





Multi-Input Functional Encryption for Inner Products: Function-Hiding Realizations and Constructions Without Pairings

Michel Abdalla^{1,2} , Dario Catalano³, Dario Fiore⁴ , Romain Gay^{1,2},
and Bogdan Ursu⁵

¹ Département informatique de l'ENS, École normale supérieure, CNRS,
PSL University, 75005 Paris, France
{michel.abdalla,romain.gay}@ens.fr

² INRIA, Paris, France

³ Dipartimento di Matematica e Informatica, Università di Catania, Catania, Italy
catalano@dmi.unict.it

⁴ IMDEA Software Institute, Madrid, Spain
dario.fiore@imdea.org

⁵ KIT, Karlsruhe, Germany
bogdan.ursu@kit.edu

Abstract. We present new constructions of multi-input functional encryption (MIFE) schemes for the inner-product functionality that improve the state of the art solution of Abdalla *et al.* (Eurocrypt 2017) in two main directions.

First, we put forward a novel methodology to convert single-input functional encryption for inner products into multi-input schemes for the same functionality. Our transformation is surprisingly simple, general and efficient. In particular, it does not require pairings and it can be instantiated with *all* known single-input schemes. This leads to two main advances. First, we enlarge the set of assumptions this primitive can be based on, notably, obtaining new MIFEs for inner products from plain DDH, LWE, and Decisional Composite Residuosity. Second, we obtain the first MIFE schemes from standard assumptions where decryption works efficiently even for messages of super-polynomial size.

Our second main contribution is the first function-hiding MIFE scheme for inner products based on standard assumptions. To this end, we show how to extend the original, pairing-based, MIFE by Abdalla *et al.* in order to make it function hiding, thus obtaining a function-hiding MIFE from the MDDH assumption.

1 Introduction

Functional Encryption (FE) [8, 15, 16] is an emerging cryptographic paradigm that allows fine-grained access control over encrypted data. Functional encryption schemes come equipped with a key generation mechanism that allows the

owner of a master secret key to generate decryption keys that have a somewhat restricted capability. Namely, each decryption key sk_f is associated with a function f and using sk_f to decrypt a ciphertext $\text{Enc}(x)$ allows for recovering $f(x)$, with the guarantee that no more information about x is revealed. The basic notion of functional encryption considers functionalities where all the inputs are provided and encrypted by a single party. The more general case of multi-input functionalities is captured by the notion of multi-input functional encryption (MIFE, for short) [12]. Informally, this notion can be thought of as an FE scheme where n encryption slots are explicitly given, in the sense that a user who is assigned the i -th slot can, independently, create a ciphertext $\text{Enc}(x_i)$ from his own plaintext x_i . Given ciphertexts $\text{Enc}(x_1), \dots, \text{Enc}(x_n)$, one can use a secret key sk_f to retrieve $f(x_1, \dots, x_n)$, similarly to the basic FE notion. This multi-input capability makes MIFE particularly well suited for many real life scenarios (such as data mining over encrypted data or multi-client delegation of computation) where the (encrypted) data may come from different and unrelated sources.

The security requirement for both FE and MIFE imposes that decryption keys should be collusion resistant. This means that a group of users, holding different decryption keys, should not be able to gain information about the encrypted messages, beyond the union of what they can individually learn. More precisely, the standard notion of security for functional encryption is *indistinguishability*. Informally, this states that an adversary that obtains the secret keys corresponding to functions f_1, \dots, f_n should not be able to decide which of the challenge messages x_0, x_1 was encrypted, as long as $f_i(x_0) = f_i(x_1)$ for all i . This indistinguishability notion has been put forward in [8, 15] and it has been shown inadequate for certain cases (see [8, 15] for details). They also proposed an alternative simulation-based security notion which is also problematic as, for instance, it cannot be satisfied in general.

As an additional security property, functional encryption schemes might also be required to guarantee so-called *function hiding*. Intuitively, this means that a secret key sk_f should not reveal information about the function f it encodes, beyond what is implicitly leaked by $f(x)$. Slightly more in detail, in the indistinguishability setting, this is formalized by imposing that the adversary should not be able to decide for which of the challenge functions $f_i^{(0)}, f_i^{(1)}$ it is holding secret keys, as long as $f_i^{(0)}(x_0) = f_i^{(1)}(x_1)$ for all i . Over the last few years, functional encryption has attracted a lot of interest, both in its basic and in its multi-input incarnations. Known results can be broadly categorized as focusing on (1) feasibility results for general functionalities, and on (2) concrete, efficient realizations for restricted functionalities of practical interest.

For the specific case of MIFE, which is the focus of this paper, constructions of the first type [6, 7, 9, 12] all rely on quite unstable assumptions, such as indistinguishability obfuscation or multilinear maps¹. The only known construction of the second category has been recently proposed by Abdalla *et al.* in [4].

¹ Here we only consider schemes where *unbounded* collusions are allowed. See [9] and references therein for the bounded collusions case.

