



HAL
open science

Dynamic Decentralized Functional Encryption

Jérémy Chotard, Edouard Dufour-Sans, Romain Gay, Duong Hieu Phan,
David Pointcheval

► **To cite this version:**

Jérémy Chotard, Edouard Dufour-Sans, Romain Gay, Duong Hieu Phan, David Pointcheval. Dynamic Decentralized Functional Encryption. CRYPTO 2020 - 40th Annual International Cryptology Conference, Aug 2020, Santa Barbara / Virtual, United States. pp.747-775, 10.1007/978-3-030-56784-2_25 . hal-02947359

HAL Id: hal-02947359

<https://inria.hal.science/hal-02947359>

Submitted on 5 Nov 2020

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Dynamic Decentralized Functional Encryption

Jérémy Chotard^{1,2,3}, Edouard Dufour-Sans^{2,3,4}, Romain Gay⁵, Duong Hieu Phan¹, and David Pointcheval^{2,3}

¹ XLIM, University of Limoges, CNRS, Limoges, France

² DIENS, École normale supérieure, CNRS, PSL University, Paris, France

³ INRIA, Paris, France

⁴ Carnegie Mellon University, Pittsburgh, USA

⁵ Cornell Tech, New York, USA

Abstract. We introduce Dynamic Decentralized Functional Encryption (DDFE), a generalization of Functional Encryption which allows multiple users to join the system dynamically, without relying on a trusted third party or on expensive and interactive Multi-Party Computation protocols.

This notion subsumes existing multi-user extensions of Functional Encryption, such as Multi-Input, Multi-Client, and Ad Hoc Multi-Input Functional Encryption.

We define and construct schemes for various functionalities which serve as building-blocks for latter primitives and may be useful in their own right, such as a scheme for dynamically computing sums in any Abelian group. These constructions build upon simple primitives in a modular way, and have instantiations from well-studied assumptions, such as DDH or LWE.

Our constructions culminate in an Inner-Product scheme for computing weighted sums on aggregated encrypted data, from standard assumptions in prime-order groups in the Random Oracle Model.

Keywords. Dynamic, Decentralized, Functional Encryption, Inner Product.

1 Introduction

At TCC'11, Boneh, Sahai, and Waters [BSW11] formalized Functional Encryption (FE), a new paradigm of Public-Key Encryption that allows the owner of the secret key to generate restricted keys, enabling third parties to recover *function evaluations* of the plaintext from a ciphertext. The formalization of FE gave many researchers a common framework in which to consider their schemes: the nuances between Identity-Based Encryption (IBE), Hierarchical IBE, Fuzzy IBE, and different forms of Attribute-Based Encryption (ABE) [BF01, BBG05, SW05, GPSW06] could now be captured simply by specifying which functionality the scheme aims to implement. The set of algorithms to be implemented and the indistinguishability game in which to prove security were now standard.

But for all its successes, Functional Encryption has two, somewhat related, important limitations: (1) In many contexts, FE encourages centralization. In his 2015 position paper *The Moral Character of Cryptographic Work*, Rogaway pointed out that a switch from Public-Key Encryption to Identity-Based Encryption would represent "a radical change in the trust model", as the authority with knowledge of the master secret key would have the ability to fully recover every message encrypted under its public key, even though those messages would be intended for a variety of parties. This criticism can be extended to many other functionalities of Functional Encryption. (2) The kind of controlled computation enabled by Functional Encryption does not extend to computations involving data from multiple parties. This is limiting because a significant component of the public's privacy concerns today is related to data being made available to a third-party for the advertised purpose of retrieving some form of intelligence of the public's needs, from the computation of simple statistics to the training of advanced machine learning models. This means FE is not an appropriate framework for addressing this pressing issue.

1.1 Our Contributions

1. First, we fill the gap left by the definition of FE by introducing a new primitive we term Dynamic Decentralized Functional Encryption (DDFE). DDFE allows aggregating data coming from different parties, does not require a trusted party with a master secret key, and accounts for participants

wanting to join at various stages during the lifetime of a system. Previous extensions of FE, which we review in more detail in Section 1.2, either failed to address the concerns we raised above, or partially forwent the generality that made the success of FE as a framework for describing cryptographic schemes. We give a formal definition of DDFE as well as a security definition.

2. We define All-or-Nothing Encapsulation (AoNE), a functionality of DDFE which we found to be a critical building-block when constructing useful DDFE schemes later in this work. AoNE allows a participant to send its data to be aggregated with other data coming from a group of participants agreeing on a label ℓ . Only if all those participants choose to send data for aggregation with the same group under the same label will the data of all participants be revealed, otherwise, nothing is revealed. We provide two constructions of AoNE. The first one is generic from any IBE, but has individual ciphertexts that grow linearly in the number of participants in an aggregation, which is not ideal. The second construction is specific and achieves constant size ciphertexts. It relies on bilinear maps, and we prove its security under the DBDH assumption in the Random Oracle Model (ROM).
3. We define and provide a construction of DSum, a functionality of DDFE which is both interesting in its own right and a useful building-block for other constructions. DSum operates over any Abelian group and allows multiple parties to send an element from that group for aggregation with a set of participants agreeing on a label ℓ . Once every participant has sent data for aggregation with that set and that label, the sum (or rather the repeated group operation) of the data is revealed. We provide a generic construction of DSum from Non-Interactive Key Exchange (NIKE), AoNE DDFE and Pseudo-Random Functions (PRF).
4. We define and provide a construction of Inner-Product DDFE (IP-DDFE), which allows for more complex patterns of aggregation than DSum. In IP-DDFE, participants can contribute to the generation of functional decryption keys that enable individuals to compute weighted sums of plaintext data, with the weights being encoded in the key. Our construction relies on AoNE, DSum, Single-Input Inner-Product Functional Encryption, and PRFs, and we prove that it is selectively secure under the DDH assumption in the ROM.

1.2 Related Work

Fully Homomorphic Encryption (FHE) [Gen09] is commonly cited as a cryptographic solution to issues involving computations on encrypted data at large. We stress here that FHE shines when computation delegation is intended. That is, it is useful when a client, owning some data it wishes to protect the confidentiality of, wants a server to perform computations on their data without seeing the data. This scenario arises when the computation depends on parameters known only to the server (as in the case of Information Retrieval), or when the client wants to leverage the computational power of the server.

In the scenario we are concerned with, however, the server directly learns something about the aggregated data, without interacting with them. This stands in contrast with FHE, where the parties need to engage in extra rounds of interaction to perform a joint decryption of the encrypted data.

FE enables the server to recover information as controlled by the client through key delegation, while FHE does not limit the types of computations the server can perform, but prevents the server from accessing any data. Given these advantages, we naturally focus on extending the line of works involving FE.

Note that FHE was also initially defined for a single data owner, and was later extended to multiple users under the name Multi-Key FHE [LTV12].

Private Stream Aggregation (PSA). This notion, initially termed Privacy-Preserving Aggregation of Time-Series Data, is an early primitive for non-interactive aggregation of multi-party data introduced by Shi *et al.* [SCR⁺11]. Unlike our DDFE schemes, PSA, under its standard definitions, relies on a

trusted third-party distributing the participant’s secret keys, cannot accommodate new participants, and does not allow the participants to choose which functions can be computed by whom via functional decryption key derivation. Most PSA schemes in the literature focus on computing (non-weighted) sums of the participants’ data [CSS12, JL13, BJJ16]. Note that Private Stream Aggregation usually relies on a Differential Privacy component as an added privacy protection, while we leave the addition of a Differential Privacy layer in DDFE for future work.

Multi-Authority Functional Encryption (MAFE) was introduced by Chandran *et al.* [CGJS15]. Like DDFE, it is a strongly decentralized variant of Functional Encryption. It allows for encrypting messages for sets of authorities along with an access policy. These authorities can then generate keys for individual identities. Armed with a single ciphertext and a set of functional decryption keys from the appropriate authorities, the decrypter can recover a function of the plaintext that is specified by the access policy on the identities for which the functional keys were computed. Unlike DDFE, MAFE does not account for the possibility of multiple ciphertexts being decrypted together, and having their data interact with one another.

Multi-Client Functional Encryption (MCFE) was defined in [GGG⁺14, GKL⁺13] along with Multi-Input Functional Encryption (MIFE), and also enables computing functions of multiple parties’ data in the presence of a trusted third-party distributing both the parties’ secret keys and functional decryption keys. That is, both MIFE and MCFE extend Functional Encryption to a setting where the input is spread across different sources. Each source can encrypt its data independently, and the ciphertexts can then be aggregated and decrypted with functional decryption keys. Generation of the latter still requires a trusted authority, which owns a so-called master secret key: a single point of failure for the cryptosystem.

As opposed to MIFE, the encryption algorithm of an MCFE takes an additional input, referred to as a label, which enforces a finer-grained control on access to the encrypted data. Unlike in MIFE, where individual ciphertexts can be arbitrarily combined, in MCFE, only ciphertexts generated for the same label can be used together to decrypt. This limits how much information is revealed by each functional decryption key, thereby strengthening security. Typically, labels are used as timestamps. In this context, a functional decryption key can only compute, say, statistics on aggregated data *for the same time frame*.

Any MCFE for a given functionality directly implies an MIFE for the same functionality, by simply using a fixed label for all encryptions¹. Reciprocally, an MIFE for general functions would directly imply an MCFE for general functions, since the label can be part of the plaintext, and the function can check that every slot used the same label. However, this is not true for the case of smaller classes of functions for which there are practical schemes, such as Inner-Products.

The first construction of MIFE for inner products was given in [AGRW17], from standard assumptions in pairing groups. This was later improved by [ACF⁺18], which gave a generic construction from any single-input FE for inner products. The first construction of MCFE from standard assumptions was given by Chotard *et al.* [CDG⁺18] for computing inner products, although the security they achieved admits several limitations compared to the standard MCFE security definition.

Decentralized Multi-Client Functional Encryption (DMCFE). Chotard *et al.* [CDG⁺18] also defined a new variant of MCFE, called Decentralized MCFE (DMCFE), for which they gave Inner-Product instantiations from pairings. The DMCFE variant did away with the trusted third-party, as it enabled participants to choose their own secret keys and generate functional decryption keys non-interactively.

¹ Note that this was not true for MCFE as originally defined in [GGG⁺14], as that definition had strictly increasing timestamps for labels. But followup works on MCFE have usually allowed any bitstring to be used as a label, opening the primitive to the possibility of repetitions.

However, it still had an interactive setup, with no easy way of adding new participants, and it suffered from the same security caveats as the MCFE it was a variant of.

In a follow-up work, [LT19] provided a construction in the standard model from the LWE assumption, which still suffers from the same security restrictions as [CDG⁺18]. The works [ABKW19, ABG19] improved the security guarantees obtained, the former using the DDH assumption in the ROM, the latter using a generic construction from any single-input FE for inner products. Both schemes however have individual ciphertexts of size proportional to the total number of users. Thus, we use different techniques to obtain the desirable security notion without having asymptotically large ciphertexts.

Ad Hoc Multi-Input Functional Encryption. In [ACF⁺20], the authors define the notion of Ad Hoc Multi-Input Functional Encryption, where users can join the system on-the-fly, and functional decryption keys can be generated in a decentralized way, by each client, without interaction. They give a feasibility result for all functions, and a practical construction for inner products.

The definition of DDFE we put forth is more general than [ACF⁺20]. For instance, in our definition, the algorithm that generates functional decryption key does not necessarily require a specified group of users: schemes with potentially more flexibility than Ad Hoc MIFE can be captured by our definition.

Moreover, their scheme for inner product cannot handle labels, which implies that ciphertexts computed by each client individually can be mix and matched arbitrarily. As explained above, this implies that each functional decryption key reveals large amounts of information on the encrypted values, and renders the security vacuous whenever sufficiently many keys are issued. Labels help mitigate this leakage by enforcing a better granularity on the way the encrypted data is accessed.

Besides, the security model of [ACF⁺20] does not explicitly address the information that can be leaked when decrypting partial ciphertexts, that is, ciphertexts coming from an incomplete group of users. Preventing the adversary from recovering information on partial ciphertexts is made more challenging in our construction, which handles labels.

1.3 Outline

We first provide a definition of DDFE in Section 2, along with a security definition and functionalities of interest. In Section 3, we recall some useful preliminaries and definitions. We then showcase our constructions: a generic construction of AoNE is presented in Section 4, while a specific instantiation is studied in Section 5. We use it modularly in Section 6 to construct a DSum scheme. In Section 7, we capitalize on both those primitives to construct a DDFE scheme for the Inner-Product functionality.

2 Definitions and Security Models

In this section, we provide the formal definition of our new primitive of *Dynamic Decentralized Functional Encryption* (DDFE), together with several security models. Then, we list a few instantiations with some concrete functionalities.

2.1 Notations

In the following, $[n]$ will denote the set of integers $\{1, \dots, n\}$. For any set \mathcal{A} , $\mathcal{L}(\mathcal{A})$ will denote the set of finite lists of elements of \mathcal{A} , while $\mathcal{S}(\mathcal{A})$ will denote the set of finite subsets of \mathcal{A} . Unlike sets, lists are ordered and may contain repeated elements.

2.2 Dynamic Decentralized Functional Encryption

In defining DDFE, one of our key concerns is generality: we want to achieve for multi-user primitives what Functional Encryption did for single-user primitives. We resist as much as possible the temptation to let the idiosyncrasies of the functionalities we present and implement in this work leak into the

definition of DDFE itself. Perhaps the best example of this is in the role of the label. We believe labels, as used in MCFE, are useful for practical use, because in limiting what can be decrypted, they limit data leakage and make it possible to consider using the same primitive over a long time. However, we recognize that some primitives which are of practical use without labels may arise, that some schemes using labels may want to have them interact in more complex ways than perfect matching, and that there is value in our definitions being able to capture existing work. In Section 2.3, we give more details on how our umbrella notion captures a large set of existing primitives, ranging from Public-Key Encryption to Ad Hoc Multi-Input Function Encryption as introduced in Agrawal *et al.* [ACF⁺20].

Definition 1 (Dynamic Decentralized Functional Encryption). *A dynamic decentralized functional encryption scheme over a set of public keys \mathcal{PK} for functionality $\mathcal{F} : \mathcal{L}(\mathcal{PK} \times \mathcal{K}) \times \mathcal{L}(\mathcal{PK} \times \mathcal{M}) \rightarrow \{0, 1\}^*$ consists of five algorithms:*

- $\text{Setup}(\lambda)$: Generates and outputs public parameters pp . Those parameters are implicit arguments to all the other algorithms;
- $\text{KeyGen}()$: Generates and outputs a party's public key $\text{pk} \in \mathcal{PK}$ and the corresponding secret key sk_{pk} ;
- $\text{Encrypt}(\text{sk}_{\text{pk}}, m)$: Takes as input a party's secret key sk_{pk} , a value $m \in \mathcal{M}$ to encrypt, and outputs a ciphertext ct_{pk} ;
- $\text{DKeyGen}(\text{sk}_{\text{pk}}, k)$: Takes as input a party's secret key sk , a key space object k , and outputs a functional decryption key $\text{dk}_{\text{pk}, k}$;
- $\text{Decrypt}((\text{dk}_{\text{pk}, k_{\text{pk}}})_{\text{pk} \in \mathcal{U}_K}, (\text{ct}_{\text{pk}})_{\text{pk} \in \mathcal{U}_M})$: Takes as input a finite list of functional decryption keys $(\text{dk}_{\text{pk}, k_{\text{pk}}})_{\text{pk} \in \mathcal{U}_K}$, a finite list of ciphertexts $(\text{ct}_{\text{pk}})_{\text{pk} \in \mathcal{U}_M}$, where $\mathcal{U}_M, \mathcal{U}_K \in \mathcal{L}(\mathcal{PK})$ are the lists of senders and receivers, respectively. It outputs a value $y \in \{0, 1\}^*$.

We call a DDFE scheme *Public-Key* if its encryption algorithm does not make use of the secret key sk_{pk} .

Correctness: We require that, for all security parameters $\lambda \in \mathbb{N}$, for all polynomial size lists $\mathcal{U}_M, \mathcal{U}_K \in \mathcal{L}(\mathcal{PK})$ of public keys issued by $\text{KeyGen}()$, $(\text{pk}, k_{\text{pk}})_{\text{pk} \in \mathcal{U}_K} \in \mathcal{L}(\mathcal{PK} \times \mathcal{K})$ and $(\text{pk}, m_{\text{pk}})_{\text{pk} \in \mathcal{U}_M} \in \mathcal{L}(\mathcal{PK} \times \mathcal{M})$, it holds that the probability for

$$\text{Decrypt}((\text{dk}_{\text{pk}, k_{\text{pk}}})_{\text{pk} \in \mathcal{U}_K}, (\text{ct}_{\text{pk}})_{\text{pk} \in \mathcal{U}_M}) = F((\text{pk}, k_{\text{pk}})_{\text{pk} \in \mathcal{U}_K}, (\text{pk}, m_{\text{pk}})_{\text{pk} \in \mathcal{U}_M})$$

is 1, taken over $\text{pp} \leftarrow \text{Setup}(\lambda)$, $\text{dk}_{\text{pk}, k_{\text{pk}}} \leftarrow \text{DKeyGen}(\text{sk}_{\text{pk}}, k_{\text{pk}})$, for all $\text{pk} \in \mathcal{U}_K$, and $\text{ct}_{\text{pk}} \leftarrow \text{Encrypt}(\text{sk}_{\text{pk}}, m_{\text{pk}})$ for all $\text{pk} \in \mathcal{U}_M$.

We stress that each user is identified by a public key pk , which it can generate on its own with the associated secret key, using KeyGen . Anyone can thus dynamically join the system, by publishing its public key.

Remark 2 (Empty keys). Note that, unlike with standard, Single-Input FE, we do not require the empty key ϵ to be in \mathcal{K} , because we operate over lists of elements of $\mathcal{PK} \times \mathcal{K}$, so we simply define ϵ as the empty list.

In both Single-Input Functional Encryption and DDFE, the empty key serves to capture all the information about the plaintext that intentionally leaks from every ciphertext (see [BSW11, Section 2]). In Single-Input FE, this is typically only used to highlight the fact that encryption leaks the length of the message.

It is crucial to the security of any Functional Encryption scheme which accepts messages of variable lengths and leaks the length of the message, for otherwise it would be easy to win the IND security game by querying QLeftRight for two messages of different lengths (see Definition 17). With the leakage clearly stated in the functionality of the scheme, such a query would trigger the condition in the game's Finalize , and it would cause the adversary's guess to be discarded.

But in the case of DDFE, more information is usually publicly associated with a ciphertext that simply

its length. For instance, the set of users the data should be aggregated with, or the aggregation label, are typically public. Besides, it happens that the leakage of a set of ciphertexts is more than the cumulative leakage of the individual ciphertexts. Our AoNE and DSum schemes have this property, and it is expressed by their functionality outputting the relevant information when evaluated on the empty key with a (possibly non-singleton) list of ciphertexts.

Definition 3 (IND-Security Game for DDFE). *Let us consider a DDFE scheme. No adversary \mathcal{A} should be able to win the following security game against a challenger \mathcal{C} , with unlimited and adaptive access to the oracles QNewHonest, QEncrypt, QLeftRight, QDKeyGen, and QCorrupt described below:*

- *Initialize: the challenger \mathcal{C} runs the setup algorithm $\text{pp} \leftarrow \text{Setup}(\lambda)$ and chooses a random bit $b \xleftarrow{\$} \{0, 1\}$. It provides pp to the adversary \mathcal{A} ;*
- *Participant creation queries QNewHonest: the challenger \mathcal{C} runs the key generation algorithm $(\text{pk}, \text{sk}_{\text{pk}}) \leftarrow \text{KeyGen}()$ to simulate a new participant, stores the association $(\text{pk}, \text{sk}_{\text{pk}})$ and returns pk to the adversary;*
- *Encryption queries QEncrypt(pk, m): Recovers the secret key sk associated to pk and outputs the ciphertext $\text{ct} \leftarrow \text{Encrypt}(\text{sk}, m)$. If pk is not associated with any secret key, nothing is returned;*
- *Challenge queries QLeftRight(pk, m^0, m^1): runs and forwards the output of QEncrypt(pk, m^b). Wlog. we assume $m^0 \neq m^1$.*
- *Functional decryption key queries QDKeyGen(pk, k): Recovers the secret key sk associated to pk and outputs the functional decryption key $\text{dk}_k \leftarrow \text{DKeyGen}(\text{sk}, k)$. If pk is not associated with any secret key, nothing is returned;*
- *Corruption queries QCorrupt(pk): Recovers the secret key sk associated to pk and outputs it. If pk is not associated with any secret key, nothing is returned;*
- *Finalize: \mathcal{A} provides its guess b' on the bit b , and this procedure outputs the result β of the security game, according to the analysis given below.*

The output β of the game depends on some conditions, where \mathcal{HS} is the set of honest participants at the end of the game (the set of public keys generated via QNewHonest-queries and not corrupted via QCorrupt). Finalize outputs the bit $\beta = (b' = b)$, unless the following condition (*) is satisfied, in which case Finalize outputs a random bit β .

