan Boneh and Matthew Franklin. Identity-based encryption from the weil pairing. *SIAM Journal on Computing*, 32(3):586–615, 2003.
5. Dan Boneh, Xavier Boyen, and Eu-Jin Goh. Hierarchical identity based encryption with constant size ciphertext. In *Advances in Cryptology-EUROCRYPT 2005*, pages 440–456. Springer, 2005.
6. Michel Abdalla, Mihir Bellare, Dario Catalano, Eike Kiltz, Tadayoshi Kohno, Tanja Lange, John Malone-Lee, Gregory Neven, Pascal Paillier, and Haixia Shi. Searchable Encryption Revisited: Consistency Properties, Relation to Anonymous IBE, and Extensions. *J. Cryptology*, 21(3):350–391, 2008.

7. Michel Abdalla, Florian Bourse, Angelo De Caro, and David Pointcheval. Simple functional encryption schemes for inner products. In *Public-Key Cryptography - PKC 2015 - 18th IACR International Conference on Practice and Theory in Public-Key Cryptography, Gaithersburg, MD, USA, March 30 - April 1, 2015, Proceedings*, pages 733–751, 2015.
8. Shweta Agrawal, Benoit Libert, and Damien Stehlé. Fully secure functional encryption for inner products, from standard assumptions. In *Advances in Cryptology - CRYPTO 2016 - 36th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 14-18, 2016, Proceedings, Part III*, pages 333–362, 2016.
9. Dan Boneh, Giovanni Di Crescenzo, Rafail Ostrovsky, and Giuseppe Persiano. Public Key Encryption with Keyword Search. In *Advances in Cryptology - EUROCRYPT 2004, International Conference on the Theory and Applications of Cryptographic Techniques, Interlaken, Switzerland, May 2-6, 2004, Proceedings*, pages 506–522, 2004.
10. Dan Boneh, Ananth Raghunathan, and Gil Segev. Function-Private Identity-Based Encryption: Hiding the Function in Functional Encryption. In *Advances in Cryptology - CRYPTO 2013 - 33rd Annual Cryptology Conference, Santa Barbara, CA, USA, August 18-22, 2013. Proceedings, Part II*, pages 461–478, 2013.
11. Dan Boneh, Ananth Raghunathan, and Gil Segev. Function-Private Subspace-Membership Encryption and Its Applications. In *Advances in Cryptology - ASIACRYPT 2013 - 19th International Conference on the Theory and Application of Cryptology and Information Security, Bengaluru, India, December 1-5, 2013, Proceedings, Part I*, pages 255–275, 2013.
12. Vincenzo Iovino, Qiang Tang, and Karol Zebrowski. On the power of public-key function-private functional encryption. In *Cryptology and Network Security - 15th International Conference, CANS 2016, Milan, Italy, November 14-16, 2016, Proceedings*, pages 585–593, 2016.
13. Shashank Agrawal, Shweta Agrawal, Saikrishna Badrinarayanan, Abishek Kumarasubramanian, Manoj Prabhakaran, and Amit Sahai. Function private functional encryption and property preserving encryption : New definitions and positive results. *IACR Cryptology ePrint Archive*, 2013:744, 2013.
14. Reza Curtmola, Juan A. Garay, Seny Kamara, and Rafail Ostrovsky. Searchable symmetric encryption: improved definitions and efficient constructions. In *Proceedings of the 13th ACM Conference on Computer and Communications Security, CCS 2006, Alexandria, VA, USA, Ioctober 30 - November 3, 2006*, pages 79–88, 2006.
15. Emily Shen, Elaine Shi, and Brent Waters. Predicate Privacy in Encryption Systems. In *Theory of Cryptography, 6th Theory of Cryptography Conference, TCC 2009, San Francisco, CA, USA, March 15-17, 2009. Proceedings*, pages 457–473, 2009.
16. Allison Bishop, Abhishek Jain, and Lucas Kowalczyk. Function-hiding inner product encryption. In *Advances in Cryptology - ASIACRYPT 2015 - 21st International Conference on the Theory and Application of Cryptology and Information Security, Auckland, New Zealand, November 29 - December 3, 2015, Proceedings, Part I*, pages 470–491, 2015.
17. Pratish Datta, Ratna Dutta, and Sourav Mukhopadhyay. Functional encryption for inner product with full function privacy. In *Public-Key Cryptography - PKC 2016 - 19th IACR International Conference on Practice and Theory in Public-Key Cryptography, Taipei, Taiwan, March 6-9, 2016, Proceedings, Part I*, pages 164–195, 2016.