There, they propose a (secret-key) MIFE scheme for the inner product functionality that relies on the standard k -linear assumption in (prime-order) bilinear groups². Remarkably, their scheme allows for unbounded collusions and supports any (polynomially bounded) number of encryption slots. On the negative side, as in previous discrete-log-based constructions of functional inner-product encryption schemes, it employs an inefficient decryption procedure that requires to extract discrete logarithms and thus imposes serious restrictions on the size of supported messages. Moreover, the scheme is not function hiding as decryption requires the function f to be provided explicitly in the clear.

1.1 Our Contributions

In this paper we propose new constructions of multi-input functional encryption schemes for the inner product functionality that address the aforementioned shortcomings of the state-of-the-art solution of Abdalla et al. [4].

MIFE for inner products without pairings. Our first contribution consists of (secret-key) MIFE schemes for inner products based on a variety of assumptions, notably *without the need of bilinear maps*, and where *decryption works efficiently*, even for messages of super-polynomial size. We achieve this result by proposing a generic construction of MIFE from any single-input FE (for inner products) in which the encryption algorithm is linearly-homomorphic. Our transformation is surprisingly simple, general and efficient. In particular, it does not require pairings (as in the case of [4]), and it can be instantiated with *all* known single-input functional encryption schemes (e.g., [1, 2, 5]). This allows us to obtain new MIFE for inner products from plain DDH, composite residuosity, and LWE. Beyond the obvious advantage of enlarging the set of assumptions on which MIFE can be based, this result yields schemes that can be used with a much larger message space. Indeed, dropping the bilinear groups requirement allows us to employ schemes where the decryption time is polynomial, rather than exponential, in the message bit size. From a more theoretical perspective, our results also show that, contrary to what was previously conjectured [4], MIFE for inner product does not need any (qualitatively) stronger assumption than their single-input counterpart.

OUR SOLUTION, IN MORE DETAIL. To better describe our solution, let us first explain the basic ideas behind Abdalla *et al.*'s scheme [4]. Informally, the latter builds upon a clever two-step decryption blueprint. The ciphertexts $\text{ct}_1 = \text{Enc}(\mathbf{x}_1), \dots, \text{ct}_n = \text{Enc}(\mathbf{x}_n)$ (corresponding to slots $1, \dots, n$) are all created using different instances of a single-input FE. Decryption is performed in two stages. One first decrypts each single ct_i separately using the secret key sk_{y_i} of the underlying single-input FE, and then the outputs of these decryptions are added up to get the final result.

² As discussed in detail in [4], we stress that in the public key setting, MIFE for inner products is both easy to achieve (from its single-input counterpart) and of very limited interest, because of its inherent leakage.

The main technical challenge of this approach is that the stage one of the above decryption algorithm leaks information on each partial inner product $\langle \mathbf{x}_i, \mathbf{y}_i \rangle$. To avoid this leakage, their idea is to let source i encrypt its plaintext vector \mathbf{x}_i augmented with some fixed (random) value u_i , which is part of the secret key. Moreover, $\mathbf{sk}_{\mathbf{y}_i}$ are built by running the single-input FE key generation algorithm on input $\mathbf{y}_i || r$, i.e., the vector \mathbf{y}_i augmented with fresh randomness r .

By these modifications, and skipping many technical details, stage-one decryption then consists of using pairings to compute, in \mathbb{G}_T , the values³ $[\langle \mathbf{x}_i, \mathbf{y}_i \rangle + u_i r]_T$ for every slot i . From these quantities, the result $[\langle \mathbf{x}, \mathbf{y} \rangle]_T$ is obtained as

$$\prod_{i=1}^n [\langle \mathbf{x}_i, \mathbf{y}_i \rangle + u_i r]_T \cdot [-(\sum_{i=1}^n u_i) r]_T$$

which can be easily computed if $[-(\sum_{i=1}^n u_i) r]_T$ is included in the secret key.

Intuitively, the scheme is secure as the quantities $[u_i r]_T$ are all pseudo random (under the DDH assumption) and thus hide all the partial information $[\langle \mathbf{x}_i, \mathbf{y}_i \rangle + u_i r]_T$ may leak. Notice that, in order for this argument to go through, it is crucial that the quantities $[\langle \mathbf{x}_i, \mathbf{y}_i \rangle + u_i r]_T$ are all encoded in the exponent, and thus decoding is possible only for small norm exponents. Furthermore, this technique seems to inherently require pairings, as both u_i and r have to remain hidden while allowing to compute an encoding of their product at decryption time. This is why the possibility of a scheme without pairings was considered as “quite surprising” in [4].

We overcome these difficulties via a new FE to MIFE transform, which manages to avoid leakage in a much simpler and efficient way. Our transformation works in two steps. First, we consider a simplified scheme where only one ciphertext query is allowed and messages live in the ring \mathbb{Z}_L , for some integer L . In this setting, we build the following multi-input scheme. For each slot i the (master) secret key for slot i consists of one random vector $\mathbf{u}_i \in \mathbb{Z}_L^m$. Encrypting \mathbf{x}_i merely consists in computing $\mathbf{c}_i = \mathbf{x}_i + \mathbf{u}_i \bmod L$. The secret key for function $\mathbf{y} = (\mathbf{y}_1, \dots, \mathbf{y}_n)$, is just $z_{\mathbf{y}} = \sum_{i=1}^n \langle \mathbf{u}_i, \mathbf{y}_i \rangle \bmod L$. To decrypt, one computes

$$\langle \mathbf{x}, \mathbf{y} \rangle \bmod L = \langle (\mathbf{c}_1, \dots, \mathbf{c}_n), \mathbf{y} \rangle - z_{\mathbf{y}} \bmod L$$

Security comes from the fact that, if only one ciphertext query is allowed, the above can be seen as the functional encryption equivalent of the one-time pad⁴.