The condition (*) is true if there exist two lists of public keys $\mathcal{U}_M, \mathcal{U}_K \in \mathcal{L}(\mathcal{PK})$, two lists of messages $(\mathbf{m}^0 = (\text{pk}, m_{\text{pk}}^0)_{\text{pk} \in \mathcal{U}_M}, \mathbf{m}^1 = (\text{pk}, m_{\text{pk}}^1)_{\text{pk} \in \mathcal{U}_M})$, and a list of keys $\mathbf{k} = (\text{pk}, k_{\text{pk}})_{\text{pk} \in \mathcal{U}_K}$, such that $F(\mathbf{k}, \mathbf{m}^0) \neq F(\mathbf{k}, \mathbf{m}^1)$, with

- $m_{\text{pk}}^0 = m_{\text{pk}}^1$, for all $\text{pk} \in \mathcal{U}_M$ such that $\text{pk} \notin \mathcal{HS}$;
- QLeftRight($\text{pk}, m_{\text{pk}}^0, m_{\text{pk}}^1$) or QEncrypt(pk, m_{pk})-queries have been asked for all $\text{pk} \in \mathcal{U}_M \cap \mathcal{HS}$;
- QDKeyGen(pk, k_{pk})-queries have been asked for all $\text{pk} \in \mathcal{U}_K \cap \mathcal{HS}$.

We say DDFE is IND-secure if for any adversary \mathcal{A} ,

$$\text{Adv}_{\text{DDFE}}^{\text{IND}}(\mathcal{A}) = |2 \times \Pr[\beta = 1] - 1|$$

is negligible.

Intuitively, condition (*) means that the adversary can trivially recover b and win the game, which is thus not a real attack, hence a meaningless output with a random bit. Otherwise, $\beta = 0$ is a wrong guess and $\beta = 1$ is a correct guess during a meaningful attack. As usual, we are interested in adversaries with non-negligible advantage. Note however that the condition of trivial win cannot, in general, be checked in polynomial time. This is because there are exponentially many choices that can be made for the various lists, including the participant public keys and the values of the messages. Even if we impose strict requirements on the functionality, such as the presence of a label and a set of participants, it might not be possible to guarantee that the condition can be checked in polynomial time without a direct analysis of the structure of the functionality. There may exist functionalities for which such a

check is a computationally hard problem. The issue of how to efficiently check for violations is thus left to the cryptosystem designers and provers. In the following, we will consider functionalities for which this condition can be efficiently checked.

Now we present several weaker variants of the above security notion.

Definition 4 (sym-IND-Security Game for DDFE). We define a symmetric-key variant of the above security game in which the Finalize procedure outputs 0 if the adversary makes a query of the form (pk, m_0, m_1) to QLeftRight and queries the same pk to QCorrupt. This means that the secret key sk_{pk} not only allows users to encrypt on behalf of party pk , but also empowers them to decrypt the ciphertext generated by party pk . Thus, the challenge messages m_0 and m_1 have to be the same to avoid the adversary trivially recovering the random bit β . That is, the oracle QEncrypt should be used instead of QLeftRight.

Definition 5 (sel-IND-Security Game for DDFE). We define a selective variant of the above security game in which the adversary is forced to send all its queries to QNewHonest, upon which it receives the corresponding public keys. Then it sends all its queries to the oracles QEncrypt, QLeftRight, QDKeyGen and QCorrupt in one shot, and receives all of the outputs at once.

Note that our security notions is strong, in the sense that it allows the adversary to generate malicious public keys on its own. The challenger does not know the corresponding secret keys (which may not exist) for such public keys. More precisely, we allow dishonest key registrations, as originally introduced in [CKS08] in the context of NIKE.

2.3 Versatility of the Notion of DDFE

The notion of DDFE captures many existing primitives. We go over some such primitives and provide details here.

We first show that the notion of public-key encryption is captured by DDFE. That is, we can cast the former as a DDFE for a specific functionality that we present here. Apart from being a warm-up before delving into more advanced primitives, this shows that DDFE is not fundamentally restricted to secret-key primitives.

Public-Key Encryption. Here, the message space $\mathcal{M} = \{0, 1\}^* \times \mathcal{PK}$ comprises pairs of plaintext and public keys. The key space is restricted to the identity function over the plaintexts: $\{f_{id}\}$. The functionality takes as input the list of pairs (pk, m_{pk}) from all senders $pk \in \mathcal{U}_M$. In our case, the list \mathcal{U}_M will contain only one user pk_1 who wishes to send the plaintext $pt \in \{0, 1\}^*$ to user pk_2 . This information is contained in the message $m_{pk_1} = (pt, pk_2)$.

The functionality also takes the list of pairs (pk, k_{pk}) from all receivers $pk \in \mathcal{U}_K$. In our case, the list \mathcal{U}_K only contains the recipient pk'_2 . The associated key space object is the identity function f_{id} , which is the only function available here.

The functionality outputs the plaintext if the intended recipient is the actual recipient. That is $F((pk'_2, f_{id}), (pk_1, (pt, pk_2))) = pt$ if $pk_2 = pk'_2$, \perp otherwise. On any input that does not have that format (for instance on lists \mathcal{U}_M and \mathcal{U}_K of more than one element), the functionality will also output \perp .

The above example can be generalized straightforwardly to capture single-input Functional Encryption [BSW11], by considering a larger key space $\{f\}$ that is not only restricted to the identity function.

Decentralized Attribute-Based Encryption. The notion of DDFE can also capture existing decentralized primitives, such as the notion of decentralized Attribute-Based Encryption introduced in [LW11], as shown below. It also captures the more general Multi-Authority Functional Encryption [CGJS15].

Here, the message space $\mathcal{M} = \{0, 1\}^* \times \mathcal{P} \times \mathcal{L}(\mathcal{PK})$ comprises tuples, each of which contains a plaintext, a predicate, and a list of public keys. The key space $\mathcal{K} = \mathcal{A} \times \mathcal{ID}$ comprises pairs of an attribute and an identifier.

The functionality takes as input the list of pairs $(\mathbf{pk}, m_{\mathbf{pk}})$ from all senders $\mathbf{pk} \in \mathcal{U}_M$. In our case, the list \mathcal{U}_M will contain only one user \mathbf{pk} who wishes to send the plaintext $\mathbf{pt} \in \{0, 1\}^*$ to any user with proper credentials, that is, whose attributes satisfy an access policy expressed by a predicate $P \in \mathcal{P}$. This predicate takes as inputs attributes that are handled by different authorities, listed in \mathcal{U} . All of this information is contained in the message $m_{\mathbf{pk}} = (\mathbf{pt}, P, \mathcal{U})$.

The functionality also takes the list of pairs $(\mathbf{pk}, k_{\mathbf{pk}})$ from all receivers $\mathbf{pk} \in \mathcal{U}_K$. In our case, the list \mathcal{U}_K contains the authorities involved. For each authority, the associated key space object is an attribute, and a global identifier.

The functionality is defined as $F((\mathbf{pk}_i, (\mathbf{att}_i, \mathbf{GID}_i))_{\mathbf{pk}_i \in \mathcal{U}_K}, (\mathbf{pk}, (\mathbf{pt}, P, \mathcal{U}))) = \mathbf{pt}$ if $\mathcal{U} = \mathcal{U}_K$, all the identifiers \mathbf{GID}_i are the same, and the predicate P on the attributes \mathbf{att}_i evaluates to true. If these conditions are not met, or if the input does not have the right syntax (e.g. the list \mathcal{U}_M has more than one element), the functionality outputs \perp .

Ad Hoc Multi-Input FE. We now show that DDFE captures more advanced decentralized primitives, such as Ad Hoc Multi-Input FE, introduced in [ACF⁺20].

Here, the message space $\mathcal{M} = \{0, 1\}^*$, the key space \mathcal{K} comprises pairs (f, \mathcal{U}) where $f : \{0, 1\}^\ell \rightarrow \{0, 1\}^*$ is an ℓ -ary function for arbitrary $\ell \in \mathbb{N}$, and \mathcal{U} is a list of ℓ users.

The functionality takes as input the list of pairs $(\mathbf{pk}, m_{\mathbf{pk}})$ from all senders $\mathbf{pk} \in \mathcal{U}_M$, and the list of pairs $(\mathbf{pk}, k_{\mathbf{pk}})$ from all receivers $\mathbf{pk} \in \mathcal{U}_K$. If all the key space objects agree on a function on the inputs of the list of users \mathcal{U}_M , the functionality outputs the evaluation of the function: $F((\mathbf{pk}_i, (f_i, \mathcal{U}_i))_{\mathbf{pk}_i \in \mathcal{U}_K}, (\mathbf{pk}_j, m_j)_{\mathbf{pk}_j \in \mathcal{U}_M}) = f(m_1, \dots, m_\ell)$ if $f_i = f$ and $\mathcal{U}_i = \mathcal{U}_M$ for all i , and $|\mathcal{U}_M| = \ell$. It outputs \perp otherwise.

Limitations of DDFE. Whereas the notion of DDFE is a strong generalization of preexisting decentralized variants of Functional Encryption, capturing functionalities not covered by Ad Hoc MIFE or MAFE, it does not cover everything. Function Private [BS15] and Delegatable [CGJS15] variants of Functional Encryption have been introduced, and our definitions leave room for similar variants of DDFE. Some important cryptographic protocols, such as Private Information Retrieval, Oblivious Pseudo Random Functions, or Non-Interactive Key Exchange, similarly cannot be written as DDFE functionalities. DDFE fails to capture key exchange because its definition doesn't allow us to express cryptographic properties of a function evaluation: the idea that the result of an evaluation would "look random" cannot be written as a functionality. It also cannot capture the aforementioned two party interactive protocols because it is non-interactive by nature, while interactivity is a core requirement for PIR and OPRFs, to ensure the protocol is not run more times than any party wishes for.

2.4 DDFE Functionalities

We now give some examples of concrete functionalities. The first two will be of independent interest, but also layers to improve the security and the functionalities of the later Inner-Product DDFE constructions.

All-or-Nothing Encapsulation (AoNE) allows several parties of a group to encapsulate individual messages, that can all be extracted by anybody if and only if all the parties of this group have sent their contributions. Otherwise, the messages remain hidden. The set \mathcal{U}_M of public keys describes the group of parties and the label ℓ imposes a constraint on which encapsulations can be considered together: if for a given pair (\mathcal{U}_M, ℓ) all the parties in \mathcal{U}_M send their encapsulations, all the messages can be recovered by anybody, otherwise the messages remain hidden. Note that all the players have to agree on the pair (\mathcal{U}_M, ℓ) for their encapsulation, and any encapsulation naturally leaks that pair (\mathcal{U}_M, ℓ) .

Definition 6 (All-or-Nothing Encapsulation). *AoNE* is defined on messages of length L as follows:

$$\mathcal{K} = \emptyset \qquad \mathcal{M} = \{0, 1\}^L \times \mathcal{S}(\mathcal{PK}) \times \{0, 1\}^*.$$

Then, $F(\epsilon, (\mathbf{pk}, (x, \mathcal{U}, \ell))) = (\mathcal{U}, \ell)$ and

$$F(\epsilon, (\mathbf{pk}, m_{\mathbf{pk}})_{\mathbf{pk} \in \mathcal{U}_M}) = \begin{cases} (\mathbf{pk}, x_{\mathbf{pk}})_{\mathbf{pk} \in \mathcal{U}_M} & \text{if condition } (*) \\ \perp & \text{otherwise.} \end{cases}$$

and *AoNE* condition $(*)$ is: $\exists \ell \in \{0, 1\}^*, \forall \mathbf{pk} \in \mathcal{U}_M, m_{\mathbf{pk}} = (x_{\mathbf{pk}}, \mathcal{U}_M, \ell)$.

Decentralized Sum (DSum) allows several parties of a group to commit to values, so that their sum is automatically revealed when all the parties of this group have sent their contributions. Otherwise, the values remain hidden. The set \mathcal{U}_M of public keys describes the group of parties and the label ℓ imposes a constraint on which values can be added together: if for a given pair (\mathcal{U}_M, ℓ) all the parties in \mathcal{U}_M send their values, the sum can be recovered by anybody, otherwise the individual values remain hidden. As above, all the players have to agree on the pair (\mathcal{U}_M, ℓ) for their encryption, and any encryption naturally leaks that pair (\mathcal{U}_M, ℓ) . The terminology *sum* is an abuse, as it works for any Abelian group.

Definition 7 ($(\mathbb{A}, +)$ -Decentralized Sum). *DSum* is defined for any Abelian group $(\mathbb{A}, +)$ as follows:

$$\mathcal{K} = \emptyset \qquad \mathcal{M} = \mathbb{A} \times \mathcal{S}(\mathcal{PK}) \times \{0, 1\}^*.$$

Then, $F(\epsilon, (\mathbf{pk}, (x, \mathcal{U}, \ell))) = (\mathcal{U}, \ell)$ and

$$F(\epsilon, (\mathbf{pk}, m_{\mathbf{pk}})_{\mathbf{pk} \in \mathcal{U}_M}) = \begin{cases} \sum_{\mathbf{pk} \in \mathcal{U}_M} x_{\mathbf{pk}} & \text{if condition } (*) \\ \perp & \text{otherwise.} \end{cases}$$

and *DSum* condition $(*)$ is: $\exists \ell \in \{0, 1\}^*, \forall \mathbf{pk} \in \mathcal{U}_M, m_{\mathbf{pk}} = (x_{\mathbf{pk}}, \mathcal{U}_M, \ell)$.

Inner-Product DDFE (IP-DDFE). We now present a more advanced functionality for Inner Products. It allows senders with public key \mathbf{pk} , as part of a group \mathcal{U}_M , to encrypt inputs $\mathbf{x}_{\mathbf{pk}}$ under a label ℓ . But they maintain control on which computations will be performed on their inputs, as they all have to agree on the weights $\mathbf{y}_{\mathbf{pk}}$ to produce the functional decryption key that allows the inner-product. The set \mathcal{U}_M of public keys describes the group of parties and the label ℓ imposes a constraint on which values can be aggregated together, the set \mathcal{U}_K of public keys describes the support of the inner-product, and $(\mathbf{y}_{\mathbf{pk}})_{\mathbf{pk}}$ specifies the weights. If $\mathcal{U}_M = \mathcal{U}_K$ and all the ciphertexts are provided (by all the senders on the same pair (\mathcal{U}_M, ℓ)), with the appropriate functional decryption key (with the same $(\mathcal{U}_K, (\mathbf{y}_{\mathbf{pk}})_{\mathbf{pk}})$), one can get the inner-product value, otherwise the individual values remain hidden. As above, all the players have to agree on the pair (\mathcal{U}_M, ℓ) for their encryption, and any encryption naturally leaks that pair (\mathcal{U}_M, ℓ) . Similarly, all the players have to agree on $(\mathcal{U}_K, (\mathbf{y}_{\mathbf{pk}})_{\mathbf{pk}})$ for the functional decryption key, otherwise they are useless.

Because our construction is based on prime-order groups, we need to impose a bound on the messages and the keys to guarantee that we can perform the discrete logarithm efficiently and recover the result of the functional evaluation in polynomial time.

Definition 8 (Inner-Product DDFE.). *IP-DDFE* is defined for a dimension $d \in \mathbb{N}$ and a bound $B \in \mathbb{N}$, and the sets \mathcal{U}_M and \mathcal{U}_K must perfectly match:

$$\mathcal{K} = \{(\mathbf{y}_{\mathbf{pk}}, \mathbf{pk})_{\mathbf{pk} \in \mathcal{U}_K} \text{ where } \mathbf{y}_{\mathbf{pk}} \in [-B, B]^d \text{ and } \mathcal{U}_K \in \mathcal{S}(\mathcal{PK})\}$$

$$\mathcal{M} = [-B, B]^d \times \mathcal{S}(\mathcal{PK}) \times \{0, 1\}^*.$$

Then, $F(\epsilon, (\mathbf{pk}, (\mathbf{x}, \mathcal{U}, \ell))) = (\mathcal{U}, \ell)$ and

$$F(\epsilon, ((\mathbf{pk}, k_{\mathbf{pk}})_{\mathbf{pk} \in \mathcal{U}_K}, (\mathbf{pk}, m_{\mathbf{pk}})_{\mathbf{pk} \in \mathcal{U}_M})) = \begin{cases} \sum_{\mathbf{pk} \in \mathcal{U}_K} \mathbf{x}_{\mathbf{pk}}^\top \mathbf{y}_{\mathbf{pk}} & \text{if condition } (*) \\ \perp & \text{otherwise.} \end{cases}$$

and *IP-DDFE* condition $(*)$ is:

- $\mathcal{U}_K = \mathcal{U}_M$
- $\exists (\mathbf{y}_{\mathbf{pk}})_{\mathbf{pk} \in \mathcal{U}_K} \in \mathcal{S}([-B, B]^d), \forall \mathbf{pk}' \in \mathcal{U}_K, k_{\mathbf{pk}'} = (\mathbf{y}_{\mathbf{pk}}, \mathbf{pk})_{\mathbf{pk} \in \mathcal{U}_K}$
- $\exists \ell \in \{0, 1\}^*, \forall \mathbf{pk} \in \mathcal{U}_K, m_{\mathbf{pk}} = (\mathbf{x}_{\mathbf{pk}}, \mathcal{U}_M, \ell)$

We stress that in all the above definition, F should always be understood to be equal to \perp on inputs on which it was not explicitly defined above.

3 Notations and Assumptions

3.1 Groups

Prime Order Groups. We use a prime-order group generator GGen , a probabilistic polynomial time (PPT) algorithm that on input the security parameter 1^λ returns a description $\mathcal{G} = (\mathbb{G}, p, P)$ of an additive cyclic group \mathbb{G} of order p for a 2λ -bit prime p , whose generator is P .

We use implicit representations of group elements as introduced in [EHK⁺13]. For $a \in \mathbb{Z}_p$, define $[a] = aP \in \mathbb{G}$ as the *implicit representation* of a in \mathbb{G} . More generally, for a matrix $\mathbf{A} = (a_{ij}) \in \mathbb{Z}_p^{n \times m}$ we define $[\mathbf{A}]$ as the implicit representation of \mathbf{A} in \mathbb{G} :

$$[\mathbf{A}] := \begin{pmatrix} a_{11}P & \dots & a_{1m}P \\ \vdots & & \vdots \\ a_{n1}P & \dots & a_{nm}P \end{pmatrix} \in \mathbb{G}^{n \times m}$$

We will always use this implicit notation of elements in \mathbb{G} , i.e., we let $[a] \in \mathbb{G}$ be an element in \mathbb{G} . Note that from a random $[a] \in \mathbb{G}$, it is generally hard to compute the value a (discrete logarithm problem in \mathbb{G}). Obviously, given $[a], [b] \in \mathbb{G}$ and a scalar $x \in \mathbb{Z}_p$, one can efficiently compute $[ax] \in \mathbb{G}$ and $[a + b] = [a] + [b] \in \mathbb{G}$.

Pairing-Friendly Groups. We also use a pairing-friendly group generator PGen , a PPT algorithm that on input 1^λ returns $\mathcal{PG} = (\mathbb{G}_1, \mathbb{G}_2, \mathbb{G}_T, p, P_1, P_2, e)$, a description of asymmetric pairing-friendly groups where $\mathbb{G}_1, \mathbb{G}_2, \mathbb{G}_T$ are additive cyclic groups of order p for a 2λ -bit prime p , P_1 and P_2 are generators of \mathbb{G}_1 and \mathbb{G}_2 , respectively, and $e : \mathbb{G}_1 \times \mathbb{G}_2 \rightarrow \mathbb{G}_T$ is an efficiently computable (non-degenerate) bilinear map. Define $P_T := e(P_1, P_2)$, which is a generator of \mathbb{G}_T . We again use implicit representation of group elements. For $s \in \{1, 2, T\}$ and $a \in \mathbb{Z}_p$, define $[a]_s = aP_s \in \mathbb{G}_s$ as the implicit representation of a in \mathbb{G}_s . Given $[a]_1, [b]_2$, one can efficiently compute $[ab]_T$ using the pairing e . For two matrices \mathbf{A}, \mathbf{B} with matching dimensions define $e([\mathbf{A}]_1, [\mathbf{B}]_2) := [\mathbf{AB}]_T \in \mathbb{G}_T$.