18. Ronald Cramer and Victor Shoup. A practical public key cryptosystem provably secure against adaptive chosen ciphertext attack. In *Advances in Cryptology - CRYPTO '98, 18th Annual International Cryptology Conference, Santa Barbara, California, USA, August 23-27, 1998, Proceedings*, pages 13–25, 1998.
19. Ananth Raghunathan, Gil Segev, and Salil P. Vadhan. Deterministic public-key encryption for adaptively chosen plaintext distributions. In *Advances in Cryptology - EUROCRYPT 2013, 32nd Annual International Conference on the Theory and Applications of Cryptographic Techniques, Athens, Greece, May 26-30, 2013. Proceedings*, pages 93–110, 2013.
20. Zvika Brakerski and Gil Segev. Function-Private Functional Encryption in the Private-Key Setting. In *Theory of Cryptography - 12th Theory of Cryptography Conference, TCC 2015, Warsaw, Poland, March 23-25, 2015, Proceedings, Part II*, pages 306–324, 2015.
21. Oded Goldreich and Rafail Ostrovsky. Software protection and simulation on oblivious RAMs. *J. ACM*, 43(3):431–473, 1996.
22. Dawn Xiaodong Song, David Wagner, and Adrian Perrig. Practical techniques for searches on encrypted data. In *2000 IEEE Symposium on Security and Privacy, Berkeley, California, USA, May 14-17, 2000*, pages 44–55, 2000.
23. Yan-Cheng Chang and Michael Mitzenmacher. Privacy preserving keyword searches on remote encrypted data. In *Applied Cryptography and Network Security, Third International Conference, ACNS 2005, New York, NY, USA, June 7-10, 2005, Proceedings*, pages 442–455, 2005.

24. Melissa Chase and Seny Kamara. Structured encryption and controlled disclosure. In *Advances in Cryptology - ASIACRYPT 2010 - 16th International Conference on the Theory and Application of Cryptology and Information Security, Singapore, December 5-9, 2010. Proceedings*, pages 577–594, 2010.
25. Seny Kamara, Charalampos Papamanthou, and Tom Roeder. Dynamic searchable symmetric encryption. In *the ACM Conference on Computer and Communications Security, CCS'12, Raleigh, NC, USA, October 16-18, 2012*, pages 965–976, 2012.
26. Dan Boneh, Xavier Boyen, and Hovav Shacham. Short group signatures. In *Advances in Cryptology - CRYPTO 2004, 24th Annual International Cryptology Conference, Santa Barbara, California, USA, August 15-19, 2004, Proceedings*, pages 41–55, 2004.
27. Hovav Shacham. A cramer-shoup encryption scheme from the linear assumption and from progressively weaker linear variants. *IACR Cryptology ePrint Archive*, 2007:74, 2007.
28. Craig Gentry, Chris Peikert, and Vinod Vaikuntanathan. Trapdoors for hard lattices and new cryptographic constructions. In *Proceedings of the 40th Annual ACM Symposium on Theory of Computing, Victoria, British Columbia, Canada, May 17-20, 2008*, pages 197–206, 2008.
29. Xavier Boyen and Brent Waters. Anonymous hierarchical identity-based encryption (without random oracles). In *Advances in Cryptology-CRYPTO 2006*, pages 290–307. Springer, 2006.
30. Shweta Agrawal, Dan Boneh, and Xavier Boyen. Lattice basis delegation in fixed dimension and shorter-ciphertext hierarchical IBE. In *Advances in Cryptology - CRYPTO 2010, 30th Annual Cryptology Conference, Santa Barbara, CA, USA, August 15-19, 2010. Proceedings*, pages 98–115, 2010.
31. Kaoru Kurosawa and Le Trieu Phong. Leakage resilient IBE and IPE under the DLIN assumption. In *Applied Cryptography and Network Security - 11th International Conference, ACNS 2013, Banff, AB, Canada, June 25-28, 2013. Proceedings*, pages 487–501, 2013.
32. Vipul Goyal, Aayush Jain, and Adam O’Neill. Multi-input functional encryption with unbounded-message security. In *Advances in Cryptology - ASIACRYPT 2016 - 22nd International Conference on the Theory and Application of Cryptology and Information Security, Hanoi, Vietnam, December 4-8, 2016, Proceedings, Part II*, pages 531–556, 2016.