Next, to guarantee security in the more challenging setting where many ciphertext queries are allowed, we just add a layer of (functional) encryption on top of the above one-time encryption. More specifically, we encrypt each \mathbf{c}_i using a FE (supporting inner products) that is both linearly homomorphic and whose message space is compatible with L . So, given ciphertexts $\{\text{ct}_i = \text{Enc}(\mathbf{c}_i)\}$

³ Here we implicitly adopt the, by now standard, bracket notation from [11].

⁴ We remark that a similar information theoretic construction was put forward by Wee in [17], as a warm-up scheme towards an FE for inner products achieving simulation security.

and secret key $\text{sk}_{\mathbf{y}} = (\{\text{sk}_{\mathbf{y}_i}\}_i, z_{\mathbf{y}})$, one can first obtain $\{\langle \mathbf{c}_i, \mathbf{y}_i \rangle = \text{Dec}(\text{ct}_i, \text{sk}_{\mathbf{y}_i})\}$, and then extract the result as $\langle \mathbf{x}, \mathbf{y} \rangle = \sum_{i=1}^n \langle \mathbf{c}_i, \mathbf{y}_i \rangle - \langle \mathbf{u}, \mathbf{y} \rangle$.

Our transformation actually comes in two flavors: the first one addresses the case where the underlying FE computes inner products over some finite ring \mathbb{Z}_L ; the second one instead considers FE schemes that compute bounded-norm inner products over the integers. In both cases the transformations are generic enough to be instantiated with known single-input FE schemes for inner products. This gives us new MIFE relying on plain DDH [1], LWE [5] and Composite residuosity [2, 5]. Moreover, the proposed transform is security-preserving in the sense that, if the underlying FE achieves adaptive security, so does our resulting MIFE.

Function-Hiding MIFE for inner products. Our second contribution are new MIFE schemes for inner products that achieve function hiding. Our constructions build on the pairing-based solution from [4] and, as such, they also rely on pairings. More precisely, we propose transformations that, starting from the MIFE from [4], build function hiding MIFEs using single input FE for inner products as additional building block. Ours transforms are generic with respect to this latter component, in the sense that they can be instantiated using any single input FE satisfying some natural additional requirements (details of which are given in Sect. 4).

Our methods build from the two-layer encryption technique recently developed by Lin [13] to generically achieve function hiding in the context of (single input) FE for inner products. Intuitively, Lin’s idea consists in doing similar operations both at encryption and at key derivation time. Starting from two independent instances of the underlying FE, an “inner” one and an “outer” one, the idea is to encrypt the plaintext \mathbf{x} in two steps. One first uses the “inner” FE to compute $\text{ct}_1 = \text{Enc}(\text{msk}_1, \mathbf{x})$ and then “extracts” the key corresponding to ct_1 , i.e., $\text{ct}_2 = \text{KeyGen}(\text{msk}_2, \text{ct}_1)$. Key derivation is done similarly, one first computes $\text{sk}_1 = \text{KeyGen}(\text{msk}_1, \mathbf{y})$ and then encrypts sk_1 using the outer scheme, i.e., $\text{sk}_2 = \text{Enc}(\text{msk}_2, \text{sk}_1)$.

If one encodes ciphertexts in \mathbb{G}_1 and secret keys in \mathbb{G}_2 , then one can use pairings to compute an encoding, in \mathbb{G}_T , of $[\langle \text{ct}_2, \text{sk}_2 \rangle]_T$. Since decryption essentially performs inner product, the latter computation actually decrypts also the inner ct_1 component using secret key sk_1 , thus yielding an encoding of $\langle \mathbf{x}, \mathbf{y} \rangle$. Moreover, since now \mathbf{y} is encrypted, the FE security also provides function hiding⁵.

An obvious drawback of Lin’s transformation is that, when applied generically, it would induce an extra-level of multilinearity in the process. This means that, starting from a pairing-free FE for inner products, one ends up with a scheme that is function hiding but also pairing-based.

We propose similar two-layer encryption techniques that *do not*, inherently, induce extra levels of multi-linearity with respect to those of the underlying primitives. Our transforms achieve this by using the MIFE from [4] as inner scheme and, several instances of, a single input FE, one for each encryption slot, as outer

⁵ Actually the transform sketched here only manages to guarantee a weaker form of function hiding. However this can be generically turned into standard function hiding [14], as described in the full version of the paper [3].

schemes. In particular, by carefully exploiting the specific algebraic properties of the MIFE, we manage to achieve function hiding from the Matrix Decisional Diffie Hellman assumption over *standard* bilinear groups (i.e., without resorting to multi-linear maps). Specifically, our schemes come in two flavors: a simpler one for selective security and a more convoluted one achieving adaptive security. A high level overview of our technique appears in Sect. 4. The MIFE schemes from Lin [13] are selectively secure and function-hiding, but are based on multi-linear maps ($d - 1$ slots require a multilinear map of degree d). In comparison, our schemes support a polynomial number of inputs and achieve adaptive-security, while using only pairings and while being based only on standard assumptions.