3.2 Intractability Assumptions

Definition 9 (Computational Diffie-Hellman Assumption). *The CDH assumption states that, in a prime-order group $\mathcal{G} \xleftarrow{\$} \text{GGen}(1^\lambda)$, no PPT adversary can compute $[xy]$, from $[x]$ and $[y]$ for $x, y \xleftarrow{\$} \mathbb{Z}_p$, with non-negligible success probability.*

Equivalently, this assumption states it is hard to compute $[a^2]$ from $[a]$ for $a \xleftarrow{\$} \mathbb{Z}_p$. This comes from the fact that $4[xy] = [(x + y)^2] - [(x - y)^2]$.

Definition 10 (Decisional Diffie-Hellman Assumption). *The DDH assumption states that, in a group $\mathcal{G} \xleftarrow{\$} \text{GGen}(1^\lambda)$, no PPT adversary can distinguish between the two following distributions with non-negligible advantage: $\{([a], [r], [ar]) \mid a, r \xleftarrow{\$} \mathbb{Z}_p\}$ and $\{([a], [r], [s]) \mid a, r, s \xleftarrow{\$} \mathbb{Z}_p\}$.*

Equivalently, this assumption states it is hard to distinguish, knowing $[a]$, a random element from the span of $[a]$ for $\mathbf{a} = \begin{pmatrix} 1 \\ a \end{pmatrix}$, from a random element in \mathbb{G}^2 : $[\mathbf{a}] \cdot r = [ar] = \begin{pmatrix} [r] \\ [ar] \end{pmatrix} \approx \begin{pmatrix} [r] \\ [s] \end{pmatrix}$.

Definition 11 (Decisional Bilinear Diffie Hellman Assumption). *The DBDH assumption states that, in a pairing group $\mathcal{PG} \stackrel{\$}{\leftarrow} \text{PGGen}(1^\lambda)$, for any PPT adversary, the following advantage is negligible, where the probability distribution is over $a, b, c, s \stackrel{\$}{\leftarrow} \mathbb{Z}_p$:*

$$\text{Adv}_{\mathcal{PG}}^{\text{DBDH}}(\mathcal{A}) = |\Pr[1 \leftarrow \mathcal{A}(\mathcal{PG}, [a]_1, [b]_1, [b]_2, [c]_2, [abc]_T)] - \Pr[1 \leftarrow \mathcal{A}(\mathcal{PG}, [a]_1, [b]_1, [b]_2, [c]_2, [s]_T)]|.$$

Definition 12 (Q -fold DBDH). *For any integer Q , the Q -fold DBDH assumption states for any PPT adversary, the following advantage is negligible, where the probability distribution is over $a, b, c_i, s_i \stackrel{\$}{\leftarrow} \mathbb{Z}_p$:*

$$\text{Adv}_{\mathcal{PG}}^{Q\text{-DBDH}}(\mathcal{A}) = |\Pr[1 \leftarrow \mathcal{A}(\mathcal{PG}, [a]_1, [b]_1, [b]_2, \{[c_i]_2, [abc_i]_T\}_{i \in [Q]})] - \Pr[1 \leftarrow \mathcal{A}(\mathcal{PG}, [a]_1, [b]_1, [b]_2, \{[c_i]_2, [s_i]_T\}_{i \in [Q]})]|.$$

This Q -fold DBDH assumption is equivalent to classical DBDH assumption:

Lemma 13 (Random Self Reducibility of DBDH). *For any adversary \mathcal{A} against the Q -fold DBDH, running within time t , there exists an adversary \mathcal{B} running within time $t + 2Q(t_{\mathbb{G}_T} + t_{\mathbb{G}_2})$, where $t_{\mathbb{G}_T}$ and $t_{\mathbb{G}_2}$ denote respectively the time for an exponentiation in \mathbb{G}_T and \mathbb{G}_2 (we only take into account the time for exponentiations here), such that*

$$\text{Adv}_{\mathcal{PG}}^{Q\text{-DBDH}}(\mathcal{A}) \leq \text{Adv}_{\mathcal{PG}}^{\text{DBDH}}(\mathcal{B}).$$

Proof. Upon receiving a DBDH challenge $(\mathcal{PG}, [a]_1, [b]_1, [b]_2, [c]_2, [s]_T)$, \mathcal{B} samples $\alpha_i, c'_i \stackrel{\$}{\leftarrow} \mathbb{Z}_p$ computes $[c_i]_2 := [\alpha_i \cdot c]_2 + [c'_i]_2$, $[s_i]_T := [\alpha_i \cdot s]_T + [c_i \cdot ab]_T$ for all $i \in [Q]$, and gives the challenge $(\mathcal{PG}, [a]_1, [b]_1, [b]_2, \{[c_i]_2, [s_i]_T\}_{i \in [Q]})$ to \mathcal{A} . \square

3.3 Non-Interactive Key Exchange

We give a definition of Non-Interactive Key Exchange below. This is a rephrasing of the m-CKS-heavy model (with dishonest key registrations) as originally introduced in [CKS08] and further refined in [FHKP13].

Definition 14 (Non-Interactive Key Exchange). *A NIKE scheme consists of three PPT algorithms:*

- **Setup**(λ): *Generates and outputs public parameters pp . Those parameters are implicit arguments to all the other algorithms;*
- **KeyGen**(\cdot): *Generates and outputs a party's public key $\text{pk} \in \mathcal{PK}$ and the corresponding secret key sk_{pk} ;*
- **SharedKey**($\text{pk}, \text{sk}_{\text{pk}'}$): *Takes as input a public key and a secret key corresponding to a different public key. Deterministically outputs a shared key K .*

Correctness: We require that, for all security parameters $\lambda \in \mathbb{N}$, it holds that:

$$\Pr [\text{SharedKey}(\text{pk}, \text{sk}_{\text{pk}'}) = \text{SharedKey}(\text{pk}', \text{sk}_{\text{pk}})] = 1,$$

where the probability is taken over $\text{pp} \leftarrow \text{Setup}(\lambda)$, $(\text{pk}, \text{sk}_{\text{pk}}) \leftarrow \text{KeyGen}(\cdot)$, $(\text{pk}', \text{sk}_{\text{pk}'}) \leftarrow \text{KeyGen}(\cdot)$.

Definition 15 (Security Game for NIKE). *Let us consider a NIKE scheme. No adversary \mathcal{A} should be able to win the following security game against a challenger \mathcal{C} , with unlimited and adaptive access to the oracles QNewHonest , QReveal , QTest , and QCorrupt described below:*

- *Initialize:* the challenger \mathcal{C} runs the setup algorithm $\text{pp} \leftarrow \text{Setup}(\lambda)$ and chooses a random bit $b \stackrel{\$}{\leftarrow} \{0, 1\}$. It initializes the set \mathcal{H} of honest participants to \emptyset . It provides pp to the adversary \mathcal{A} ;

- Participant creation queries $\text{QNewHonest}()$: the challenger \mathcal{C} runs the KeyGen algorithm $(\text{pk}, \text{sk}_{\text{pk}}) \leftarrow \text{KeyGen}()$ to simulate a new participant, stores the association $(\text{pk}, \text{sk}_{\text{pk}})$ in the set \mathcal{H} of honest keys, and returns pk to the adversary;
- Reveal queries $\text{QReveal}(\text{pk}, \text{pk}')$: Requires that at least one of pk and pk' be in \mathcal{H} . Without loss of generality assume it is pk . The challenger returns $\text{SharedKey}(\text{pk}', \text{sk}_{\text{pk}})$;
- Test queries $\text{QTest}(\text{pk}, \text{pk}')$: Requires that both pk and pk' were generated via QNewHonest .
 - If $b = 0$, the challenger returns $\text{SharedKey}(\text{pk}', \text{sk}_{\text{pk}})$;
 - If $b = 1$, the challenger returns a (uniformly) random value, which it stores so it can consistently answer further queries to $\text{QTest}(\text{pk}, \text{pk}')$ or $\text{QTest}(\text{pk}', \text{pk})$
- Corruption queries $\text{QCorrupt}(\text{pk})$: Recovers the secret key sk associated to pk from \mathcal{H} and outputs it, then removes the key-pair from \mathcal{H} . If pk is not associated with any secret key (i.e. it is not in \mathcal{H}), then nothing is returned;
- Finalize: \mathcal{A} provides its guess b' on the bit b , and this procedure outputs the result β of the security game, according to the analysis given below which aims at preventing trivial wins.

Finalize outputs the bit $\beta = (b' = b)$ unless a QCorrupt query was made for any public key which was involved in a query to QTest , or a QReveal query was made for a pair of public keys which was also involved in a QTest query, in which case a random bit β is returned.

We say *NIKE* is secure if for any adversary \mathcal{A} , the following advantage is negligible:

$$\text{Adv}_{\text{NIKE}}(\mathcal{A}) = 2 \times |\Pr[\beta = 1] - 1/2|.$$

Definition 16 (Static Security Game for NIKE). We define a static variant of the security game above in which the adversary does not have access to the QCorrupt oracle, which means all parties created by the challenger will remain honest, and the only corrupt parties are entirely managed by the adversary.

3.4 Definition of Symmetric Key Encryption

A symmetric key encryption $\text{SKE} = (\text{SEnc}, \text{SDec})$ with key space \mathcal{K} is defined as:

- $\text{SEnc}(K, m)$: given a key K and a message m , outputs a ciphertext ct ;
- $\text{SDec}(K, \text{ct})$: given a key K and a ciphertext ct , outputs a plaintext.

Correctness. For all m in the message space and all K in the key space, we must have the equality $\text{SDec}(K, \text{SEnc}(K, m)) = m$.

Security. We say SKE is secure if for any PPT adversary \mathcal{A} , the following advantage is negligible:

$$\text{Adv}_{\text{SKE}}(\mathcal{A}) = \left| 2 \times \Pr \left[b' = b : \begin{array}{l} K \xleftarrow{\$} \mathcal{K}, b \xleftarrow{\$} \{0, 1\} \\ b' \leftarrow \mathcal{A}(1^\lambda)^{\text{QLeftRight}(\cdot, \cdot)} \end{array} \right] - 1 \right|,$$

where the oracle QLeftRight , when queried on m_0, m_1 , returns $\text{SEnc}(K, m_b)$.

One-Time Security. We say SKE is One-Time Secure if the above security holds for only one QLeftRight -oracle query. Note that if the key space is larger than the message space, one can simply use the one-time pad to build a One-Time Secure symmetric encryption. Otherwise, a pseudo-random generator can stretch the key to the right length.

3.5 Single-Input Functional Encryption

For some of our constructions, we will need a instantiation of single-input Functional Encryption (for a specific functionalities). A Functional encryption scheme for a family of functions \mathcal{F} consists of the following PPT algorithms:

- $\text{KeyGen}(\lambda)$: on input a security parameter, it outputs a master secret key msk and a public key pk .
- $\text{Encrypt}(\text{pk}, m)$: outputs a ciphertext ct .
- $\text{DKeyGen}(\text{msk}, f)$: on input the master secret key and a function $f \in \mathcal{F}$, it outputs a decryption key dk_f .
- $\text{Dec}(\text{ct}, \text{dk}_f)$: deterministic algorithm that returns a message or a rejection symbol \perp if it fails.

Correctness. For any message m , and any function f in the family \mathcal{F} , we have: $\Pr[\text{Dec}(\text{ct}, \text{dk}_f) = f(m)] = 1$, where the probability is taken over $(\text{msk}, \text{mpk}) \leftarrow \text{Setup}(\lambda)$, $\text{ct} \leftarrow \text{Encrypt}(\text{msk}, m)$, and $\text{dk}_f \leftarrow \text{DKeyGen}(\text{msk}, f)$.

Indistinguishability. The security notion is defined by a classical indistinguishability game:

Definition 17 (IND-Security Game for FE). *Let FE be a functional encryption scheme. No adversary \mathcal{A} should be able to win the following security game:*

- *Initialize:* runs $(\text{msk}, \text{mpk}) \leftarrow \text{Setup}(\lambda)$, choose a random bit $b \xleftarrow{\$} \{0, 1\}$ and returns mpk to \mathcal{A} .
- $\text{QLeftRight}(m_0, m_1)$: on input two messages (m_0, m_1) , returns $\text{Enc}(\text{msk}, m_b)$.
- $\text{QDKeyGen}(f)$: on input a function $f \in \mathcal{F}$, returns $\text{DKeyGen}(\text{msk}, f)$.
- *Finalize:* from the guess b' of \mathcal{A} on the bit b , it outputs the bit $\beta = (b' = b)$ unless some f was queried to QDKeyGen and (m_0, m_1) was queried to QLeftRight such that $f(m_0) \neq f(m_1)$, in which case it outputs a uniformly random bit β .

The adversary \mathcal{A} has unlimited and adaptive access to the left-right encryption oracle QLeftRight , and to the key generation oracle QDKeyGen . We say FE is IND-secure if for any adversary \mathcal{A} , $\text{Adv}_{\text{FE}}^{\text{IND}}(\mathcal{A}) = |2 \times \Pr[\beta = 1] - 1|$ is negligible.

We can also define a weaker selective variant, where pairs (m_0, m_1) to QLeftRight -queries are known from the beginning.

Identity-Based Encryption. Here we define the functionality that corresponds to Identity-Based Encryption, originally envisioned in [Sha84], and first realized in [BF01, Coc01]. The functionality is described by an identity space \mathcal{I} , which can be of exponential size. Each function is described by an identity $\text{id} \in \mathcal{I}$, and given as input a pair (m, id') where m is a payload, and $\text{id}' \in \mathcal{I}$ is an identity, the function outputs m if $\text{id} = \text{id}'$, nothing otherwise.

Inner Product Functionality. For any dimension $d \in \mathbb{N}$ and cyclic group \mathbb{G} of prime order p , the inner product functionality corresponds to the set of functions described by a vector $\mathbf{y} \in \mathbb{Z}_p^d$ that on input a vector $[\mathbf{x}] \in \mathbb{G}^d$, outputs $[\mathbf{x}^\top \mathbf{y}]$. FE schemes for the inner-product functionality were originally introduced in [ABDP15], later in [ALS16] with adaptive security.

We will make use of the following property, satisfied by several FE schemes, including [ABDP15, ALS16]. For concreteness we recall the scheme from [ALS16] in Appendix A.

Property 18 (Linear Homomorphism). An FE scheme for the inner product functionality (IP-FE.Setup, IP-FE.Encrypt, IP-FE.DKeyGen, IP-FE.Dec) satisfies the linear homomorphism property if there exists a PPT algorithm Add such that for all $\mathbf{x}, \mathbf{x}' \in \mathbb{Z}_p^d$, the following are identically distributed:

$$\begin{aligned} & \left(\text{IP-FE.Encrypt}(\text{IP-FE.pk}, \mathbf{x}), \text{IP-FE.Encrypt}(\text{IP-FE.pk}, \mathbf{x} + \mathbf{x}') \right) \\ & \quad \text{and} \\ & \left(\text{IP-FE.Encrypt}(\text{IP-FE.pk}, \mathbf{x}), \text{Add} \left(\text{IP-FE.Encrypt}(\text{IP-FE.pk}, \mathbf{x}), \mathbf{x}' \right) \right), \end{aligned}$$

where $(\text{IP-FE.pk}, \text{IP-FE.sk}) \leftarrow \text{IP-FE.Setup}(\lambda)$.

4 All-or-Nothing Encapsulation from IBE

4.1 Technical Overview

Our generic construction only requires an IBE. Messages are encrypted under the public key of each member of the group successively, using the set of participants \mathcal{U}_M and the label ℓ as the identity. The $|\mathcal{U}_M|$ -layers deep encryption is accompanied by the functional decryption key of the IBE for the encrypting participant and the same identity. The only way to recover the messages is to gather all the decryption keys in order to decrypt all layers of IBE encryption: this requires having access to all the ciphertexts. IBE is a well-studied primitive, which admits constructions from multiple hardness assumptions, including pairings [BF01], LWE [GPV08], or more recently the CDH assumption [DG17]. This directly implies feasibility of AoNE from these assumptions. To keep the size of the ciphertext polynomial in the number of users, we use rate-1 IBE, using hybrid encryption. In Section 5 we give a more efficient construction directly from pairings, inspired by the IBE from [BF01].

4.2 A Generic Construction of All-or-Nothing Encapsulation

Our construction uses an Identity-Based encryption scheme IBE.

- $\text{Setup}(\lambda)$: Return $pp \leftarrow \text{IBE.Setup}(\lambda)$
- $\text{KeyGen}()$: Return $(\text{pk}, \text{sk}_{\text{pk}}) \leftarrow \text{IBE.KeyGen}()$.
- $\text{Encrypt}(\text{sk}_{\text{pk}}, m)$: Parse $m = (x_{\text{pk}}, \mathcal{U}_M, \ell)$ where $x_{\text{pk}} \in \{0, 1\}^L$, $\mathcal{U}_M \in \mathcal{S}(\mathcal{PK})$, and $\ell \in \{0, 1\}^*$. If $\text{pk} \notin \mathcal{U}_M$, return \perp . Let $n = |\mathcal{U}_M|$ be the cardinal of \mathcal{U}_M , and, for some universally accepted order, number the elements in \mathcal{U}_M as $\mathcal{U}_M = \{\text{pk}_1, \dots, \text{pk}_n\}$.
Let $\alpha_{\text{pk},0} = x_{\text{pk}}$, and for i going from 1 to n , compute

$$\alpha_{\text{pk},i} := \text{IBE.Encrypt}(\text{pk}_i, (\alpha_{\text{pk},i-1}, \mathcal{U}_M || \ell)).$$

We write $\alpha_{\text{pk},\mathcal{U}_M,\ell} = \alpha_{\text{pk},n}$. Compute $\gamma_{\text{pk},\mathcal{U}_M,\ell} = \text{IBE.DKeyGen}(\text{sk}_{\text{pk}}, \mathcal{U}_M || \ell)$.

Return $(\alpha_{\text{pk},\mathcal{U}_M,\ell}, \gamma_{\text{pk},\mathcal{U}_M,\ell}, \mathcal{U}_M, \ell)$.

- $\text{DKeyGen}(\text{sk}, k)$: There are no keys in this functionality, so no DKeyGen;
- $\text{Decrypt}(\epsilon, (\text{ct}_{\text{pk}})_{\text{pk} \in \mathcal{U}_M})$: Parse the ciphertexts, for all $\text{pk} \in \mathcal{U}_M$, as

$$\text{ct}_{\text{pk}} = (\alpha_{\text{pk},\mathcal{U}_M,\ell}, \gamma_{\text{pk},\mathcal{U}_M,\ell}, \mathcal{U}_M, \ell),$$

with common (\mathcal{U}_M, ℓ) (otherwise return \perp). For each $\text{pk} \in \mathcal{U}_M$, we recover x_{pk} as follows: with $\mathcal{U}_M = \{\text{pk}_1, \dots, \text{pk}_n\}$, recompute the $\alpha_{\text{pk},i}$ for i going from n to 0 as $\alpha_{\text{pk},n} = \alpha_{\text{pk},\mathcal{U}_M,\ell}$ and $\alpha_{\text{pk},i} = \text{IBE.Decrypt}(\gamma_{\text{pk}_i,\mathcal{U}_M,\ell}, \alpha_{\text{pk},i+1})$. Output $(\text{pk}, x_{\text{pk}})_{\text{pk} \in \mathcal{U}_M}$.

Correctness: Correctness follows immediately from the correctness of IBE.

Remark 19 (Rate-1 IBE). To avoid ciphertexts having length exponential in $|\mathcal{U}_M|$, we require that IBE has rate-1 encryption. That is, the ciphertext has the same size as the plaintext plus a polynomial in the security parameter. This can be obtained generically via hybrid encryption: the IBE is used to encrypt a symmetric key, that is used to encrypt the actual message. Assuming such properties of the IBE scheme, our ciphertexts have length linear in $|\mathcal{U}_M|$.

Remark 20. The astute reader will have noticed that the $\gamma_{\text{pk},\mathcal{U}_M,\ell}$ seem to be playing the role of a Functional Decryption Key. Indeed, AoNE could have been defined with keys allowing decryption of the ciphertext if the appropriate key shares (i.e., for that pair (\mathcal{U}_M, ℓ)) are contributed by all parties. However, our applications of AoNE are such that we would always end up giving out the key share with the corresponding ciphertext, so we gave a definition which is more practical for our uses and may allow constructions in settings where the alternative with keys would be harder to design.

Remark 21. Note that while we show here how to construct AoNE from IBE, it's also possible to construct IBE from AoNE. A possible construction uses only two AoNE identities, one of which creates AoNE ciphertexts that serve as IBE ciphertexts, while the other creates AoNE ciphertexts that serve as IBE functional keys. The secret key for the first identity is made public (it is part of the IBE's public key) while that for the latter remains private. Identities are encoded as labels, and groups are always chosen as the pair of identities. Now to recover the message behind a ciphertext, even generated with the known AoNE secret key of the first identity, one needs an AoNE ciphertext from the second identity for the same label/identity, which effectively acts as an IBE secret key.

4.3 Security Proof

Theorem 22 (IND-Security of AoNE). *The All-or-Nothing Encapsulation scheme described in Section 4.2 is IND-secure (as per Definition 3) assuming the IBE scheme is IND-secure (as per Definition 17).*

The proof can be found in Appendix B.1.

5 All-or-Nothing Encapsulation from Bilinear Maps

5.1 Technical Overview

This construction is essentially an instantiation of the generic construction given in Section 4.2 using Boneh and Franklin's IBE [BF01]. However, we make a few optimizations exploiting the structure of the Boneh-Franklin IBE (BF) to achieve short ciphertexts. First, we use the IBE as a Key-Encapsulation Mechanism to generate a symmetric key, which we then use to encrypt the message. Second, we exploit the randomness reusability of El Gamal-like schemes, from which BF benefits, to only commit to a randomness once. The size difference between the message and the ciphertext in BF comes entirely from that commitment to randomness, so sharing it across all encryptions removes the dependence on the size of the set of participants in the size of the ciphertext.

We provide a direct security analysis of the resulting scheme in Section 5.3.

5.2 A Construction of All-or-Nothing Encapsulation from Bilinear Maps

Our construction uses pairing-friendly groups, a hash function modeled as a random oracle in the security analysis, and a (One-Time Secure) symmetric encryption scheme.

- **Setup**(λ): Generate $\mathcal{PG} = (\mathbb{G}_1, \mathbb{G}_2, \mathbb{G}_T, p, P_1, P_2, e) \xleftarrow{\$} \text{PGGen}(1^\lambda)$, a full domain hash function \mathcal{H} from $\{0, 1\}^*$ into \mathbb{G}_1 , and return $\text{pp} = (\mathcal{PG}, \mathcal{H})$. For the sake of clarity, for any input x , we will denote $\mathcal{H}(x) = h_x P_1 = [h_x]_1$, where h_x is the unknown discrete logarithm.
- **KeyGen**(\cdot): Sample $t_{\text{pk}} \xleftarrow{\$} \mathbb{Z}_p$ and return $(\text{pk}, \text{sk}_{\text{pk}}) = ([t_{\text{pk}}]_2, t_{\text{pk}})$.
- **Encrypt**(sk_{pk}, m): Parse $\text{sk}_{\text{pk}} = t_{\text{pk}} \in \mathbb{Z}_p$ and $m = (x_{\text{pk}}, \mathcal{U}_M, \ell)$ where $x_{\text{pk}} \in \{0, 1\}^L$, $\mathcal{U}_M \in \mathcal{S}(\mathcal{PK})$, and $\ell \in \{0, 1\}^*$. If $\text{pk} \notin \mathcal{U}_M$, return \perp . Otherwise, sample $r_{\text{pk}} \xleftarrow{\$} \mathbb{Z}_p$ and compute the symmetric key $K_{\text{pk}, \mathcal{U}_M, \ell}$ as

$$e \left(\mathcal{H}(\mathcal{U}_M || \ell), r_{\text{pk}} \cdot \left(\sum_{\text{pk}' \in \mathcal{U}_M} \text{pk}' \right) \right) = \left[h_{\mathcal{U}_M || \ell} \cdot r_{\text{pk}} \cdot \sum_{\text{pk}' \in \mathcal{U}_M} t_{\text{pk}'} \right]_T,$$

and use it to encrypt x_{pk} as $c_{\text{pk}} = \text{SEnc}(K_{\text{pk}, \mathcal{U}_M, \ell}, x_{\text{pk}})$. Compute its share $S_{\text{pk}, \mathcal{U}_M, \ell} = t_{\text{pk}} \cdot \mathcal{H}(\mathcal{U}_M || \ell) = [t_{\text{pk}} \cdot h_{\mathcal{U}_M || \ell}]_1$, and output the ciphertext $\text{ct}_{\text{pk}} = (c_{\text{pk}}, [r_{\text{pk}}]_2, S_{\text{pk}, \mathcal{U}_M, \ell}, \mathcal{U}_M, \ell)$.

- **DKeyGen**(sk, k): There are no keys in this functionality, so no **DKeyGen**;
- **Decrypt**($\epsilon, (\text{ct}_{\text{pk}})_{\text{pk} \in \mathcal{U}_M}$): Parse the ciphertexts, for all $\text{pk} \in \mathcal{U}_M$, as $\text{ct}_{\text{pk}} = (c_{\text{pk}}, [r_{\text{pk}}]_2, S_{\text{pk}, \mathcal{U}_M, \ell}, \mathcal{U}_M, \ell)$, with common (\mathcal{U}_M, ℓ) . For each $\text{pk} \in \mathcal{U}_M$, compute

$$K_{\text{pk}, \mathcal{U}_M, \ell} = e \left(\sum_{\text{pk}' \in \mathcal{U}_M} S_{\text{pk}', \mathcal{U}_M, \ell}, [r_{\text{pk}}]_2 \right) = \left[h_{\mathcal{U}_M || \ell} \cdot r_{\text{pk}} \cdot \sum_{\text{pk}' \in \mathcal{U}_M} t_{\text{pk}'} \right]_T$$

and recover x_{pk} as $x_{\text{pk}} = \text{SDec}(K_{\text{pk}, \mathcal{U}_M, \ell}, c_{\text{pk}})$.

Correctness: First, note that the use of $K_{pk, \mathcal{U}_M, \ell}$ is consistent across `Encrypt` and `Decrypt`. Then, the two evaluations correspond to $\left[h_{\mathcal{U}_M \| \ell} \cdot r_{pk} \cdot \sum_{pk' \in \mathcal{U}_M} t_{pk'} \right]_T$. Now correctness immediately follows from the correctness of the underlying symmetric encryption scheme.

Remark 23. Note that the sum $\sum_{pk' \in \mathcal{U}_M} S_{pk', \mathcal{U}_M, \ell}$ is common to all ciphertexts for the same pair (\mathcal{U}_M, ℓ) and can thus be precomputed and reused, such that n messages can be recovered in time $\mathcal{O}(n)$.

5.3 Security Proof

Theorem 24 (IND-Security of AoNE). *The All-or-Nothing Encapsulation scheme described in Section 5.2 is IND-secure (as per Definition 3) under the DBDH assumption, in the random oracle model.*

The proof can be found in Appendix B.2.

6 Decentralized Sum

6.1 Technical Overview

The starting point of our construction is the "Sum-of-PRFs" technique used by Chase and Chow [CC09]. The technique aims to enable a set of parties to evaluate local PRFs for a common label ℓ , such that the sum of their local PRFs is zero. It relies on shared seeds between each pair of participants, that are computed on-the-fly using Non-Interactive Key Exchange. Those PRFs can then be added to each participant's input, masking the individual contribution but revealing their sum, because adding the masked ciphertexts causes the PRF evaluation to cancel out.

Remarkably, this is not enough to achieve IND security in the DDFE setting. As such, the random mask would be a deterministic function of the set of participants \mathcal{U}_M and the label ℓ . So, repeated `QLeftRight` queries to the same pair (\mathcal{U}_M, ℓ) with different messages would enable an adversary to break security, simply by subtracting two ciphertexts associated with the same pair (\mathcal{U}_M, ℓ) so as to remove the identical masks. This issue can be addressed with a layer of AoNE encryption. Since our AoNE construction is asymmetric and its encryption is randomized, the layer prevents the adversary from combining ciphertexts for the same pair (\mathcal{U}_M, ℓ) in a meaningful way. Only when all the ciphertexts are revealed can the adversary remove the AoNE layer, and get access to the underlying ciphertexts. In that case however, the information recovered by the adversary is part of the information revealed by the functionality. For instance, the adversary can subtract two deterministic ciphertexts to obtain the different of the underlying messages. This information can also be learnt by subtracting two sums that are revealed by correctness of the scheme. In general, we show that when the AoNE layer can be removed, the Finalize condition imposes sufficient constraints on the adversary's queries that trivial attacks are no longer on the table.

Moreover, the AoNE layer that lets us achieve full IND security, instead of having to settle for sym-IND security, since, as explained, the AoNE is an asymmetric form of encryption.

6.2 A Generic Construction of Decentralized Sum DDFE for $(\mathbb{A}, +)$

For our construction, we assume a NIKE scheme `NIKE`, an All-or-Nothing Encapsulation scheme `AoNE` for messages of length the size of an element of \mathbb{A} , and a PRF family $(\mathcal{F}_K)_K$ that takes keys from the NIKE and messages from $\{0, 1\}^*$ and outputs pseudo-random elements in \mathbb{A} .

- `Setup`(λ): Run `NIKE.pp` \leftarrow `NIKE.Setup`(λ), `AoNE.pp` \leftarrow `AoNE.Setup`(λ), and output `pp` = (`NIKE.pp`, `AoNE.pp`);

- **KeyGen()**: Run the **KeyGen** algorithms from the **NIKE** and the **AoNE**:

$$(\text{NIKE.pk}, \text{NIKE.sk}_{\text{pk}}) \leftarrow \text{NIKE.KeyGen}(),$$

$$(\text{AoNE.pk}, \text{AoNE.sk}_{\text{pk}}) \leftarrow \text{AoNE.KeyGen}(),$$

and output the key pair

$$(\text{pk}, \text{sk}_{\text{pk}}) = ((\text{NIKE.pk}, \text{AoNE.pk}), (\text{NIKE.sk}_{\text{pk}}, \text{AoNE.sk}_{\text{pk}}));$$

- **Encrypt**(sk_{pk}, m): Parse m as (x, \mathcal{U}_M, ℓ) , with $x \in \mathbb{A}$, $\mathcal{U}_M \in \mathcal{S}(\mathcal{PK})$, and $\ell \in \{0, 1\}^*$. Let pk be our encryptor's public key². If $\text{pk} \notin \mathcal{U}_M$, then return \perp . Otherwise, for all $\text{pk}' = (\text{NIKE.pk}', \text{AoNE.pk}') \in \mathcal{U}_M$ such that $\text{pk}' \neq \text{pk}$, compute $K_{\text{pk}, \text{pk}'} = \text{NIKE.SharedKey}(\text{NIKE.sk}_{\text{pk}}, \text{NIKE.pk}')$ and $r_{\text{pk}, \text{pk}', \mathcal{U}_M, \ell} = \mathcal{F}_{K_{\text{pk}, \text{pk}'}}(\mathcal{U}_M || \ell)$. Compute $c_{\text{pk}} = x + \sum_{\text{pk}' < \text{pk}} r_{\text{pk}, \text{pk}', \mathcal{U}_M, \ell} - \sum_{\text{pk}' > \text{pk}} r_{\text{pk}, \text{pk}', \mathcal{U}_M, \ell}$, where the sums are on $\text{pk}' \in \mathcal{U}_M$, on which a total ordering is defined. Return

$$\text{ct}_{\text{pk}} = (\text{AoNE.Encrypt}(\text{AoNE.sk}_{\text{pk}}, (c_{\text{pk}}, \mathcal{U}_M, \ell)), \mathcal{U}_M, \ell);$$

- **DKeyGen**(sk, k): There are no keys in this functionality, so no **DKeyGen**;
- **Decrypt**($\epsilon, (\text{ct}_{\text{pk}})_{\text{pk} \in \mathcal{U}_M}$): Get $(c_{\text{pk}})_{\text{pk} \in \mathcal{U}_M} = \text{AoNE.Decrypt}(\epsilon, (\text{ct}_{\text{pk}})_{\text{pk} \in \mathcal{U}_M})$, and return $\sum_{\text{pk} \in \mathcal{U}_M} c_{\text{pk}}$.

Correctness: The (c_{pk}) should be consistent between **Encrypt** and **Decrypt** by correctness of **AoNE**. Besides:

$$\begin{aligned} \sum_{\text{pk} \in \mathcal{U}_M} \text{ct}_{\text{pk}} &= \sum_{\text{pk} \in \mathcal{U}_M} \left(x_{\text{pk}} + \sum_{\text{pk}' < \text{pk}, \text{pk}' \in \mathcal{U}_M} r_{\text{pk}, \text{pk}', \mathcal{U}_M, \ell} - \sum_{\text{pk}' > \text{pk}} r_{\text{pk}, \text{pk}', \mathcal{U}_M, \ell} \right) \\ &= \sum_{\text{pk} \in \mathcal{U}_M} x_{\text{pk}} + \sum_{\substack{\text{pk}, \text{pk}' \in \mathcal{U}_M \\ \text{pk}' < \text{pk}}} r_{\text{pk}, \text{pk}', \mathcal{U}_M, \ell} - r_{\text{pk}, \text{pk}', \mathcal{U}_M, \ell} = \sum_{\text{pk} \in \mathcal{U}_M} x_{\text{pk}} \end{aligned}$$

by correctness of **NIKE**.

6.3 Security Proof

Theorem 25 (IND-Security of DSum). *The Decentralized Sum scheme described in Section 6.2 is IND-secure (as per Definition 3) so long as NIKE is IND-secure (as per Definition 15) and $(\mathcal{F}_K)_K$ is a secure PRF family.*

The proof can be found in Appendix B.3.

7 Inner-Product DDFE

7.1 Technical Overview

Our starting point is Chotard *et al.*'s Inner-Product MCFE [CDG⁺18]: as they do, we use a Random Oracle to generate shared randomness across participants for a given label ℓ (in our case a (\mathcal{U}_M, ℓ) pair). However, their construction has several drawbacks, which we overcome:

1. Their security game requires that if one ciphertext is queried for a label ℓ , all such ciphertexts must be queried (for the same label ℓ and for all other honest parties). We lift this requirement by protecting ciphertexts with a layer of **AoNE**.

² Depending on the details of **NIKE** and **AoNE** it may be necessary to explicitly include pk in sk_{pk} to ensure the following check can be performed.

2. Their **Encrypt** algorithm is a deterministic function of the message and the label ℓ , and thus they do not tolerate repeated queries to the same participant for the same label. We address this by adding a layer of IP-FE, which randomizes ciphertexts. IP-FE keys are provided in our **KeyGen** algorithm, and they are protected by an AoNE layer to ensure ciphertexts can only be decrypted once the all the partial functional decryption keys are present.
3. Their scheme, being MCFE, only works in the context of a fixed group. We show how using a PRF to dynamically generate independent secret keys for different groups removes this constraint.
4. To enable non-interactive generation of functional decryption keys in DMCFE, they introduce pairings, and perform message-related operations in \mathbb{G}_1 while key-related operations take place in \mathbb{G}_2 . Instead, we use our DSum to enforce proper key aggregation, which simplifies the scheme to a pairing-free group³.
 DSum has the added benefit that it is a DDFE functionality and thus non-interactive, meaning our Inner-Product scheme is also non-interactive, while their DMCFE has an interactive setup.

7.2 A construction of IP-DDFE

To build our IP-DDFE, we use a cyclic group \mathbb{G} of prime order p where DDH holds, a random oracle $\mathcal{H} : \{0, 1\}^* \rightarrow \mathbb{G}$, an single-input FE for the inner product functionality, where each function is described by a vector $\mathbf{y} \in \mathbb{Z}_p^d$, and on input a vector $[\mathbf{x}] \in \mathbb{G}^d$, outputs $[\mathbf{x}^\top \mathbf{y}]$. We require that IP-FE is IND secure, and satisfies Property 18. We also use an All-or-Nothing Encapsulation scheme AoNE, a Distributed Sum DSum over $(\mathbb{Z}_p, +)$, and a PRF family $(\mathcal{F}_K)_K$ that outputs in \mathbb{Z}_p^d .

- **Setup**(λ): Generate $\mathcal{G} = (\mathbb{G}, p, P) \xleftarrow{\$} \text{GGen}(1^\lambda)$. Generate a full domain hash function $\mathcal{H} : \{0, 1\}^* \rightarrow \mathbb{G}$. Compute $\text{AoNE.pp} \leftarrow \text{AoNE.Setup}(\lambda)$ and $\text{DSum.pp} \leftarrow \text{DSum.Setup}(\lambda)$. Return:

$$\text{pp} = (\mathcal{G}, \mathcal{H}, \text{NIKE.pp}, \text{AoNE.pp}).$$

- **KeyGen**(λ): Sample the keys
 - a PRF key K ,
 - IP-FE keys $(\text{IP-FE.pk}, \text{IP-FE.sk}_{\text{pk}}) \leftarrow \text{IP-FE.KeyGen}(\mathbb{G}, d)$,
 - AoNE keys $(\text{AoNE.pk}, \text{AoNE.sk}_{\text{pk}}) \leftarrow \text{AoNE.KeyGen}()$,
 - and DSum keys $(\text{DSum.pk}, \text{DSum.sk}_{\text{pk}}) \leftarrow \text{DSum.KeyGen}()$.
 Set the public key $\text{pk} = (\text{IP-FE.pk}, \text{AoNE.pk}, \text{DSum.pk})$ and the secret key $\text{sk}_{\text{pk}} = (K, \text{IP-FE.sk}_{\text{pk}}, \text{AoNE.sk}, \text{DSum.sk})$. Return the key pair $(\text{pk}, \text{sk}_{\text{pk}})$.
- **Encrypt**(sk_{pk}, m): Parse m as $(\mathbf{x}, \mathcal{U}_M, \ell)$, where $\mathbf{x} \in \mathbb{Z}_p^d$, $\mathcal{U}_M \in \mathcal{S}(\mathcal{PK})$, $\ell \in \{0, 1\}^*$, $\mathbf{s}_{\text{pk}, \mathcal{U}_M} = \mathcal{F}_K(\mathcal{U}_M) \in \mathbb{Z}_p^d$, $[h_\ell] = \mathcal{H}(\ell) \in \mathbb{G}$, and

$$c_{\text{pk}} \leftarrow \text{IP-FE.Encrypt}(\text{IP-FE.sk}_{\text{pk}}, [\mathbf{x}] + \mathbf{s}_{\text{pk}, \mathcal{U}_M} \cdot [h_\ell]).$$

Return

$$ct_{\text{pk}} = (\text{AoNE.Encrypt}(\text{AoNE.sk}_{\text{pk}}, (c_{\text{pk}}, (\text{AoNE.pk}')_{\text{pk}' \in \mathcal{U}_M}, \text{"ct"} \parallel \ell)), \mathcal{U}_M, \ell);$$

- **DKeyGen**(sk_{pk}, k): Parse k as $(\mathbf{y}_{\text{pk}'}, \text{pk}')_{\text{pk}' \in \mathcal{U}_K}$. Compute $\mathbf{s}_{\text{pk}, \mathcal{U}_K} = \mathcal{F}_K(\mathcal{U}_K) \in \mathbb{Z}_p^d$ and

$$d_{\text{pk}, k} = \text{DSum.Encrypt}(\text{DSum.sk}_{\text{pk}}, (\mathbf{y}_{\text{pk}'}^T \mathbf{s}_{\text{pk}, \mathcal{U}_K}, (\text{DSum.pk}')_{\text{pk}' \in \mathcal{U}_K}, k)).$$

Compute $d''_{\text{pk}, k} = \text{IP-FE.DKeyGen}(\text{IP-FE.sk}_{\text{pk}}, \mathbf{y}_{\text{pk}'})$ and

$$d'_{\text{pk}, k} \leftarrow \text{AoNE.Encrypt}(\text{AoNE.sk}_{\text{pk}}, (d''_{\text{pk}, k}, (\text{AoNE.pk}')_{\text{pk}' \in \mathcal{U}_K}, \text{"key"} \parallel k))$$

and return $\text{dk}_{\text{pk}, k} = (d_{\text{pk}, k}, d'_{\text{pk}, k})$;

³ Of course, our DSum and our IP-DDFE themselves use AoNE, which may rely on pairings if instantiated with our construction from Section 5.