GENERALITY OF OUR APPROACH. As mentioned above, our function-hiding transforms are not entirely generic as they impose restrictions on the underlying MIFE. These restrictions, while compatible with the pairing-based realization from [4], do not cope well with our newly constructed MIFEs without pairings. Very informally, this is due to the fact that our transform relies on the two-step decryption blueprint in which one learns $[\langle \mathbf{x}_i, \mathbf{y}_i \rangle + z_i]$, and each z_i is “sufficiently” random to guarantee security in the MIFE security experiment. Specifically, in Abdalla et al.’s scheme $z_i = u_i r$ whereas in our new scheme $z_i = \langle \mathbf{u}_i, \mathbf{y}_i \rangle$. While the latter value is sufficiently random in the MIFE indistinguishability experiment, this is no longer the case in the function-hiding experiment, where the adversary asks for pairs of keys $(\mathbf{y}^0, \mathbf{y}^1)$, and $z_i = \langle \mathbf{u}_i, \mathbf{y}_i^\beta \rangle$ may actually leak information about which of the two keys was chosen (i.e. information about the value of the bit β). With a different interpretation, if one sees $[\langle \mathbf{x}_i, \mathbf{y}_i \rangle + z_i]$ as a secret sharing of $\langle \mathbf{x}, \mathbf{y} \rangle$, then in our new scheme this secret sharing depends on the function \mathbf{y} whereas in [4] this is function independent and more suitable for function-hiding. We believe that coming up with more powerful transforms, capable of exploiting the potential of our efficient MIFEs, is a very natural and interesting open problem.

CONCURRENT WORK ON FUNCTION-HIDING. Concurrently and independently of our work, Datta et al. [10] proposed a multi-input function-hiding scheme for inner products. Their construction uses the framework of dual pairing vector spaces and require the use of pairings. They achieve slightly shorter ciphertexts and decryption keys (ciphertexts are shorter by 2 group elements, while decryption keys require $2n+1$ less group elements). However, this comes at the expense of a larger master secret key, which contains $4n(m^2 - 1)$ more group elements (a quadratic blow-up in m).

Interestingly, Datta et al. [10] also provide a technique based on pseudo-random functions to extend their multi-input function-hiding scheme to an unbounded number of slots. Although their techniques also appear to be applicable to our schemes, hence capable of extending both the pairing-free and the pairing-based constructions to the unbounded setting, we leave it as future work.

2 Preliminaries

Notation. We denote with $\lambda \in \mathbb{N}$ a security parameter. A *probabilistic polynomial time* (PPT) algorithm \mathcal{A} is a randomized algorithm for which there exists a polynomial $p(\cdot)$ such that for every input x the running time of $\mathcal{A}(x)$ is bounded by $p(|x|)$. We say that a function $\varepsilon : \mathbb{N} \rightarrow \mathbb{R}^+$ is *negligible* if for every positive polynomial $p(\lambda)$ there exists $\lambda_0 \in \mathbb{N}$ such that for all $\lambda > \lambda_0$: $\varepsilon(\lambda) < 1/p(\lambda)$. If S is a set, $x \leftarrow_{\mathbb{R}} S$ denotes the process of selecting x uniformly at random in S . If \mathcal{A} is a probabilistic algorithm, $y \leftarrow_{\mathbb{R}} \mathcal{A}(\cdot)$ denotes the process of running \mathcal{A} on some appropriate input and assigning its output to y . For a positive integer n , we denote by $[n]$ the set $\{1, \dots, n\}$. We denote vectors $\mathbf{x} = (x_i)$ and matrices $\mathbf{A} = (a_{i,j})$ in bold. For a set S (resp. vector \mathbf{x}) $|S|$ (resp. $|\mathbf{x}|$) denotes its cardinality (resp. number of entries). Also, given two vectors \mathbf{x} and \mathbf{x}' we denote by $\mathbf{x} \parallel \mathbf{x}'$ their concatenation. By \equiv , we denote the equality of statistical distributions, and for any $\varepsilon > 0$, we denote by \approx_{ε} the ε -statistical difference of two distributions.

2.1 Definitions for Multi-Input Functional Encryption

In this section we recall the definitions of multi-input functional encryption [12] specialized to the private-key setting, as this is the one relevant for our constructions.