- **Decrypt**(($\mathbf{dk}_{\mathbf{pk}',k_{\mathbf{pk}'}}$) $_{\mathbf{pk}' \in \mathcal{U}_K}$, ($\mathbf{ct}_{\mathbf{pk}}$) $_{\mathbf{pk} \in \mathcal{U}_M}$): If $\mathcal{U}_M \neq \mathcal{U}_K$ return \perp . Now let $\mathcal{U} = \mathcal{U}_M = \mathcal{U}_K$. Let $k \in \mathcal{K}$ be such that $k = k_{\mathbf{pk}}$ for all $\mathbf{pk} \in \mathcal{U}$. If there is no such k return \perp . Parse $\mathbf{dk}_{\mathbf{pk},k}$ as $(d_{\mathbf{pk},k}, d'_{\mathbf{pk},k})$ for all $\mathbf{pk} \in \mathcal{U}$.

Get

$$(c_{\mathbf{pk}})_{\mathbf{pk} \in \mathcal{U}} = \text{AoNE.Decrypt}(\epsilon, (\mathbf{ct}_{\mathbf{pk}})_{\mathbf{pk} \in \mathcal{U}})$$

and

$$(d''_{\mathbf{pk},k})_{\mathbf{pk} \in \mathcal{U}} = \text{AoNE.Decrypt}(\epsilon, (d'_{\mathbf{pk},k})_{\mathbf{pk} \in \mathcal{U}}).$$

Then compute $s_k = \sum_{\mathbf{pk} \in \mathcal{U}} \mathbf{y}_{\mathbf{pk}}^T \mathbf{s}_{\mathbf{pk},\mathcal{U}} = \text{DSum.Decrypt}(\epsilon, (d_{\mathbf{pk},k})_{\mathbf{pk} \in \mathcal{U}})$.

For all $\mathbf{pk} \in \mathcal{U}$, compute $z_{\mathbf{pk}} \in \mathbb{G}$ as

$$z_{\mathbf{pk}} \leftarrow \text{IP-FE.Decrypt}(d''_{\mathbf{pk},k}, c_{\mathbf{pk}}).$$

Let $\ell \in \{0, 1\}^*$ such that all $\mathbf{ct}_{\mathbf{pk}}$ for $\mathbf{pk} \in \mathcal{U}$ contain ℓ . If there is no such ℓ , return \perp . Otherwise, compute $[h_\ell] = \mathcal{H}(\ell) \in \mathbb{G}$ and return the discrete logarithm in base $[1]$ of

$$\left(\sum_{\mathbf{pk} \in \mathcal{U}} z_{\mathbf{pk}} \right) - s_k \cdot [h_\ell].$$

Correctness: We write $\mathbf{s}_{\mathbf{pk},\mathcal{U}} = \mathcal{F}_{K_{\mathbf{pk}}}(\mathcal{U})$. By correctness of AoNE, the use of $c_{\mathbf{pk}}$ in **Encrypt** and in **Decrypt** is consistent, as well as the use of $d''_{\mathbf{pk},k}$ in **DKeyGen** and **Decrypt**; By correctness of DSum, we have $s_k = \sum_{\mathbf{pk} \in \mathcal{U}} \mathbf{y}_{\mathbf{pk}}^T \mathbf{s}_{\mathbf{pk},\mathcal{U}}$; By correctness property of IP-FE, we have $z_{\mathbf{pk}} = [\mathbf{y}_{\mathbf{pk}}^T \mathbf{x}_{\mathbf{pk}} + \mathbf{y}_{\mathbf{pk}}^T \mathbf{s}_{\mathbf{pk},\mathcal{U}} h_\ell]$. Thus we eventually compute and return the discrete logarithm of

$$\begin{aligned} \left(\sum_{\mathbf{pk} \in \mathcal{U}} z_{\mathbf{pk}} \right) - s_k \cdot [h_\ell] &= \left(\sum_{\mathbf{pk} \in \mathcal{U}} [\mathbf{y}_{\mathbf{pk}}^T \mathbf{x}_{\mathbf{pk}} + \mathbf{y}_{\mathbf{pk}}^T \mathbf{s}_{\mathbf{pk},\mathcal{U}} h_\ell] \right) - \left(\sum_{\mathbf{pk} \in \mathcal{U}} \mathbf{y}_{\mathbf{pk}}^T \mathbf{s}_{\mathbf{pk},\mathcal{U}} \right) \cdot [h_\ell] \\ &= \left[\sum_{\mathbf{pk} \in \mathcal{U}} \mathbf{y}_{\mathbf{pk}}^T \mathbf{x}_{\mathbf{pk}} + \mathbf{y}_{\mathbf{pk}}^T \mathbf{s}_{\mathbf{pk},\mathcal{U}} h_\ell \right] - \left[\sum_{\mathbf{pk} \in \mathcal{U}} \mathbf{y}_{\mathbf{pk}}^T \mathbf{s}_{\mathbf{pk},\mathcal{U}} h_\ell \right] = \left[\sum_{\mathbf{pk} \in \mathcal{U}} \mathbf{y}_{\mathbf{pk}}^T \mathbf{x}_{\mathbf{pk}} \right] \end{aligned}$$

7.3 Security Proof

Theorem 26 (sel-sym-IND-Security of our IP-DDFE). *The Inner-Product DDFE scheme described in Section 7.2 is sel-sym-IND-secure (as per Definition 5) under the DDH assumption, assuming IP-FE is sel-IND secure, the AoNE scheme is sel-sym-IND-secure, the DSum scheme is sel-sym-IND-secure, and $(\mathcal{F}_K)_K$ is a secure PRF family.*

The proof can be found in Appendix B.4.

Acknowledgments.

This work was supported in part by the European Community's Seventh Framework Programme (FP7/2007-2013 Grant Agreement no. 339563 – CryptoCloud), the European Community's Horizon 2020 Project FENTEC (Grant Agreement no. 780108), the Google PhD fellowship, and the French FUI ANBLIC Project. This work was partially done while the third author was visiting ENS, Paris, and UC Berkeley, California.

References

- ABDP15. Michel Abdalla, Florian Bourse, Angelo De Caro, and David Pointcheval. Simple functional encryption schemes for inner products. In Jonathan Katz, editor, *PKC 2015*, volume 9020 of *LNCS*, pages 733–751. Springer, Heidelberg, March / April 2015.
- ABG19. Michel Abdalla, Fabrice Benhamouda, and Romain Gay. From single-input to multi-client inner-product functional encryption. In Steven D. Galbraith and Shihoh Moriai, editors, *ASIACRYPT 2019, Part III*, volume 11923 of *LNCS*, pages 552–582. Springer, Heidelberg, December 2019.
- ABKW19. Michel Abdalla, Fabrice Benhamouda, Markulf Kohlweiss, and Hendrik Waldner. Decentralizing inner-product functional encryption. In Dongdai Lin and Kazue Sako, editors, *PKC 2019, Part II*, volume 11443 of *LNCS*, pages 128–157. Springer, Heidelberg, April 2019.
- ACF⁺18. Michel Abdalla, Dario Catalano, Dario Fiore, Romain Gay, and Bogdan Ursu. Multi-input functional encryption for inner products: Function-hiding realizations and constructions without pairings. In Hovav Shacham and Alexandra Boldyreva, editors, *CRYPTO 2018, Part I*, volume 10991 of *LNCS*, pages 597–627. Springer, Heidelberg, August 2018.
- ACF⁺20. Shweta Agrawal, Michael Clear, Ophir Frieder, Sanjam Garg, Adam O’Neill, and Justin Thaler. Ad hoc multi-input functional encryption. In *11th Innovations in Theoretical Computer Science Conference (ITCS 2020)*. Schloss Dagstuhl-Leibniz-Zentrum für Informatik, 2020.
- AGRW17. Michel Abdalla, Romain Gay, Mariana Raykova, and Hoeteck Wee. Multi-input inner-product functional encryption from pairings. In Jean-Sébastien Coron and Jesper Buus Nielsen, editors, *EUROCRYPT 2017, Part I*, volume 10210 of *LNCS*, pages 601–626. Springer, Heidelberg, April / May 2017.
- ALS16. Shweta Agrawal, Benoît Libert, and Damien Stehlé. Fully secure functional encryption for inner products, from standard assumptions. In Matthew Robshaw and Jonathan Katz, editors, *CRYPTO 2016, Part III*, volume 9816 of *LNCS*, pages 333–362. Springer, Heidelberg, August 2016.
- BBG05. Dan Boneh, Xavier Boyen, and Eu-Jin Goh. Hierarchical identity based encryption with constant size ciphertext. In Ronald Cramer, editor, *EUROCRYPT 2005*, volume 3494 of *LNCS*, pages 440–456. Springer, Heidelberg, May 2005.
- BF01. Dan Boneh and Matthew K. Franklin. Identity-based encryption from the Weil pairing. In Joe Kilian, editor, *CRYPTO 2001*, volume 2139 of *LNCS*, pages 213–229. Springer, Heidelberg, August 2001.
- BJL16. Fabrice Benhamouda, Marc Joye, and Benoît Libert. A new framework for privacy-preserving aggregation of time-series data. *ACM Trans. Inf. Syst. Secur.*, 18(3):10:1–10:21, 2016.
- BS15. Zvika Brakerski and Gil Segev. Function-private functional encryption in the private-key setting. In Yevgeniy Dodis and Jesper Buus Nielsen, editors, *TCC 2015, Part II*, volume 9015 of *LNCS*, pages 306–324. Springer, Heidelberg, March 2015.
- BSW11. Dan Boneh, Amit Sahai, and Brent Waters. Functional encryption: Definitions and challenges. In Yuval Ishai, editor, *TCC 2011*, volume 6597 of *LNCS*, pages 253–273. Springer, Heidelberg, March 2011.
- CC09. Melissa Chase and Sherman S. M. Chow. Improving privacy and security in multi-authority attribute-based encryption. In Ehab Al-Shaer, Somesh Jha, and Angelos D. Keromytis, editors, *ACM CCS 2009*, pages 121–130. ACM Press, November 2009.
- CDG⁺18. Jérémy Chotard, Edouard Dufour Sans, Romain Gay, Duong Hieu Phan, and David Pointcheval. Decentralized multi-client functional encryption for inner product. In Thomas Peyrin and Steven Galbraith, editors, *ASIACRYPT 2018, Part II*, volume 11273 of *LNCS*, pages 703–732. Springer, Heidelberg, December 2018.
- CGJS15. Nishanth Chandran, Vipul Goyal, Aayush Jain, and Amit Sahai. Functional encryption: Decentralised and delegatable. Cryptology ePrint Archive, Report 2015/1017, 2015. <http://eprint.iacr.org/2015/1017>.
- CKS08. David Cash, Eike Kiltz, and Victor Shoup. The twin Diffie-Hellman problem and applications. In Nigel P. Smart, editor, *EUROCRYPT 2008*, volume 4965 of *LNCS*, pages 127–145. Springer, Heidelberg, April 2008.
- Coc01. Clifford Cocks. An identity based encryption scheme based on quadratic residues. In *IMA international conference on cryptography and coding*, pages 360–363. Springer, 2001.
- CSS12. T.-H. Hubert Chan, Elaine Shi, and Dawn Song. Privacy-preserving stream aggregation with fault tolerance. In Angelos D. Keromytis, editor, *FC 2012*, volume 7397 of *LNCS*, pages 200–214. Springer, Heidelberg, February / March 2012.
- DG17. Nico Döttling and Sanjam Garg. Identity-based encryption from the Diffie-Hellman assumption. In Jonathan Katz and Hovav Shacham, editors, *CRYPTO 2017, Part I*, volume 10401 of *LNCS*, pages 537–569. Springer, Heidelberg, August 2017.
- EHK⁺13. Alex Escala, Gottfried Herold, Eike Kiltz, Carla Ràfols, and Jorge Villar. An algebraic framework for Diffie-Hellman assumptions. In Ran Canetti and Juan A. Garay, editors, *CRYPTO 2013, Part II*, volume 8043 of *LNCS*, pages 129–147. Springer, Heidelberg, August 2013.
- FHKP13. Eduarda S. V. Freire, Dennis Hofheinz, Eike Kiltz, and Kenneth G. Paterson. Non-interactive key exchange. In Kaoru Kurosawa and Goichiro Hanaoka, editors, *PKC 2013*, volume 7778 of *LNCS*, pages 254–271. Springer, Heidelberg, February / March 2013.
- Gen09. Craig Gentry. Fully homomorphic encryption using ideal lattices. In Michael Mitzenmacher, editor, *41st ACM STOC*, pages 169–178. ACM Press, May / June 2009.
- GGG⁺14. Shafi Goldwasser, S. Dov Gordon, Vipul Goyal, Abhishek Jain, Jonathan Katz, Feng-Hao Liu, Amit Sahai, Elaine Shi, and Hong-Sheng Zhou. Multi-input functional encryption. In Phong Q. Nguyen and Elisabeth Oswald, editors, *EUROCRYPT 2014*, volume 8441 of *LNCS*, pages 578–602. Springer, Heidelberg, May 2014.

- GKL⁺13. S. Dov Gordon, Jonathan Katz, Feng-Hao Liu, Elaine Shi, and Hong-Sheng Zhou. Multi-input functional encryption. Cryptology ePrint Archive, Report 2013/774, 2013. <http://eprint.iacr.org/2013/774>.
- GPSW06. Vipul Goyal, Omkant Pandey, Amit Sahai, and Brent Waters. Attribute-based encryption for fine-grained access control of encrypted data. In Ari Juels, Rebecca N. Wright, and Sabrina De Capitani di Vimercati, editors, *ACM CCS 2006*, pages 89–98. ACM Press, October / November 2006. Available as Cryptology ePrint Archive Report 2006/309.
- GPV08. Craig Gentry, Chris Peikert, and Vinod Vaikuntanathan. Trapdoors for hard lattices and new cryptographic constructions. In Richard E. Ladner and Cynthia Dwork, editors, *40th ACM STOC*, pages 197–206. ACM Press, May 2008.
- JL13. Marc Joye and Benoît Libert. A scalable scheme for privacy-preserving aggregation of time-series data. In Ahmad-Reza Sadeghi, editor, *FC 2013*, volume 7859 of *LNCS*, pages 111–125. Springer, Heidelberg, April 2013.
- LT19. Benoît Libert and Radu Titiu. Multi-client functional encryption for linear functions in the standard model from LWE. In Steven D. Galbraith and Shihō Moriai, editors, *ASIACRYPT 2019, Part III*, volume 11923 of *LNCS*, pages 520–551. Springer, Heidelberg, December 2019.
- LTV12. Adriana López-Alt, Eran Tromer, and Vinod Vaikuntanathan. On-the-fly multiparty computation on the cloud via multikey fully homomorphic encryption. In Howard J. Karloff and Toniann Pitassi, editors, *44th ACM STOC*, pages 1219–1234. ACM Press, May 2012.
- LW11. Allison B. Lewko and Brent Waters. Decentralizing attribute-based encryption. In Kenneth G. Paterson, editor, *EUROCRYPT 2011*, volume 6632 of *LNCS*, pages 568–588. Springer, Heidelberg, May 2011.
- SCR⁺11. Elaine Shi, T.-H. Hubert Chan, Eleanor G. Rieffel, Richard Chow, and Dawn Song. Privacy-preserving aggregation of time-series data. In *NDSS 2011*. The Internet Society, February 2011.
- Sha84. Adi Shamir. Identity-based cryptosystems and signature schemes. In G. R. Blakley and David Chaum, editors, *CRYPTO'84*, volume 196 of *LNCS*, pages 47–53. Springer, Heidelberg, August 1984.
- SW05. Amit Sahai and Brent R. Waters. Fuzzy identity-based encryption. In Ronald Cramer, editor, *EUROCRYPT 2005*, volume 3494 of *LNCS*, pages 457–473. Springer, Heidelberg, May 2005.

A Single-Input FE for Inner Products

Here we recall the IPFE from [ALS16] on a cyclic group \mathbb{G} . Its IND security is proven in [ALS16], under the DDH assumption in \mathbb{G} .

- IP-FE.KeyGen($\mathbb{G}, d \in \mathbb{N}$): $\mathbf{a} \xleftarrow{\$} \text{DDH}$, $\mathbf{U} \xleftarrow{\$} \mathbb{Z}_p^{d \times 2}$, $\text{pk} = ([\mathbf{a}], [\mathbf{U}\mathbf{a}])$, $\text{msk} = \mathbf{U}$. Return (pk, msk) .
- IP-FE.Enc($\text{pk}, \mathbf{x} \in \mathbb{Z}_p^d$): $r \xleftarrow{\$} \mathbb{Z}_p$, return $\begin{bmatrix} \mathbf{a}r \\ \mathbf{x} + \mathbf{U}\mathbf{a}r \end{bmatrix} \in \mathbb{G}^{d+2}$.
- IP-FE.DKeyGen($\text{msk}, \mathbf{y} \in \mathbb{Z}_p^d$): return $\begin{pmatrix} -\mathbf{U}^\top \mathbf{y} \\ \mathbf{y} \end{pmatrix} \in \mathbb{Z}_p^{d+2}$.
- IP-FE.Dec($\text{pk}, [\mathbf{c}], \mathbf{k}$): return $[\mathbf{c}]^\top \mathbf{k} \in \mathbb{G}$.

B Security Proofs

B.1 Theorem 22 (IND-Security of AoNE)

The All-or-Nothing Encapsulation scheme described in Section 4.2 is IND-secure (as per Definition 3) assuming the IBE scheme is IND-secure (as per Definition 17).

Proof. Let q_p, q_c denote (polynomial) upper bounds on the number of adversary queries to the QNewHonest oracle, and the number of *unique* pairs (\mathcal{U}_M, ℓ) for which the adversary sends at least one QEncrypt or QLeftRight query, respectively. We define the following games for $i \in \{0, \dots, q_c\}$:

Game \mathbf{G}_i : The challenger does as specified in Definition 3, except for queries to QLeftRight. Queries to QLeftRight take as arguments a public key pk and two messages $m_0 = (x_0, \mathcal{U}_{M,0}, \ell_0)$ and $m_1 = (x_1, \mathcal{U}_{M,1}, \ell_1)$. Note that by functionality and by description of the scheme, the response reveals $\mathcal{U}_{M,b}, \ell_b$, so if the adversary wants to avoid the Finalize condition ignoring its guess, it must have $\ell_0 = \ell_1 = \ell$ and $\mathcal{U}_{M,0} = \mathcal{U}_{M,1} = \mathcal{U}_M$. Now let $(\mathcal{U}_{M,j}, \ell_j)$ be the j 'th such pair queried to QEncrypt or QLeftRight. In \mathbf{G}_i , the challenger will respond to QLeftRight queries by encrypting m_0 if $i < j$ and m_b otherwise, where $b \xleftarrow{\$} \{0, 1\}$ is the random bit chosen by the challenger.

Note that in \mathbf{G}_0 , all challenge ciphertexts contain the left message, while in \mathbf{G}_{q_c} all challenge ciphertexts contain the right message. Thus we only need to show that $\mathbf{G}_{i-1} \sim_c \mathbf{G}_i$ for all $i \in [q_c]$.

$\mathbf{G}_{i-1} \sim_c \mathbf{G}_i$: We proceed by contradiction, and, from a PPT adversary \mathcal{A} which can distinguish between \mathbf{G}_{i-1} and \mathbf{G}_i with noticeable advantage, we construct a PPT algorithm \mathcal{B} which breaks the IND-security of IBE with noticeable advantage.

\mathcal{B} starts playing the IBE IND-security game and gets a public key IBE.pk . We need to choose the participant whose AoNE public key will be $\text{AoNE.pk} = \text{IBE.pk}$ carefully, because we won't be able to answer QDKeyGen requests for them. The key is to notice that if the adversary is going to distinguish between \mathbf{G}_{i-1} and \mathbf{G}_i , the two need to be different, meaning the adversary \mathcal{A} needs to make at least one query to QLeftRight on $(\mathcal{U}_{M,i}, \ell_i)$ with $x'_0 \neq x'_1$ with noticeable probability, and that, conditioned on that event, \mathcal{A} retains noticeable advantage. We can thus safely assume that \mathcal{A} will make such a query, and abort otherwise. From then on, if it were the case that for every $\text{pk} \in \mathcal{U}_M$, the adversary either makes a QEncrypt or QLeftRight query on (\mathcal{U}_M, ℓ) or pk is eventually not honest⁴, then the condition in the Finalize part of the security game (see Definition 3) would notice that $x'_0 \neq x'_1$ and set the adversary's guess at random, rendering the adversary's efforts fruitless. We can thus safely assume that there is a $\text{pk}^* \in \mathcal{U}_M$ such that pk^* will be created via QNewHonest⁵, and the adversary will not query QEncrypt or QLeftRight on (\mathcal{U}_M, ℓ) for pk^* or query QCorrupt on pk^* .

We proceed by guessing which query to QNewHonest will eventually be pk^* . We cannot simply guess a member of \mathcal{U}_M because we do not know anything about \mathcal{U}_M during the initialization phase of the game, and by the time the i th (\mathcal{U}_M, ℓ) pair is queried, it is possible that many queries have been made to the QNewHonest oracle, and at that point it would be too late to embed the IBE public key in the adversary's view. Instead, we guess the index $t^* \in [q_p]$ of the query to QNewHonest which eventually yields a public key pk_{t^*} which we hope matches pk^* . At that index, we respond with $\text{pk}_{t^*} = \text{IBE.pk}$. Because the adversary will only make polynomially many queries to QNewHonest, our advantage is only polynomially degraded by this guess and the reduction remains valid.

Having done this, we can naturally answer most queries involving pk_{t^*} by using the oracles of the IND security game of IBE and the fact that our IBE is public key. That is, we answer all QEncrypt and most (see below) QLeftRight queries by running IBE.Encrypt ourselves and making IBE.QDKeyGen queries.

The exceptions are QLeftRight queries to any $\text{pk} \in \mathcal{U}_M$ for (\mathcal{U}_M, ℓ) . Let $n = |\mathcal{U}_M|$ and $\zeta \in [n]$ be such that pk^* is the ζ th public key in \mathcal{U}_M for the universally agreed upon order. In responding to $\text{QLeftRight}(\text{pk}, (x_0, \mathcal{U}_M, \ell), (x_1, \mathcal{U}_M, \ell))$, we will compute two sequences of α 's as follows: for $s \in \{0, 1\}$, $k \in [\zeta - 1]$, let $\alpha_{\text{pk},0}^s = x_s$ and $\alpha_{\text{pk},k} = \text{IBE.Encrypt}(\text{pk}_k, (\alpha_{\text{pk},k-1}, \mathcal{U}_M || \ell))$. Now compute $\alpha_{\text{pk},\zeta} = \text{IBE.QLeftRight}((\alpha_{\text{pk},\zeta-1}^0, \mathcal{U}_M || \ell), (\alpha_{\text{pk},\zeta-1}^1, \mathcal{U}_M || \ell))$, and compute the rest of the α 's and the resulting ciphertext as per AoNE.Encrypt .

When $\text{IBE.b} = 0$, the adversary \mathcal{A} is playing \mathbf{G}_{i-1} . When $\text{IBE.b} = 1$, the adversary \mathcal{A} is playing \mathbf{G}_i . We only need to check that we do not violate the Finalize condition of the IBE IND-security game. But this is clear because the only IBE.QLeftRight query we make is for (\mathcal{U}_M, ℓ) , and for that pair we never get a AoNE.QLeftRight or AoNE.QEncrypt query so we never make an IBE.QDKeyGen query. This concludes our proof. \square

B.2 Theorem 24 (IND-Security of AoNE)

The All-or-Nothing Encapsulation scheme described in Section 5.2 is IND-secure (as per Definition 3) under the DBDH assumption, in the random oracle model.

⁴ Note that here there are two ways for pk to be dishonest: either the adversary has the challenger create pk via QNewHonest and later corrupts it via QCorrupt, or the adversary generates pk on its own.

⁵ Note that here, and in subsequent proofs, we implicitly ignore the very real possibility that the adversary sends a query for a set \mathcal{U}_M for which a *later* query to QNewHonest generates a $\text{pk} \in \mathcal{U}_M$. Because this happens with negligible probability it is safe to abort when this situation materializes.

Proof. Let q_r , q_p , and q_c denote (polynomial) upper bounds on the number of *unique* queries sent by the adversary to the Random Oracle (both directly and indirectly through queries to QEncrypt or QLeftRight), the QNewHonest oracle, and the challenge oracle QLeftRight, respectively. We define the following games for $i \in [q_r]$:

Game $\mathbf{G}_{i,0}$: The challenger plays as does the challenger in Definition 3, except for queries to QLeftRight. Queries to QLeftRight take as arguments a public key pk and two messages $m_0 = (x_0, \mathcal{U}_{M,0}, \ell_0)$ and $m_1 = (x_1, \mathcal{U}_{M,1}, \ell_1)$. Note that by functionality and by description of the scheme, the response reveals $\mathcal{U}_{M,b}, \ell_b$, so if the adversary wants to avoid the Finalize condition ignoring its guess, it must have $\ell_0 = \ell_1 = \ell$ and $\mathcal{U}_{M,0} = \mathcal{U}_{M,1} = \mathcal{U}_M$. Now let j be such that (\mathcal{U}_M, ℓ) is the j th unique query to the random oracle. In $\mathbf{G}_{i,0}$, the challenger will respond to QLeftRight queries by encrypting m_0 if $i \leq j$ and m_1 otherwise.

Game $\mathbf{G}_{i,1}$: This game is similar to $\mathbf{G}_{i,0}$, except for the fact that, using (\mathcal{U}_M, ℓ) to denote the i th unique query to the Random Oracle, in all QLeftRight queries for (\mathcal{U}_M, ℓ) such that $x_0 \neq x_1$, the challenger uses an ephemeral random value K to compute $c_{\text{pk}} = \text{SEnc}(K, x_0)$ instead of using $K_{\text{pk}, \mathcal{U}_M, \ell}$ as described in Encrypt.

Game $\mathbf{G}_{i,2}$: This game is similar to $\mathbf{G}_{i,1}$, except for the fact that, using (\mathcal{U}_M, ℓ) to denote the i th unique query to the Random Oracle, in all QLeftRight queries for (\mathcal{U}_M, ℓ) , the challenger will encrypt the right message: $c_{\text{pk}} = \text{SEnc}(K, x_1)$ when $x_0 \neq x_1$ and $c_{\text{pk}} = \text{SEnc}(K_{\text{pk}, \mathcal{U}_M, \ell}, x_1)$ otherwise.

Game $\mathbf{G}_{i,3}$: This game is similar to $\mathbf{G}_{i,2}$, except for the fact that, using (\mathcal{U}_M, ℓ) to denote the i th unique query to the Random Oracle, in all QLeftRight queries for (\mathcal{U}_M, ℓ) , the challenger goes back to using an honestly computed $K_{\text{pk}, \mathcal{U}_M, \ell}$.

Note that in $\mathbf{G}_{1,0}$, all challenge ciphertexts contain the left message, while in $\mathbf{G}_{q_r,3}$ all challenge ciphertexts contain the right message, and $\mathbf{G}_{i,3} = \mathbf{G}_{i+1,0}$, for $i \in [q_r - 1]$. Thus, we need only prove that $\mathbf{G}_{1,0}$ is computationally indistinguishable from $\mathbf{G}_{q_r,3}$ which we do by showing that $\mathbf{G}_{i,0}$ is computationally indistinguishable from $\mathbf{G}_{i,3}$ for all $i \in [q_r]$, and then using the Hybrid Lemma. We want to show that $\mathbf{G}_{i,0} \sim_c \mathbf{G}_{i,1} \sim_c \mathbf{G}_{i,2} \sim_c \mathbf{G}_{i,3} = \mathbf{G}_{i+1,0}$.

First, $\mathbf{G}_{i,1} \sim_c \mathbf{G}_{i,2}$ follows immediately from the One-Time Security of our symmetric encryption scheme. The challenge is in proving $\mathbf{G}_{i,0} \sim_c \mathbf{G}_{i,1}$ and $\mathbf{G}_{i,2} \sim_c \mathbf{G}_{i,3}$. The problems are basically the same: they consist in switching between an honestly computed key $K_{\text{pk}, \mathcal{U}_M, \ell}$ and a random key K in one direction or the other, so without loss of generality we focus on proving $\mathbf{G}_{i,0} \sim_c \mathbf{G}_{i,1}$.

$\mathbf{G}_{i,0} \sim_c \mathbf{G}_{i,1}$: We proceed by contradiction, and, from a PPT adversary \mathcal{A} which can distinguish between $\mathbf{G}_{i,0}$ and $\mathbf{G}_{i,1}$ with noticeable advantage, we construct a PPT algorithm \mathcal{B} which breaks the q_r -fold DBDH assumption with noticeable advantage. By Lemma 13 this is enough to show that our scheme is secure under DBDH. The core idea of the reduction is to embed the q_r -fold DBDH tuple $([a]_1, [b]_1, [b]_2, \{[c_i]_2, [v_i]_T\}_{i \in [q_r]})$ in the view of the adversary as follows: $[a]_1$ will serve as the random oracle response to a query for $\mathcal{U}_M || \ell$, $[b]_2$ will be the public key of a participant in \mathcal{U}_M , and the c_i 's will be the randomnesses used in the encryption (denoted by r_{pk} in our construction).

We need to choose the participant whose public key will be $[b]_2$ carefully, because its $S_{\text{pk}, \mathcal{U}_M, \ell}$ on a query to QEncrypt or QLeftRight with (\mathcal{U}_M, ℓ) will be $[ab]_1$, which we are unable to compute from the Q -fold DBDH tuple (otherwise we could pair it with $[c_i]_2$ and trivially break the assumption). The key is to notice that if the adversary is going to distinguish between $\mathbf{G}_{i,0}$ and $\mathbf{G}_{i,1}$, the two need to be different, meaning the adversary \mathcal{A} needs to make at least one query to QLeftRight on (\mathcal{U}_M, ℓ) with $x'_0 \neq x'_1$ with noticeable probability, and that, conditioned on that event, \mathcal{A} retains noticeable advantage. We can thus safely assume that \mathcal{A} will make such a query, and abort otherwise. From then on, if it were the case that for every $\text{pk} \in \mathcal{U}_M$, the adversary either makes a QEncrypt or QLeftRight query on (\mathcal{U}_M, ℓ) or pk is eventually not honest⁶, then the condition in the Finalize part of the security game

⁶ Note that here there are two ways for pk to be dishonest: either the adversary has the challenger create pk via QNewHonest and later corrupts it via QCorrupt, or the adversary generates pk on its own.

(see Definition 3) would notice that $x'_0 \neq x'_1$ and set the adversary's guess at random, rendering the adversary's efforts fruitless. We can thus safely assume that there is a $\text{pk}^* \in \mathcal{U}_M$ such that pk^* will be created via `QNewHonest`, and the adversary will not query `QEncrypt` or `QLeftRight` on (\mathcal{U}_M, ℓ) for pk^* or query `QCorrupt` on pk^* .

We proceed by guessing which query to `QNewHonest` will eventually be pk^* . We cannot simply guess a member of \mathcal{U}_M because we do not know anything about \mathcal{U}_M during the initialization phase of the game, and by the time the i th query to the Random Oracle is made, it is possible that many queries have been made to the `QNewHonest` oracle, and at that point it would be too late to embed the DBDH challenge in the adversary's view. Instead, we guess the index $j^* \in [q_p]$ of the query to `QNewHonest` which eventually yields a public key pk_{j^*} which we hope matches pk^* . At that index, we respond with $\text{pk}_{j^*} = [b]_2$. Because the adversary will only make polynomially many queries to `QNewHonest`, our advantage is only polynomially degraded by this guess and the reduction remains valid.

Having done this, we can no longer "naturally" respond to `QEncrypt` or `QLeftRight` queries for pk_{j^*} even on pairs of sets of participants and labels distinct from (\mathcal{U}_M, ℓ) because we do not know the discrete logarithm of pk_{j^*} (its implicit secret key), which we would need to compute $S_{\text{pk}, \mathcal{U}_M, \ell}$ as described in `Encrypt`. Instead, we sample responses to Random Oracle queries by sampling an element of \mathbb{Z}_p and raising $[1]_1$ to that randomness, and store that randomness so we can both give consistent answers to further queries and compute $S_{\text{pk}, \mathcal{U}_M, \ell}$ by raising $[b]_1$ to the appropriate power. The only exception to that rule is that the i th unique random oracle request will be answered with $[a]_1$, and subsequent requests for the same (\mathcal{U}_M, ℓ) pair will be answered consistently.

During initialization, we set a counter *count* to 0 and, when encrypting for queries to `QLeftRight` on (\mathcal{U}_M, ℓ) such that $x_0 \neq x_1$, we increment it and replace $[r_{\text{pk}}]_2$ by $[c_{\text{count}}]_2$ from our q_r -fold DBDH tuple and $e(\mathcal{H}(\mathcal{U}_M || \ell), r_{\text{pk}} \cdot \text{pk}_{j^*})$ in the computation of $K_{\text{pk}, \mathcal{U}_M, \ell}$ by $[v_{\text{count}}]_T$ from our q_r -fold DBDH tuple. We respond to `QEncrypt` queries (or `QLeftRight` queries with $x_0 = x_1$) by encrypting honestly with a uniformly sampled $r_{\text{pk}} \in \mathbb{Z}_p$. The target group element $e([a]_1, [b]_2)$ can be raised to the r_{pk} to generate honest ciphertexts; We respond to `QCorrupt` queries for public keys other than pk_{j^*} naturally, since we know the discrete logarithm of the public key. If anything goes wrong with our guesses and we get a query we can't respond to, we guess at random for the q_r -fold DBDH game and abort.

Now notice that when we get an "honest" q_r -fold DBDH tuple, that is $v_i = abc_i$ for all $i \in [q_r]$, the adversary is playing $\mathbf{G}_{i,0}$. On the other hand, when we are getting a fake tuple and we instead have $v_i = s_i$, then the randomness in s_i masks any information in the computation of $K_{\text{pk}, \mathcal{U}_M, \ell}$, causing it to be uniformly random, as it would be in $\mathbf{G}_{i,1}$. Thus, when \mathcal{A} has noticeable advantage against our AoNE, \mathcal{B} will have noticeable advantage against q_r -fold DBDH, which violates the DBDH assumption. This concludes our proof. \square

B.3 Theorem 25 (IND-Security of DSum)

The Decentralized Sum scheme described in Section 6.2 is IND-secure (as per Definition 3) so long as NIKE is IND-secure (as per Definition 15) and $(\mathcal{F}_K)_K$ is a secure PRF family.

Proof. We use a hybrid argument that goes over all the pairs (\mathcal{U}_M, ℓ) that are contained in `QLeftRight`-queries. Writing Q the number of such pairs, and ordering these pairs as they are revealed to the experiment, we define for all $i \in \{0, \dots, Q\}$ the game \mathbf{G}_i as the security game given in Definition 3, except the `QLeftRight`-queries containing one of the first i 'th pairs are always answered with the left challenge message m^0 , as opposed to m^b for the chosen random bit $b \stackrel{\$}{\leftarrow} \{0, 1\}$. Let \mathcal{A} be a PPT adversary. For any game \mathbf{G} , we write $\text{Adv}_{\mathbf{G}}(\mathcal{A})$ the advantage of \mathcal{A} in the game \mathbf{G} . Note that \mathbf{G}_0 is the security game defined in Definition 3, whereas $\text{Adv}_{\mathbf{G}_Q}(\mathcal{A}) = 0$, since the adversary's view in \mathbf{G}_Q does not depend on the random bit $b \stackrel{\$}{\leftarrow} \{0, 1\}$. Thus, it suffices to show that for all $i \in [Q]$, $\mathbf{G}_{i-1} \sim_c \mathbf{G}_i$. We denote by $(\mathcal{U}_M^*, \ell^*)$ the i 'th pair.

We stress that in Definition 3, in case of a `QLeftRight`-query, we assume the two messages to be distinct, otherwise this is actually equivalent to a `QEncrypt`-query. This remark will be crucial in the rest of the proof.

We distinguish two cases: Case 1) there exists $\text{pk} \in \mathcal{U}_M^*$ that is output of **QNewHonest**, never queried to **QCorrupt**, **QLeftRight** and **QEncrypt** are never queried on a query containing $\text{pk}, \mathcal{U}_M^*, \ell^*$. Intuitively, that means the adversary does not get a complete ciphertext for \mathcal{U}_M^*, ℓ^* . In that case, we can use the security of **AoNE**; Case 2) all $\text{pk} \in \mathcal{U}_M^*$ that are output of **QNewHonest** are either queried to **QCorrupt**, or part of query containing $(\mathcal{U}_M^*, \ell^*)$ to **QLeftRight** or **QEncrypt**. Intuitively, that means the adversary gets a complete ciphertext for \mathcal{U}_M^*, ℓ^* , as it may have generated on its own all the ciphertexts under dishonest keys.

As case 1) is easily dealt with by leveraging the security of **AoNE**, we now focus on case 2): For all $\text{pk} \in \mathcal{U}_M^*$, in case of multiple **QLeftRight**-queries for the same tuple $(\text{pk}, \mathcal{U}_M^*, \ell^*)$, there is a unique value $\Delta_{\text{pk}} \in \mathbb{A}$ such that all **QLeftRight**-queries of the form $(\text{pk}, (x_{\text{pk}}^0, \mathcal{U}_M^*, \ell^*), (x_{\text{pk}}^1, \mathcal{U}_M^*, \ell^*))$, they are such that $x_{\text{pk}}^0 - x_{\text{pk}}^b = \Delta_{\text{pk}}$, unless the **Finalize** procedure outputs a random bit, independent of \mathcal{A} 's guess. This is because when the ciphertext is complete, the adversary can legitimately decrypt and recover both the sum involving x_{pk}^b and x_{pk}^b , subtract the two sums, and obtain $x_{\text{pk}}^b - x_{\text{pk}}^b$. Thus, this value must be independent of the random bit b unless the adversary trivially wins the game, i.e. the condition (*) from Definition 1 holds, in which case **Finalize** outputs a random bit.

Moreover, when a query $(\text{pk}, (x_{\text{pk}}^0, \mathcal{U}_M^*, \ell^*), (x_{\text{pk}}^1, \mathcal{U}_M^*, \ell^*))$ is sent to **QLeftRight**, then pk cannot be queried to **QCorrupt**, unless the **Finalize** procedure outputs a random bit. Indeed, recall \mathcal{A} can legitimately recover a sum involving x_{pk}^b , since we are in case 2). If **Adv** corrupts pk , then it can compute a partial ciphertext for $x_{\text{pk}}^0 = 0_{\mathbb{A}}$, the neutral element of \mathbb{A} on $(\mathcal{U}_M^*, \ell^*)$, use this ciphertext with the other challenge ciphertexts and recover a sum involving $x_{\text{pk}}^b = 0$. Subtracting the two sums, \mathcal{A} recovers the value x_{pk}^b , which must not depend on b , unless \mathcal{A} trivially wins the game, that is, unless **Finalize** outputs a random bit. This means $x_{\text{pk}}^0 = x_{\text{pk}}^1$, which is a contradiction, since by definition of the security game, **QLeftRight** is only queried for different messages — **QEncrypt** is used otherwise.

Simply put, in case 2), when pk is queried as part of a query containing \mathcal{U}_M^*, ℓ^* , the challenger knows pk cannot be queried to **QCorrupt**. We call such pk explicitly honest. We will make crucial use of that observation in the rest of the proof. If there is no explicitly honest client, that means no queries to **QLeftRight** containing $(\mathcal{U}_M^*, \ell^*)$, then \mathbf{G}_{i-1} and \mathbf{G}_i are clearly the same. Thus, we focus on the case where there are explicitly honest clients.

Note that there must be at least two such explicitly honest clients, unless **Finalize** outputs a random bit: Suppose $(\text{pk}, (x_{\text{pk}}^0, \mathcal{U}_M^*, \ell^*), (x_{\text{pk}}^1, \mathcal{U}_M^*, \ell^*))$ is the only explicitly honest query to **QLeftRight**, then the sum \mathcal{A} can legitimately recover depends on the random bit $b \stackrel{\$}{\leftarrow} \{0, 1\}$ chosen by the experiment only in the value x_{pk}^b used in the sum. Therefore, it must be the case that $x_{\text{pk}}^0 = x_{\text{pk}}^1$, which is again a contradiction.

The rest of the proof follows a similar strategy than the adaptive security proof of [ABG19, Theorem 3.7] in the context of multi-client inner-product FE. Namely, the challenger cannot simply guess the set of explicitly honest clients, since that would incur an exponential security loss. Instead, it guesses the number κ of explicitly honest clients, and gradually introduces a κ -out-of- κ secret sharing of 0 in the output of **QLeftRight**-queries on explicitly honest clients. This is done step by step in the proof, so that the reduction has to guess only a pair of explicitly honest clients at each step, which only incurs a polynomial security loss. To prove that $\mathbf{G}_{i-1} \sim_c \mathbf{G}_i$ in case 2), we introduce the following hybrid games.