Definition 1 (Multi-input Function Encryption). Let $\mathcal{F} = \{\mathcal{F}_n\}_{n \in \mathbb{N}}$ be an ensemble where each \mathcal{F}_n is a family of n -ary functions. A function $f \in \mathcal{F}_n$ is defined as follows $f : \mathcal{X}_1 \times \dots \times \mathcal{X}_n \rightarrow \mathcal{Y}$. A multi-input functional encryption scheme *MIFE* for \mathcal{F} consists of the following algorithms:

- $\text{Setup}(1^\lambda, \mathcal{F}_n)$ takes as input the security parameter λ and a description of $\mathcal{F}_n \in \mathcal{F}$, and outputs a master public key mpk ⁶ and a master secret key msk . The master public key mpk is assumed to be part of the input of all the remaining algorithms.
- $\text{Enc}(\text{msk}, i, x_i)$ takes as input the master secret key msk , an index $i \in [n]$, and a message $x_i \in \mathcal{X}_i$, and it outputs a ciphertext ct . Each ciphertext is assumed to be associated with an index i denoting for which slot this ciphertext can be used for. When $n = 1$, the input i is omitted.
- $\text{KeyGen}(\text{msk}, f)$ takes as input the master secret key msk and a function $f \in \mathcal{F}_n$, and it outputs a decryption key sk_f .
- $\text{Dec}(\text{sk}_f, \text{ct}_1, \dots, \text{ct}_n)$ takes as input a decryption key sk_f for function f and n ciphertexts, and it outputs a value $y \in \mathcal{Y}$.

A scheme *MIFE* as defined above is correct if for all $n \in \mathbb{N}$, $f \in \mathcal{F}_n$ and all $x_i \in \mathcal{X}_i$ for $1 \leq i \leq n$, we have

$$\Pr \left[\begin{array}{l} (\text{mpk}, \text{msk}) \leftarrow \text{Setup}(1^\lambda, \mathcal{F}_n); \quad \text{sk}_f \leftarrow \text{KeyGen}(\text{msk}, f); \\ \text{Dec}(\text{sk}_f, \text{Enc}(\text{msk}, 1, x_1), \dots, \text{Enc}(\text{msk}, n, x_n)) = f(x_1, \dots, x_n) \end{array} \right] = 1,$$

⁶ In the private key setting, we think of mpk as some public parameters common to all algorithms.

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

Security notions. Here we recall the definitions of security for multi-input functional encryption. We give both one-time and many-time indistinguishability-based security definitions. Namely, we consider several security notions denoted xx -AD-IND and xx -SEL-IND, where: $xx \in \{one, many\}$. We also give simulation-based security definitions in the full version of the paper [3].

Definition 2 (xx-AD-IND-secure MIFE). For every multi-input functional encryption $MIFE$ for \mathcal{F} , every stateful adversary \mathcal{A} , every security parameter $\lambda \in \mathbb{N}$, and every $xx \in \{one, many\}$, we define the following experiments for $\beta \in \{0, 1\}$:

$$\begin{array}{l}
 \text{Experiment } \mathbf{xx-AD-IND}_{\beta}^{MIFE}(1^{\lambda}, \mathcal{A}): \\
 (\text{mpk}, \text{msk}) \leftarrow \text{Setup}(1^{\lambda}, \mathcal{F}_n) \\
 \alpha \leftarrow \mathcal{A}^{\text{KeyGen}(\text{msk}, \cdot), \text{Enc}(\cdot, \cdot, \cdot)}(\text{mpk}) \\
 \text{Output: } \alpha
 \end{array}$$

where Enc is an oracle that on input (i, x_i^0, x_i^1) outputs $\text{Enc}(\text{msk}, i, x_i^{\beta})$. Also, \mathcal{A} is restricted to only make queries f to $\text{KeyGen}(\text{msk}, \cdot)$ satisfying

$$f(x_1^{j_1, 0}, \dots, x_n^{j_n, 0}) = f(x_1^{j_1, 1}, \dots, x_n^{j_n, 1})$$

for all $j_1, \dots, j_n \in [Q_1] \times \dots \times [Q_n]$, where for all $i \in [n]$, Q_i denotes the number of encryption queries for input slot i . We denote by Q_f the number of key queries. Note that w.l.o.g. (as shown in [4, Lemma 3]), we can assume that for all $i \in [n]$, $Q_i > 0$. When $xx = one$, we also require that \mathcal{A} queries $\text{Enc}(i, \cdot, \cdot)$ once per slot, namely that $Q_i = 1$, for all $i \in [n]$.

A private-key multi-input functional encryption $MIFE$ for \mathcal{F} is xx -AD-IND-secure if every PPT adversary \mathcal{A} has advantage negligible in λ , where the advantage is defined as:

$$\begin{aligned}
 \text{Adv}_{MIFE}^{xx\text{-AD-IND}}(\lambda, \mathcal{A}) = \\
 \left| \Pr [\mathbf{xx-AD-IND}_0^{MIFE}(1^{\lambda}, \mathcal{A}) = 1] - \Pr [\mathbf{xx-AD-IND}_1^{MIFE}(1^{\lambda}, \mathcal{A}) = 1] \right|
 \end{aligned}$$

Remark 1 (winning condition). The winning condition may not always efficiently checkable because of the combinatorial explosion in the restrictions on the queries.