Game \mathbf{G}_{i-1}^* : this game is as \mathbf{G}_{i-1} , except the challenger guesses the number of explicitly honest clients. Writing q_{pk} the number of queries made to **QNewHonest**, the challenger samples $\kappa \stackrel{\$}{\leftarrow} [2, q_{\text{pk}}]$. If eventually the challenger realizes the guess was incorrect, it aborts the experiments and outputs a random bit. Otherwise, nothing changes. It is clear that $\text{Adv}_{\mathbf{G}_{i-1}^*}(\mathcal{A}) = \frac{\text{Adv}_{\mathbf{G}_{i-1}}(\mathcal{A})}{q_{\text{pk}}-1}$. For all $t \in [0, \kappa]$, we define the following games.

Game $\mathbf{G}_{i-1,t}$: this game is as \mathbf{G}_{i-1}^* , except that for the first explicitly honest client pk_1 , **QLeftRight** $(\text{pk}_1$,

$(x_{\mathbf{pk}_1}^0, (\mathcal{U}_M^*, \ell^*)), (x_{\mathbf{pk}_1}^1, (\mathcal{U}_M^*, \ell^*))$ uses

$$c_{\mathbf{pk}_1} = x_{\mathbf{pk}_1}^b + \sum_{\mathbf{pk}' < \mathbf{pk}_1, \mathbf{pk}' \in \mathcal{U}_M^*} r_{\mathbf{pk}_1, \mathbf{pk}', \mathcal{U}_M^*, \ell^*} - \sum_{\mathbf{pk}' > \mathbf{pk}_1, \mathbf{pk}' \in \mathcal{U}_M^*} r_{\mathbf{pk}_1, \mathbf{pk}', \mathcal{U}_M^*, \ell^*} - \sum_{j \in [2, t]} u_j,$$

where $u_j \xleftarrow{\$} \mathbb{A}$ for all $j \in [t]$. Similarly, for the ρ 'th honest client \mathbf{pk}_ρ with $1 < \rho \leq t$, $\text{QLeftRight}(\mathbf{pk}_\rho, (x_{\mathbf{pk}_\rho}^0, (\mathcal{U}_M^*, \ell^*)), (x_{\mathbf{pk}_\rho}^1, (\mathcal{U}_M^*, \ell^*)))$ uses

$$c_{\mathbf{pk}_\rho} = x_{\mathbf{pk}_\rho}^b + \sum_{\mathbf{pk}' < \mathbf{pk}_\rho, \mathbf{pk}' \in \mathcal{U}_M^*} r_{\mathbf{pk}_\rho, \mathbf{pk}', \mathcal{U}_M^*, \ell^*} - \sum_{\mathbf{pk}' > \mathbf{pk}_\rho, \mathbf{pk}' \in \mathcal{U}_M^*} r_{\mathbf{pk}_\rho, \mathbf{pk}', \mathcal{U}_M^*, \ell^*} + u_\rho.$$

The changes from \mathbf{G}_{i-1}^* are highlighted in gray. It is clear that $\mathbf{G}_{i-1,0}$ is the same as \mathbf{G}_{i-1}^* . To transition from $\mathbf{G}_{i-1,0}$ to $\mathbf{G}_{i-1,2}$, the challenger first guesses the first and the second explicitly honest clients, denoted \mathbf{pk}_1 and \mathbf{pk}_2 , respectively. If the guess is incorrect, the challenger aborts and return a random bit. This incurs a security loss of $\frac{q_{\mathbf{pk}}(q_{\mathbf{pk}}-1)}{2}$. Then the challenger uses the security of the NIKE to switch the $K_{\mathbf{pk}_1, \mathbf{pk}_2}$ to uniformly random. Then it uses the security of the PRF to switch the values $r_{\mathbf{pk}_1, \mathbf{pk}_2, \mathcal{U}_M^*, \ell^*}$ to uniformly random over \mathbb{A} . Since this value appears positively and negatively in the values $c_{\mathbf{pk}_1}$ and $c_{\mathbf{pk}_2}$ used by QLeftRight , we can add the offset $-u_2$ to $c_{\mathbf{pk}_1}$ and u_2 to $c_{\mathbf{pk}_2}$, as in game $\mathbf{G}_{i-1,2}$. Then, we switch back the value $r_{\mathbf{pk}_1, \mathbf{pk}_2, \mathcal{U}_M^*, \ell^*}$ from truly random to pseudo-random, and the key $K_{\mathbf{pk}_1, \mathbf{pk}_2}$ to pseudo random, using the security of the PRF and the NIKE, respectively. Note that these security notions can only be used when the guess on $\mathbf{pk}_1, \mathbf{pk}_2$ is correct, which is sufficient—when the guess is incorrect, the challenger outputs a random bit anyway, regardless of the adversary's behavior.

We now show how to transition from $\mathbf{G}_{i-1,t-1}$ to $\mathbf{G}_{i-1,t}$ for all $t \in [3, \kappa]$. It is similar to the transition from $\mathbf{G}_{i-1,0}$ to $\mathbf{G}_{i-1,2}$. Namely, the challenger guesses the first and the t 'th honest clients. If the guess is unsuccessful, the challenger aborts and outputs a random bit. As before, this incurs a security loss of $\frac{q_{\mathbf{pk}}(q_{\mathbf{pk}}-1)}{2}$. Then the challenger uses the security of the NIKE to switch the $K_{\mathbf{pk}_1, \mathbf{pk}_t}$ to uniformly random, the security of the PRF to switch the values $r_{\mathbf{pk}_1, \mathbf{pk}_t, \mathcal{U}_M^*, \ell^*}$ to uniformly random over \mathbb{A} . Since this value appears positively and negatively in the values $c_{\mathbf{pk}_1}$ and $c_{\mathbf{pk}_t}$ used by QLeftRight , we can add the offset $-u_t$ to $c_{\mathbf{pk}_1}$, which becomes:

$$x_{\mathbf{pk}_1}^b + \sum_{\mathbf{pk}' < \mathbf{pk}_1, \mathbf{pk}' \in \mathcal{U}_M^*} r_{\mathbf{pk}_1, \mathbf{pk}', \mathcal{U}_M^*, \ell^*} - \sum_{\mathbf{pk}' > \mathbf{pk}_1, \mathbf{pk}' \in \mathcal{U}_M^*} r_{\mathbf{pk}_1, \mathbf{pk}', \mathcal{U}_M^*, \ell^*} - \sum_{j \in [2, t-1]} u_j - u_t$$

and u_t to $c_{\mathbf{pk}_t}$, which becomes:

$$x_{\mathbf{pk}_t}^b + \sum_{\mathbf{pk}' < \mathbf{pk}_t, \mathbf{pk}' \in \mathcal{U}_M^*} r_{\mathbf{pk}_t, \mathbf{pk}', \mathcal{U}_M^*, \ell^*} - \sum_{\mathbf{pk}' > \mathbf{pk}_t, \mathbf{pk}' \in \mathcal{U}_M^*} r_{\mathbf{pk}_t, \mathbf{pk}', \mathcal{U}_M^*, \ell^*} + u_t,$$

as in game $\mathbf{G}_{i-1,t}$ (changes from $\mathbf{G}_{i-1,t-1}$ are highlighted in gray). Then, we switch back the value $r_{\mathbf{pk}_1, \mathbf{pk}_t, \mathcal{U}_M^*, \ell^*}$ from truly random to pseudo-random, and the key $K_{\mathbf{pk}_1, \mathbf{pk}_t}$ to pseudo-random, using the security of the PRF and the NIKE, respectively.

Now, we show that the game $\mathbf{G}_{i-1, \kappa}$ is identically distributed to \mathbf{G}_i^* , using the fact that the following are identically distributed:

$$\{u_t\}_{t \in [2, \kappa]} \quad \text{and} \quad \{u_t + \Delta_{\mathbf{pk}_t}\}_{t \in [2, \kappa]},$$

where $u_t \xleftarrow{\$} \mathbb{A}$ for all $t \in [2, \kappa]$. The leftmost distribution corresponds to $\mathbf{G}_{i-1, \kappa}$, whereas the right most distribution corresponds to \mathbf{G}_i^* . To see this, note that for all $t \in [2, \kappa]$, we have $x_{\mathbf{pk}_t}^b + \Delta_{\mathbf{pk}_t} = x_{\mathbf{pk}_t}^0$, even in case of multiple QLeftRight -queries for the same tuple $(\mathbf{pk}_t, \mathcal{U}_M^*, \ell^*)$. Moreover, by definition of the security game, we know that $\sum_{t \in [\kappa]} x_{\mathbf{pk}_t}^0 = \sum_{t \in [\kappa]} x_{\mathbf{pk}_t}^1$ unless Finalize outputs a random bit. This implies $\sum_{t \in [2, \kappa]} \Delta_{\mathbf{pk}_t} + \Delta_{\mathbf{pk}_1} = 0$. Thus, $x_{\mathbf{pk}_1}^b - \sum_{t \in [2, \kappa]} \Delta_{\mathbf{pk}_t} = x_{\mathbf{pk}_1}^b - \Delta_{\mathbf{pk}_1} = x_{\mathbf{pk}_1}^0$. This concludes the proof that $\mathbf{G}_{i-1}^* \sim_c \mathbf{G}_i^*$, therefore, $\mathbf{G}_{i-1} \sim_c \mathbf{G}_i$. \square

B.4 Theorem 26 (sel-sym-IND-Security of our IP-DDFE)

The Inner-Product DDFE scheme described in Section 7.2 is *sel-sym-IND-secure* (as per Definition 5) under the DDH assumption, assuming IP-FE is *sel-IND* secure, the AoNE scheme is *sel-sym-IND-secure*, the DSum scheme is *sel-sym-IND-secure*, and $(\mathcal{F}_K)_K$ is a secure PRF family.

Proof. We define the following games:

Game G_0 : The challenger plays as in the *sel-sym-IND* security game for DDFE.

Game G_1 : In this game, the challenger changes the way $\mathbf{s}_{\text{pk}, \mathcal{U}_M}$ and $\mathbf{s}_{\text{pk}, \mathcal{U}_K}$ are computed in QEncrypt, QLeftRight and QDKeyGen: they are now sampled uniformly at random from \mathbb{Z}_p^d , and stored for consistency (across subsequent calls to QEncrypt, QLeftRight and QDKeyGen), instead of using the PRF.

Game G_2 : In this game, the challenger changes the way it computes $d_{\text{pk}, k}$ when responding to QDKeyGen queries. Writing $k = (y_{\text{pk}'}, \text{pk}')_{\text{pk}' \in \mathcal{U}_K}$, the challenger keeps track of which $\text{pk} \in \mathcal{U}_K$ have been queried for k , and responds to queries (pk, k) for $\text{pk} \in \mathcal{U}_K$ as follows:

- If pk is *not* the last honest participant in \mathcal{U}_K for which a query remains to be made for k , then the challenger responds with:

$$\text{DSum.Encrypt}(\text{DSum.sk}_{\text{pk}}, (0, (\text{DSum.pk}')_{\text{pk}' \in \mathcal{U}_K}, k)).$$

- Otherwise, if pk is the last honest participant in \mathcal{U}_K to be queried for k , then, let \mathcal{HS} denote the subset of honest participants of \mathcal{U}_K , compute:

$$hc = \sum_{\text{pk} \in \mathcal{HS}} \mathbf{y}_{\text{pk}}^T \mathbf{s}_{\text{pk}}$$

and respond with $\text{DSum.Encrypt}(\text{DSum.sk}_{\text{pk}}, (hc, (\text{DSum.pk}')_{\text{pk}' \in \mathcal{U}_K}, k))$.

For each k , the challenger will need to memorize the last honest participant in \mathcal{U}_K queried on k so it can provide consistent responses to later queries.

Game G_3 : In this game, we change the way incomplete queries $(\text{pk}, k = (\mathbf{y}_{\text{pk}'}, \text{pk}')_{\text{pk}' \in \mathcal{U}_K})$ to QDKeyGen are answered. Such a query is called incomplete if not every $\text{pk}' \in \mathcal{U}_K$ is part of a query containing k to QDKeyGen. For these queries, QDKeyGen uses the value $d''_{\text{pk}, k} = 0$ instead of $d''_{\text{pk}, k} = \text{IP-FE.DKeyGen}(\text{IP-FE.sk}_{\text{pk}}, \mathbf{y}_{\text{pk}})$.

Game G_4 : In this game, we change the way incomplete queries $(\text{pk}, ([\mathbf{x}_{\text{pk}}^0], \mathcal{U}_M, \ell), ([\mathbf{x}_{\text{pk}}^1], \mathcal{U}_M, \ell))$ to QLeftRight are answered. A query is called incomplete if not every $\text{pk}' \in \mathcal{U}_M$ is part of a query containing (\mathcal{U}_M, ℓ) to QEncrypt or QLeftRight. In that case QLeftRight uses 0 instead of c_{pk} in the call to AoNE.Encrypt.

Game G_5 : In this game, we change the way complete queries $(\text{pk}, ([\mathbf{x}_{\text{pk}}^0], \mathcal{U}_M, \ell), ([\mathbf{x}_{\text{pk}}^1], \mathcal{U}_M, \ell))$ to QLeftRight are answered. A query is called complete if it is not incomplete. We consider repeated complete queries to the same $\text{pk}, \mathcal{U}_M, \ell$. We denote by $[\mathbf{x}_{\text{pk}, \mathcal{U}_M, \ell}^{*0}], [\mathbf{x}_{\text{pk}, \mathcal{U}_M, \ell}^{*1}]$ the pair of vectors used for the first complete query containing $\text{pk}, \mathcal{U}_M, \ell$. For any complete repeated query $(\text{pk}, ([\mathbf{x}_{\text{pk}}^0], \mathcal{U}_M, \ell), ([\mathbf{x}_{\text{pk}}^1], \mathcal{U}_M, \ell))$, QLeftRight uses $c_{\text{pk}} = \text{IP-FE.Encrypt}(\text{IP-FE.sk}_{\text{pk}}, [\mathbf{x}_{\text{pk}}^b + \mathbf{x}_{\text{pk}, \mathcal{U}_M, \ell}^{*0} - \mathbf{x}_{\text{pk}, \mathcal{U}_M, \ell}^{*b} + \mathbf{s}_{\text{pk}, \mathcal{U}_M} h_\ell])$ instead of using $c_{\text{pk}} = \text{IP-FE.Encrypt}(\text{IP-FE.sk}_{\text{pk}}, [\mathbf{x}_{\text{pk}}^b + \mathbf{s}_{\text{pk}, \mathcal{U}_M} h_\ell])$ when making the call to AoNE.Encrypt.

Game G_6 : In this game, we change the way complete queries $(\text{pk}, ([\mathbf{x}_{\text{pk}}^0], \mathcal{U}_M, \ell), ([\mathbf{x}_{\text{pk}}^1], \mathcal{U}_M, \ell))$ to QLeftRight are answered. For these, QLeftRight uses $c_{\text{pk}} = \text{IP-FE.Encrypt}(\text{IP-FE.sk}_{\text{pk}}, [\mathbf{x}_{\text{pk}}^0 + \mathbf{s}_{\text{pk}, \mathcal{U}_M} h_\ell])$ instead of $c_{\text{pk}} = \text{IP-FE.Encrypt}(\text{IP-FE.sk}_{\text{pk}}, [\mathbf{x}_{\text{pk}}^b + \mathbf{x}_{\text{pk}, \mathcal{U}_M, \ell}^{*0} - \mathbf{x}_{\text{pk}, \mathcal{U}_M, \ell}^{*b} + \mathbf{s}_{\text{pk}, \mathcal{U}_M} h_\ell])$ in the call to AoNE.Encrypt. Note that in this game, the adversary's view does not depend on b anymore.

We now show how to transition from each of those games to the next.

$\mathbf{G}_0 \sim_c \mathbf{G}_1$: This transition rests on the security of the PRF. It requires going through multiple hybrid games, and for brevity we only describe the process at a high-level. We successively switch the way each honest participant (created by `QNewHonest` and not queried to `QCorrupt`) computes its $\mathbf{s}_{\text{pk}, \mathcal{U}}$ for all \mathcal{U} (one honest participant at a time, all of the participant's \mathcal{U} 's at a time) from being honestly computed to being sampled uniformly at random in \mathbb{Z}_p^d , with one hybrid game per successive participant. If at any point, an adversary can tell the difference between two successive games with noticeable advantage, it's easy to convert this into an algorithm which distinguishes PRF outputs from real randomness with noticeable advantage, which contradicts the security claim of the PRF. The Hybrid Lemma allows us to conclude.

$\mathbf{G}_1 \sim_c \mathbf{G}_2$: We proceed by contradiction. From a PPT adversary \mathcal{A} which distinguishes between \mathbf{G}_1 and \mathbf{G}_2 with noticeable probability, we build \mathcal{B} which wins the `sel-sym-IND` security game of `DSum` as follows: use the oracles of the `DSum` challenger to naturally handle all `DSum` related operations in the IP-DDFE challenger. The exceptions are `QDKeyGen` queries: for those, write $lq = (\mathbf{y}_{\text{pk}}^T \mathbf{s}_{\text{pk}, \mathcal{U}_K}, (\text{DSum.pk}')_{\text{pk}' \in \mathcal{U}_K}, k)$ and $rq = (v, (\text{DSum.pk}')_{\text{pk}' \in \mathcal{U}_K}, k)$ where $v = hc = \sum_{\text{pk} \in \mathcal{H}_S} \mathbf{y}_{\text{pk}}^T \mathbf{s}_{\text{pk}}$ if this is the last honest participant that remained to be queried for this set \mathcal{U}_K (for some ordering on the participants), otherwise $v = 0$. Then make a query `DSum.QLeftRight` for that public key with left ciphertext lq and right ciphertext rq . When `DSum.b` = 0, \mathcal{A} is playing \mathbf{G}_1 . When `DSum.b` = 1, \mathcal{A} is playing \mathbf{G}_2 . Because we use k as a label and we chose our challenges in such a way that the sums on the left always equal the sums on the right, we never trigger the `DSum`'s Finalize conditions. This concludes this transition.

$\mathbf{G}_2 \sim_c \mathbf{G}_3$: We proceed by contradiction. From a PPT adversary \mathcal{A} which distinguishes between \mathbf{G}_2 and \mathbf{G}_3 with noticeable probability, we build \mathcal{B} which wins the `sel-sym-IND` security game of `AoNE` as follows: use the oracles of the `AoNE` challenger to naturally handle all `AoNE` related operations in the IP-DDFE challenger. The exceptions are `QDKeyGen` queries: for those, write $lq = \text{IP-FE.DKeyGen}(\text{IP-FE.sk}_{\text{pk}}, \mathbf{y}_{\text{pk}})$ and $rq = lq$ if that query is complete, $rq = 0$ otherwise. Then make a query `AoNE.QLeftRight` for that public key with left ciphertext lq and right ciphertext rq . When `AoNE.b` = 0, \mathcal{A} is playing \mathbf{G}_2 . When `AoNE.b` = 1, \mathcal{A} is playing \mathbf{G}_3 . Because we use k as a label (with the "key" prefix isolating this from encryption related queries) and we chose our challenges in such a way that the contents differ only when we know a complete request won't be made, we never trigger the `AoNE`'s Finalize conditions. This concludes this transition.

$\mathbf{G}_3 \sim_c \mathbf{G}_4$: We proceed by contradiction. From a PPT adversary \mathcal{A} which distinguishes between \mathbf{G}_3 and \mathbf{G}_4 with noticeable probability, we build \mathcal{B} which wins the `sel-sym-IND` security game of `AoNE` as follows: use the oracles of the `AoNE` challenger to naturally handle all `AoNE` related operations in the IP-DDFE challenger. The exceptions are `QEncrypt` and `QLeftRight` queries: for those, write $lq = c_{\text{pk}}$ and $rq = c_{\text{pk}}$ if that query is complete, $rq = 0$ otherwise. Then make a query `AoNE.QLeftRight` for that public key with left ciphertext lq and right ciphertext rq . When `AoNE.b` = 0, \mathcal{A} is playing \mathbf{G}_3 . When `AoNE.b` = 1, \mathcal{A} is playing \mathbf{G}_4 . Because we use \mathcal{U}_M, ℓ as a label (with the "ct" prefix isolating this from key generation related queries) and we chose our challenges in such a way that the contents differ only when we know a complete request won't be made, we never trigger the `AoNE`'s Finalize conditions. This concludes this transition.