Definition 3 (xx-SEL-IND-secure MIFE). For every multi-input functional encryption $MIFE$ for \mathcal{F} , every stateful adversary \mathcal{A} , every security parameter $\lambda \in \mathbb{N}$, and every $xx \in \{one, many\}$, we define the following experiments for $\beta \in \{0, 1\}$:

Experiment $\mathbf{xx}\text{-SEL-IND}_\beta^{\text{MIFE}}(1^\lambda, \mathcal{A})$:

$\{x_i^{j,b}\}_{i \in [n], j \in [Q_i], b \in \{0,1\}} \leftarrow \mathcal{A}(1^\lambda, \mathcal{F}_n)$
 $(\text{mpk}, \text{msk}) \leftarrow \text{Setup}(1^\lambda, \mathcal{F}_n)$
 $\text{ct}_i^j := \text{Enc}(\text{msk}, x_i^{j,\beta})$
 $\alpha \leftarrow \mathcal{A}^{\text{KeyGen}(\text{msk}, \cdot)}(\text{mpk}, \{\text{ct}_i^j\}_{i \in [n], j \in [Q_i]})$
Output: α

where \mathcal{A} is restricted to only make queries f to $\text{KeyGen}(\text{msk}, \cdot)$ satisfying

$$f(x_1^{j_1,0}, \dots, x_n^{j_n,0}) = f(x_1^{j_1,1}, \dots, x_n^{j_n,1})$$

for all $j_1, \dots, j_n \in [Q_1] \times \dots \times [Q_n]$. When $xx = \text{one}$, we also require that $Q_i = 1$, for all $i \in [n]$.

A MIFE for \mathcal{F} is $xx\text{-SEL-IND}$ -secure if every PPT adversary \mathcal{A} has negligible advantage in λ , where the advantage is defined as:

$$\text{Adv}_{\text{MIFE}, \mathcal{A}}^{xx\text{-SEL-IND}}(\lambda) = |\Pr [\mathbf{xx}\text{-SEL-IND}_0^{\text{MIFE}}(1^\lambda, \mathcal{A}) = 1] - \Pr [\mathbf{xx}\text{-SEL-IND}_1^{\text{MIFE}}(1^\lambda, \mathcal{A}) = 1]|.$$

Zero vs multiple queries in the private-key setting. A nice feature enjoyed by all the schemes in Sect. 3 is that the owner of a decryption key sk_y associated with the vector $\mathbf{y} = \mathbf{y}_1 \parallel \dots \parallel \mathbf{y}_n$ does not need to know a specific value ct_i of the ciphertext vector $\text{ct} = (\text{ct}_1, \dots, \text{ct}_n)$ in order to decrypt ct if $\mathbf{y}_i = \mathbf{0}$. In other words, Q_i can be 0 whenever $\mathbf{y}_i = \mathbf{0}$. In this case, the adversary is only allowed to obtain a secret key sk_y for a vector \mathbf{y} satisfying the condition

$$\sum_{i \in I} \langle \mathbf{x}_i^{j,0}, \mathbf{y}_i \rangle = \sum_{i \in I} \langle \mathbf{x}_i^{j,1}, \mathbf{y}_i \rangle,$$

for all queries $j \in [Q_i]$, where $I \subseteq [n]$ denotes the set of slots for which the adversary made at least one query to Enc , that is, for which $Q_i > 0$. Though we believe this feature can be useful in practice (for instance, if one of the encrypting parties decides to stop collaborating), certain applications may require at least one ciphertext for each encryption slot in order for decryption to be possible. In such cases, one can apply to our schemes the simple generic compiler given in [4, Lemma 3] to ensure that the set $I = [n]$, thus obtaining new schemes which leak no information in the setting where some $Q_i = 0$. For this reason, we assume without loss of generality that $Q_i > 0$ in all our security definitions and proofs.

2.2 Function-Hiding Multi-Input Functional Encryption

For function-hiding, we focus on indistinguishability security notions. This is because even single-input function-hiding inner-product encryption is known to be unrealizable in a simulation sense under standard assumptions.

Definition 4 (xx-SEL-Function-hiding MIFE). For every multi-input functional encryption \mathcal{MIFE} for \mathcal{F} , every security parameter λ , every stateful adversary \mathcal{A} , and every $xx \in \{one, many\}$, we define the following experiments for $\beta \in \{0, 1\}$:

Experiment $\mathbf{xx-SEL-FH-IND}_\beta^{\mathcal{MIFE}}(1^\lambda, \mathcal{A})$:

$$\{x_i^{j,b}\}_{i \in [n], j \in [Q_i], b \in \{0,1\}} \leftarrow \mathcal{A}(1^\lambda, \mathcal{F}_n)$$

$$\{f^{j,b}\}_{j \in [Q_f], b \in \{0,1\}} \leftarrow \mathcal{A}(1^\lambda, \mathcal{F}_n)$$

$$(\text{mpk}, \text{msk}) \leftarrow \text{Setup}(1^\lambda, \mathcal{F}_n)$$

$$\text{ct}_i^j \leftarrow \text{Enc}(\text{msk}, i, x_i^{j,\beta}) \quad \forall i \in [n], j \in [Q_i]$$

$$\text{sk}^j \leftarrow \text{KeyGen}(\text{msk}, f^{j,\beta}) \quad \forall j \in [Q_f]$$

$$\alpha \leftarrow \mathcal{A}(\text{mpk}, (\text{ct}_i^j)_{i \in [n], j \in [Q_i]}, (\text{sk}^j)_{j \in [Q_f]})$$

Output: α

where \mathcal{A} only makes Q_i selective queries of plaintext pairs $(x_i^{j_i,0}, x_i^{j_i,1})$ and Q_f selective queries of key pairs $(f^{j_f,0}, f^{j_f,1})$, that must satisfy:

$$f^{j_f,0}(x_1^{j_1,0}, \dots, x_n^{j_n,0}) = f^{j_f,1}(x_1^{j_1,1}, \dots, x_n^{j_n,1})$$

for all $j_1, \dots, j_n \in [Q_1] \times \dots \times [Q_n]$ and for all $j_f \in [Q_f]$.

A \mathcal{MIFE} is xx -SEL-FH-IND-secure if every PPT adversary \mathcal{A} has negligible advantage in λ , where the advantage is defined as:

$$\text{Adv}_{\mathcal{MIFE}, \mathcal{A}}^{xx\text{-SEL-FH-IND}}(\lambda) = \left| \Pr[\mathbf{xx-SEL-FH-IND}_0^{\mathcal{MIFE}}(1^\lambda, \mathcal{A}) = 1] - \Pr[\mathbf{xx-SEL-FH-IND}_1^{\mathcal{MIFE}}(1^\lambda, \mathcal{A}) = 1] \right|$$

Definition 5 (xx-AD-Function-hiding MIFE). For every multi-input functional encryption $\mathcal{MIFE} := (\text{Setup}, \text{Enc}, \text{KeyGen}, \text{Dec})$ for \mathcal{F} , every security parameter λ , every stateful adversary \mathcal{A} , and every $xx \in \{one, many\}$, we define the following experiments for $\beta \in \{0, 1\}$:

Experiment $\mathbf{xx-AD-FH-IND}_\beta^{\mathcal{MIFE}}(1^\lambda, \mathcal{A})$:

$$(\text{mpk}, \text{msk}) \leftarrow \text{Setup}(1^\lambda, \mathcal{F}_n)$$

$$\beta' \leftarrow \mathcal{A}^{\text{KeyGen}(\text{msk}, \cdot, \cdot), \text{Enc}(\text{msk}, \cdot, \cdot)}(\text{mpk})$$

Output: α

where Enc is an oracle that on input (i, x_i^0, x_i^1) outputs $\text{Enc}(\text{msk}, i, x_i^\beta)$ and KeyGen is an oracle that on input (f^0, f^1) outputs $\text{KeyGen}(\text{msk}, f^\beta)$. Additionally, \mathcal{A} queries must satisfy:

$$f^{j_f,0}(x_1^{j_1,0}, \dots, x_n^{j_n,0}) = f^{j_f,1}(x_1^{j_1,1}, \dots, x_n^{j_n,1})$$

for all $j_1, \dots, j_n \in [Q_1] \times \dots \times [Q_n]$ and for all $j_f \in [Q_f]$.

A MIFE is xx -AD-FH-IND-secure if every PPT adversary has negligible advantage in λ , where the advantage is defined as:

$$\text{Adv}_{\text{MIFE}, \mathcal{A}}^{xx\text{-AD-FH-IND}}(\lambda) = \left| \Pr[\text{xx-AD-FH-IND}_0^{\text{MIFE}}(1^\lambda, \mathcal{A}) = 1] - \Pr[\text{xx-AD-FH-IND}_1^{\text{MIFE}}(1^\lambda, \mathcal{A}) = 1] \right|$$

Definition 6 (Weak function hiding MIFE). Following the approach from [14], we define the notion of weak function hiding (denoted xx - yy -wFH-IND) in the multi-input case, which is as in Definitions 4 and 5, with the exception that the previous constraints on ciphertext and key challenges:

$$f^{j_f, 0}(x_1^{j_1, 0}, \dots, x_n^{j_n, 0}) = f^{j_f, 1}(x_1^{j_1, 1}, \dots, x_n^{j_n, 1}),$$

for all $j_1, \dots, j_n \in [Q_1] \times \dots \times [Q_n]$ and for all $j_f \in [Q_f]$

are extended with additional constraints to help with our hybrid proof:

$$f^{j_f, 0}(x_1^{j_1, 0}, \dots, x_n^{j_n, 0}) = f^{j_f, 0}(x_1^{j_1, 1}, \dots, x_n^{j_n, 1}) = f^{j_f, 1}(x_1^{j_1, 1}, \dots, x_n^{j_n, 1}),$$

for all $j_1, \dots, j_n \in [Q_1] \times \dots \times [Q_n]$ and for all $j_f \in [Q_f]$.

2.3 Inner-Product Functionality

In this paper we construct multi-input functional encryption schemes that support the following two variants of the multi-input inner product functionality:

Multi-Input Inner Product over \mathbb{Z}_L . This is a family of functions that is defined as $\mathcal{F}_{L,n}^m = \{f_{\mathbf{y}_1, \dots, \mathbf{y}_n} : (\mathbb{Z}_L^m)^n \rightarrow \mathbb{Z}_L, \text{ for } \mathbf{y}_i \in \mathbb{Z}_L^m\}$ where

$$f_{\mathbf{y}_1, \dots, \mathbf{y}_n}(\mathbf{x}_1, \dots, \mathbf{x}_n) = \sum_{i=1}^n \langle \mathbf{x}_i, \mathbf{y}_i \rangle \text{ mod } L.$$

Multi-Input Bounded-Norm Inner Product over \mathbb{Z} . This is defined as $\mathcal{F}_n^{m, X, Y} = \{f_{\mathbf{y}_1, \dots, \mathbf{y}_n} : (\mathbb{Z}^m)^n \rightarrow \mathbb{Z}\}$ where $f_{\mathbf{y}_1, \dots, \mathbf{y}_n}(\mathbf{x}_1, \dots, \mathbf{x}_n)$ is the same as above except that the result is not reduced mod L , and vectors are required to satisfy the following bounds: $\|\mathbf{x}\|_\infty < X, \|\mathbf{y}\|_\infty < Y$.