$\mathbf{G}_4 \sim_c \mathbf{G}_5$: For any complete query $(\text{pk}, (\mathbf{x}_{\text{pk}}^0, \mathcal{U}_M, \ell), (\mathbf{x}_{\text{pk}}^1, \mathcal{U}_M, \ell))$ to `QLeftRight`, we switch the value c_{pk} from $c_{\text{pk}} = \text{IP-FE.Encrypt}(\text{IP-FE.pk}, [\mathbf{x}_{\text{pk}}^b + \mathbf{s}_{\text{pk}, \mathcal{U}_M} h_\ell])$ to $c_{\text{pk}} = \text{IP-FE.Encrypt}(\text{IP-FE.pk}, [\mathbf{x}_{\text{pk}}^b + \mathbf{x}_{\text{pk}, \mathcal{U}_M, \ell}^{*0} - \mathbf{x}_{\text{pk}, \mathcal{U}_M, \ell}^{*b} + \mathbf{s}_{\text{pk}, \mathcal{U}_M} h_\ell])$, where $(\mathbf{x}_{\text{pk}, \mathcal{U}_M, \ell}^{*0}, \mathbf{x}_{\text{pk}, \mathcal{U}_M, \ell}^{*1})$ is the pair of vectors used in the first complete query to `QLeftRight` containing $(\text{pk}, \mathcal{U}_M, \ell)$. We do so using a hybrid argument that goes over the pairs (\mathcal{U}_M, ℓ) used in complete queries to `QLeftRight`. Denoting q_c the number of such pairs, we introduce hybrid games $\mathbf{G}_{4, i-1}$ for all $i \in [q_c + 1]$, which answers the first $i - 1$ 'st pairs as in \mathbf{G}_5 , and the last $q_c - i + 1$ pairs as in \mathbf{G}_4 . Clearly, $\mathbf{G}_{4,0}$ is the same as \mathbf{G}_4 , and \mathbf{G}_{4, q_c} is the same as \mathbf{G}_5 . Thus, it suffices to show that for all $i \in [q_c + 1]$, we have $\mathbf{G}_{4, i-1} \sim_c \mathbf{G}_{4, i}$. We prove this in Lemma 27.

$\mathbf{G}_5 \sim_c \mathbf{G}_6$: We will switch the generation of c_{pk} used by `QLeftRight` on complete queries from

being computed from $\text{IP-FE.Encrypt}(\text{IP-FE.sk}_{\text{pk}}, [\mathbf{x}_{\text{pk}}^b + \mathbf{x}_{\text{pk}, \mathcal{U}_M, \ell}^{*0} - \mathbf{x}_{\text{pk}, \mathcal{U}_M, \ell}^{*b} + \mathbf{s}_{\text{pk}, \mathcal{U}_M} h_\ell])$ as in \mathbf{G}_5 to $\text{IP-FE.Encrypt}(\text{IP-FE.sk}_{\text{pk}}, [\mathbf{x}_{\text{pk}}^0 + \mathbf{s}_{\text{pk}, \mathcal{U}_M} h_\ell])$ as in \mathbf{G}_6 , one pk at a time, using a series of hybrid games. Recall that $([\mathbf{x}_{\text{pk}, \mathcal{U}_M, \ell}^{*0}], [\mathbf{x}_{\text{pk}, \mathcal{U}_M, \ell}^{*1}])$ is the pair of vectors contained in the first complete query to QLeftRight that contains $\text{pk}, \mathcal{U}_M, \ell$. Let q_p denote the number of queries made to QNewHonest . For $i \in [q_p + 1]$, we define \mathbf{H}_i as follows: all ciphertexts for pk with index $j < i$ are generated using $\text{IP-FE.Encrypt}(\text{IP-FE.sk}_{\text{pk}}, [\mathbf{x}_{\text{pk}}^b + \mathbf{x}_{\text{pk}, \mathcal{U}_M, \ell}^{*0} - \mathbf{x}_{\text{pk}, \mathcal{U}_M, \ell}^{*b} + \mathbf{s}_{\text{pk}, \mathcal{U}_M} h_\ell])$, all those for $j \geq i$ are generated using $\text{IP-FE.Encrypt}(\text{IP-FE.sk}_{\text{pk}}, [\mathbf{x}_{\text{pk}}^0 + \mathbf{s}_{\text{pk}, \mathcal{U}_M} h_\ell])$. Now we simply need to show that $\mathbf{H}_i \sim_c \mathbf{H}_{i+1}$, and we proceed by contradiction. From a PPT adversary \mathcal{A} which distinguishes between the current hybrid game and the next one with noticeable probability, we build \mathcal{B} which wins the IND security game of IP-FE (see Definition 17). Let pk_i denote the participant whose ciphertexts are being switched in this transition. First, note that when a query to QCorrupt is made for pk_i , the constraints from the symmetric security game sym-IND (see Definition 4) impose that $\mathbf{x}_{\text{pk}_i, \mathcal{U}_M, \ell}^{*0} = \mathbf{x}_{\text{pk}_i, \mathcal{U}_M, \ell}^{*1}$ and $\mathbf{x}_{\text{pk}_i}^0 = \mathbf{x}_{\text{pk}_i}^1$, so

$$\mathbf{x}_{\text{pk}_i}^b + \mathbf{x}_{\text{pk}_i, \mathcal{U}_M, \ell}^{*0} - \mathbf{x}_{\text{pk}_i, \mathcal{U}_M, \ell}^{*b} = \mathbf{x}_{\text{pk}_i}^0 + \mathbf{x}_{\text{pk}_i, \mathcal{U}_M, \ell}^{*0} - \mathbf{x}_{\text{pk}_i, \mathcal{U}_M, \ell}^{*0} = \mathbf{x}_{\text{pk}_i}^0$$

and the adversary's view does not depend on the game. Thus for \mathcal{A} to distinguish between \mathbf{H}_i and \mathbf{H}_{i+1} with noticeable advantage, it must be that, with noticeable probability, \mathcal{A} does not corrupt pk_i , and, conditioned on that event, \mathcal{A} retains noticeable advantage. \mathcal{B} can thus safely assume that \mathcal{A} will not corrupt pk_i , abort and guess at random otherwise. Now proceed as follows: use the oracles of the IP-FE challenger to naturally handle all IP-FE related operations in the IP-DDFE challenger for participant with public key pk_i . The exceptions are QLeftRight queries. For these, write $lq = \mathbf{x}_{\text{pk}_i}^b - \mathbf{x}_{\text{pk}_i, \mathcal{U}_M, \ell}^{*b}$ and $rq = \mathbf{x}_{\text{pk}_i}^0 - \mathbf{x}_{\text{pk}_i, \mathcal{U}_M, \ell}^{*0}$. Then the reduction makes a query to IP-FE.QLeftRight with left ciphertext lq and right ciphertext rq , upon which it receives the ciphertext c'_{pk_i} . The reduction \mathcal{B} computes $c_{\text{pk}_i} = \text{IP-FE.Add}(c'_{\text{pk}_i}, [\mathbf{x}_{\text{pk}_i, \mathcal{U}_M, \ell}^{*0} + \mathbf{s}_{\text{pk}_i, \mathcal{U}_M} \cdot h_\ell])$, and runs AoNE.Encrypt on it. By Property 18, c_{pk_i} is distributed as $\text{IP-FE.Encrypt}(\text{IP-FE.sk}_{\text{pk}_i}, [\mathbf{x}_{\text{pk}_i}^b + \mathbf{x}_{\text{pk}_i, \mathcal{U}_M, \ell}^{*0} - \mathbf{x}_{\text{pk}_i, \mathcal{U}_M, \ell}^{*b} + \mathbf{s}_{\text{pk}_i, \mathcal{U}_M} \cdot h_\ell])$, which is as \mathbf{H}_i , when $\text{IP-FE.b} = 0$, whereas it is distributed as $\text{IP-FE.Encrypt}(\text{IP-FE.sk}_{\text{pk}_i}, [\mathbf{x}_{\text{pk}_i}^0 + \mathbf{s}_{\text{pk}_i, \mathcal{U}_M} \cdot h_\ell])$, which is as in \mathbf{H}_{i+1} when $\text{IP-FE.b} = 1$.

The reduction \mathcal{B} 's queries to its oracle will not trigger the Finalize to output a random bit in the security game for IP-FE, as long as, for any \mathbf{y}_{pk_i} that is part of a complete query $(\text{pk}_i, k = (\mathbf{y}_{\text{pk}'}, \mathcal{U}_K = \mathcal{U}_M)_{\text{pk}' \in \mathcal{U}_K})$ to QDKeyGen on pk_i , we have:

$$\mathbf{y}_{\text{pk}_i}^\top (\mathbf{x}_{\text{pk}_i}^b - \mathbf{x}_{\text{pk}_i, \mathcal{U}_M, \ell}^{*b}) = \mathbf{y}_{\text{pk}_i}^\top (\mathbf{x}_{\text{pk}_i}^0 - \mathbf{x}_{\text{pk}_i, \mathcal{U}_M, \ell}^{*0}).$$

Since $(\text{pk}_i, (\mathbf{x}_{\text{pk}_i}^0, \mathcal{U}_M, \ell), (\mathbf{x}_{\text{pk}_i}^1, \mathcal{U}_M, \ell))$ is a complete query, we have for all $\text{pk}' \in \mathcal{HS} \setminus \{\text{pk}_i\}$, a query of the form $(\text{pk}', (\mathbf{x}_{\text{pk}'}^0, \mathcal{U}_M, \ell), (\mathbf{x}_{\text{pk}'}^1, \mathcal{U}_M, \ell))$ to QLeftRight , or a query of the form $(\text{pk}', (\mathbf{x}_{\text{pk}'}^0 = \mathbf{x}_{\text{pk}'}^1, \mathcal{U}_M, \ell))$ to QEncrypt . Thus, the Finalize procedure of the IP-DDFE security game outputs a random bit unless the following holds:

$$\mathbf{y}_{\text{pk}_i}^\top \mathbf{x}_{\text{pk}_i}^0 + \sum_{\text{pk}' \in \mathcal{HS} \setminus \{\text{pk}_i\}} \mathbf{y}_{\text{pk}'}^\top \mathbf{x}_{\text{pk}'}^0 = \mathbf{y}_{\text{pk}_i}^\top \mathbf{x}_{\text{pk}_i}^1 + \sum_{\text{pk}' \in \mathcal{HS} \setminus \{\text{pk}_i\}} \mathbf{y}_{\text{pk}'}^\top \mathbf{x}_{\text{pk}'}^1,$$

and

$$\mathbf{y}_{\text{pk}_i}^\top \mathbf{x}_{\text{pk}_i, \mathcal{U}_M, \ell}^{*0} + \sum_{\text{pk}' \in \mathcal{HS} \setminus \{\text{pk}_i\}} \mathbf{y}_{\text{pk}'}^\top \mathbf{x}_{\text{pk}'}^0 = \mathbf{y}_{\text{pk}_i}^\top \mathbf{x}_{\text{pk}_i, \mathcal{U}_M, \ell}^{*1} + \sum_{\text{pk}' \in \mathcal{HS} \setminus \{\text{pk}_i\}} \mathbf{y}_{\text{pk}'}^\top \mathbf{x}_{\text{pk}'}^1,$$

which together imply that

$$\mathbf{y}_{\text{pk}_i}^\top (\mathbf{x}_{\text{pk}_i}^b - \mathbf{x}_{\text{pk}_i, \mathcal{U}_M, \ell}^{*b}) = \mathbf{y}_{\text{pk}_i}^\top (\mathbf{x}_{\text{pk}_i}^0 - \mathbf{x}_{\text{pk}_i, \mathcal{U}_M, \ell}^{*0}).$$

□

Lemma 27. For all $i \in [q_c + 1]$, we have $\mathbf{G}_{4, i-1} \sim_c \mathbf{G}_{4, i}$.

Proof. We use the following hybrid games.

$\mathbf{G}_{4,i-1,1}$: is as $\mathbf{G}_{4,i-1}$, except for all pk that are part of a complete query to QLeftRight that contains $(\mathcal{U}_{M,i}, \ell_i)$, QLeftRight answers any query containing pk using $c_{\text{pk}} = \text{IP-FE.Encrypt}(\text{IP-FE.pk}, [\mathbf{x}_{\text{pk}}^b + s_{\text{pk}, \mathcal{U}_{M,i}} h_{\ell_i} + \gamma h_{\ell_i} (\mathbf{x}_{\text{pk}, \mathcal{U}_{M,i}, \ell_i}^{*0} - \mathbf{x}_{\text{pk}, \mathcal{U}_{M,i}, \ell_i}^{*b})])$, for a $\gamma \xleftarrow{\$} \mathbb{Z}_p$ (the same γ is used for all outputs). These two games are identically distributed, as can be shown using the fact that the following are identically distributed:

$$s_{\text{pk}, \mathcal{U}_{M,i}} \quad \text{and} \quad s_{\text{pk}, \mathcal{U}_{M,i}} + \gamma(\mathbf{x}_{\text{pk}, \mathcal{U}_{M,i}, \ell_i}^{*0} - \mathbf{x}_{\text{pk}, \mathcal{U}_{M,i}, \ell_i}^{*b}),$$

where $s_{\text{pk}, \mathcal{U}_{M,i}} \xleftarrow{\$} \mathbb{Z}_p$. Note that we can use this fact since we are in a selective game, where the challenge $\mathbf{x}_{\text{pk}, \mathcal{U}_{M,i}, \ell_i}^{*0}, \mathbf{x}_{\text{pk}, \mathcal{U}_{M,i}, \ell_i}^{*b}$ is chosen before, and thus, independently, of the value $s_{\text{pk}, \mathcal{U}_{M,i}} \xleftarrow{\$} \mathbb{Z}_p$. The leftmost distribution corresponds to $\mathbf{G}_{4,i-1}$, whereas the rightmost distribution corresponds to $\mathbf{G}_{4,i-1,1}$. This is due to the fact that the extra terms $\gamma(\mathbf{x}_{\text{pk}, \mathcal{U}_{M,i}, \ell_i}^{*0} - \mathbf{x}_{\text{pk}, \mathcal{U}_{M,i}, \ell_i}^{*b})$ only appears in the outputs of QDKeyGen on complete queries as:

$$\sum_{\text{pk}' \in \mathcal{U}_{M,i}} \mathbf{y}_{\text{pk}'}^\top s_{\text{pk}', \mathcal{U}_{M,i}} + \underbrace{\gamma \sum_{\text{pk}' \in \mathcal{U}_{M,i}} \mathbf{y}_{\text{pk}'}^\top (\mathbf{x}_{\text{pk}', \mathcal{U}_{M,i}, \ell_i}^{*0} - \mathbf{x}_{\text{pk}', \mathcal{U}_{M,i}, \ell_i}^{*b})}_{=0},$$

where the sum equals $\mathbf{0}$ unless the Finalize procedure of the IP-DDFE outputs a random bit. That is, the extra terms actually only appear in the output of QLeftRight on complete queries that contain $(\mathcal{U}_{M,i}, \ell_i)$, as defined in $\mathbf{G}_{4,i-1,1}$.

$\mathbf{G}_{4,i-1,2}$: is as $\mathbf{G}_{4,i-1,1}$, except that for all pk that are part of a complete query to QLeftRight that contains $(\mathcal{U}_{M,i}, \ell_i)$, QLeftRight answers any query containing pk using $c_{\text{pk}} = \text{IP-FE.Encrypt}(\text{IP-FE.pk}, [\mathbf{x}_{\text{pk}}^b + s_{\text{pk}, \mathcal{U}_{M,i}} h_{\ell_i} + \omega_{\ell_i} (\mathbf{x}_{\text{pk}, \mathcal{U}_{M,i}, \ell_i}^{*0} - \mathbf{x}_{\text{pk}, \mathcal{U}_{M,i}, \ell_i}^{*b})])$, where $\omega_{\ell_i} \xleftarrow{\$} \mathbb{Z}_p$. This transition is justified by the DDH assumption in \mathbb{G} , which states that $([h_{\ell_i}], [\gamma], [\gamma h_{\ell_i}]) \sim_c ([h_{\ell_i}], [\gamma], [\omega_{\ell_i}])$, where $\omega_{\ell_i} \xleftarrow{\$} \mathbb{Z}_p$. The leftmost distribution corresponds to $\mathbf{G}_{4,i-1,1}$, whereas the rightmost distribution corresponds to $\mathbf{G}_{4,i-1,2}$.

$\mathbf{G}_{4,i-1,3}$: is as $\mathbf{G}_{4,i-1,2}$, except that for all pk that are part of a complete query to QLeftRight that contains $(\mathcal{U}_{M,i}, \ell_i)$, QLeftRight answers any query containing pk using $c_{\text{pk}} = \text{IP-FE.Encrypt}(\text{IP-FE.pk}, [\mathbf{x}_{\text{pk}}^b + \mathbf{x}_{\text{pk}, \mathcal{U}_{M,i}, \ell_i}^{*0} - \mathbf{x}_{\text{pk}, \mathcal{U}_{M,i}, \ell_i}^{*b} + s_{\text{pk}, \mathcal{U}_{M,i}} h_{\ell_i} + \omega_{\ell_i} (\mathbf{x}_{\text{pk}, \mathcal{U}_{M,i}, \ell_i}^{*0} - \mathbf{x}_{\text{pk}, \mathcal{U}_{M,i}, \ell_i}^{*b})])$, where $\omega_{\ell_i} \xleftarrow{\$} \mathbb{Z}_p$. These two games are identically distributed, as can be shown using the fact that the following are identically distributed:

$$\omega_{\ell_i} \quad \text{and} \quad \omega_{\ell_i} + 1,$$

where $\omega_{\ell_i} \xleftarrow{\$} \mathbb{Z}_p$. The leftmost distribution corresponds to $\mathbf{G}_{4,i-1,2}$, whereas the rightmost distribution corresponds to $\mathbf{G}_{4,i-1,3}$.

$\mathbf{G}_{4,i-1,4}$: is as $\mathbf{G}_{4,i-1,3}$, except for all pk that are part of a complete query to QLeftRight that contains $(\mathcal{U}_{M,i}, \ell_i)$, QLeftRight answers any query containing pk using $c_{\text{pk}} = \text{IP-FE.Encrypt}(\text{IP-FE.pk}, [\mathbf{x}_{\text{pk}}^b + \mathbf{x}_{\text{pk}, \mathcal{U}_{M,i}, \ell_i}^{*0} - \mathbf{x}_{\text{pk}, \mathcal{U}_{M,i}, \ell_i}^{*b} + s_{\text{pk}, \mathcal{U}_{M,i}} h_{\ell_i} + \gamma h_{\ell_i} (\mathbf{x}_{\text{pk}, \mathcal{U}_{M,i}, \ell_i}^{*0} - \mathbf{x}_{\text{pk}, \mathcal{U}_{M,i}, \ell_i}^{*b})])$. This transition is reverse to the transition from $\mathbf{G}_{4,i-1,1}$ to $\mathbf{G}_{4,i-1,2}$. Similarly, it uses the DDH assumption in \mathbb{G} .

$\mathbf{G}_{4,i}$: is as $\mathbf{G}_{4,i-1,4}$, except for all pk that are part of a complete query to QLeftRight that contains $(\mathcal{U}_{M,i}, \ell_i)$, QLeftRight answers any query containing pk using $c_{\text{pk}} = \text{IP-FE.Encrypt}(\text{IP-FE.pk}, [\mathbf{x}_{\text{pk}}^b + \mathbf{x}_{\text{pk}, \mathcal{U}_{M,i}, \ell_i}^{*0} - \mathbf{x}_{\text{pk}, \mathcal{U}_{M,i}, \ell_i}^{*b} + s_{\text{pk}, \mathcal{U}_{M,i}} h_{\ell_i}])$. This transition is reverse to the transition from $\mathbf{G}_{4,i-1}$ to $\mathbf{G}_{4,i-1,1}$. Similarly, it uses the fact that the following are identically distributed:

$$s_{\text{pk}, \mathcal{U}_{M,i}, \ell_i} \quad \text{and} \quad s_{\text{pk}, \mathcal{U}_{M,i}, \ell_i} + \gamma(\mathbf{x}_{\text{pk}, \mathcal{U}_{M,i}, \ell_i}^{*0} - \mathbf{x}_{\text{pk}, \mathcal{U}_{M,i}, \ell_i}^{*b}),$$

where $s_{\text{pk}, \mathcal{U}_{M,i}, \ell_i} \xleftarrow{\$} \mathbb{Z}_p$. This concludes the proof that $\mathbf{G}_{4,i-1} \sim_c \mathbf{G}_{4,i}$. \square

Remark 28. Note that while we have to prove a symmetric variant of security, our IP-DDFE scheme instantiated with our constructions is not actually symmetric: the asymmetry of the AoNE layer prevents decrypting a lone ciphertext, even with knowledge of the public key. However, once enough queries have been made that a complete ciphertext is available, the AoNE layer can be removed, and the message can now be recovered, meaning the caveat of symmetric security is necessary.