2.4 Computational Assumptions

Prime-order groups. Let GGen be a probabilistic polynomial time (PPT) algorithm that on input 1^λ returns a description $\mathcal{G} = (\mathbb{G}, p, g)$ of an cyclic group \mathcal{G} of order p for a 2λ -bit prime p , whose generator is g .

We use implicit representation of group elements as introduced in [11]. For $a \in \mathbb{Z}_p$, define $[a] = g^a \in \mathbb{G}$ as the *implicit representation* of a in \mathcal{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 \mathcal{G} :

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

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 $[a + b] \in \mathbb{G}$.

Matrix Diffie-Hellman Assumption for prime-order groups. We recall the definition of the Matrix Decision Diffie-Hellman (MDDH) Assumption [11].

Definition 7 (Matrix Distribution). *Let $k \in \mathbb{N}$. We call \mathcal{D}_k a matrix distribution if it outputs matrices in $\mathbb{Z}_p^{(k+1) \times k}$ of full rank k in polynomial time.*

W.l.o.g. we assume the first k rows of $\mathbf{A} \leftarrow_{\mathbb{R}} \mathcal{D}_k$ form an invertible matrix. The \mathcal{D}_k -Matrix Diffie-Hellman problem is to distinguish the two distributions $([\mathbf{A}], [\mathbf{A}\mathbf{w}])$ and $([\mathbf{A}], [\mathbf{u}])$ where $\mathbf{A} \leftarrow_{\mathbb{R}} \mathcal{D}_k$, $\mathbf{w} \leftarrow_{\mathbb{R}} \mathbb{Z}_p^k$ and $\mathbf{u} \leftarrow_{\mathbb{R}} \mathbb{Z}_p^\ell$.

Definition 8 (\mathcal{D}_k -Matrix Diffie-Hellman (\mathcal{D}_k -MDDH) assumption in prime-order groups). *Let \mathcal{D}_k be a matrix distribution. The \mathcal{D}_k -Matrix Diffie-Hellman (\mathcal{D}_k -MDDH) assumption holds relative to GGen if for all PPT adversaries \mathcal{A} ,*

$$\text{Adv}_{\text{GGen}, \mathcal{A}}^{\mathcal{D}_k\text{-mddh}}(\lambda) := |\Pr[\mathcal{A}(\mathcal{G}, [\mathbf{A}], [\mathbf{A}\mathbf{w}]) = 1] - \Pr[\mathcal{A}(\mathcal{G}, [\mathbf{A}], [\mathbf{u}]) = 1]| = \text{negl}(\lambda),$$

where probabilities are over $\mathcal{G} \leftarrow_{\mathbb{R}} \text{GGen}(1^\lambda)$, $\mathbf{A} \leftarrow_{\mathbb{R}} \mathcal{D}_k$, $\mathbf{w} \leftarrow_{\mathbb{R}} \mathbb{Z}_p^k$, $\mathbf{u} \leftarrow_{\mathbb{R}} \mathbb{Z}_p^{k+1}$.

Pairing groups. Let PGGen be a probabilistic polynomial time (PPT) algorithm that on input 1^λ returns a description $\mathcal{PG} = (\mathbb{G}_1, \mathbb{G}_2, q, g_1, g_2)$ of asymmetric pairing groups where $\mathbb{G}_1, \mathbb{G}_2, \mathbb{G}_T$ are cyclic group of order p for a 2λ -bit prime p , g_1 and g_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 $g_T := e(g_1, g_2)$, 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 $[a]_s = g_s^a \in \mathbb{G}_s$ as the implicit representation of a in \mathbb{G}_s . Given $[a]_1, [a]_2$, one can efficiently compute $[ab]_T$ using the pairing e . For two matrices \mathbf{A}, \mathbf{B} with matching dimensions define $e([\mathbf{A}]_1, [\mathbf{B}]_2) := [\mathbf{AB}]_T \in \mathbb{G}_T$.

We define the \mathcal{D}_k -MDDH assumption in pairing groups similarly than in prime-order groups (see Definition 8).

Definition 9 (\mathcal{D}_k -MDDH assumption in pairing groups). *Let \mathcal{D}_k be a matrix distribution. The \mathcal{D}_k -MDDH assumption holds relative to PGGen in \mathbb{G}_s , for $s \in \{1, 2, T\}$, if for all PPT adversaries \mathcal{A} , the following is $\text{negl}(\lambda)$:*

$$\text{Adv}_{\mathbb{G}_s, \mathcal{A}}^{\mathcal{D}_k\text{-mddh}}(\lambda) := |\Pr[\mathcal{A}(\mathcal{PG}, [\mathbf{A}]_s, [\mathbf{A}\mathbf{w}]_s) = 1] - \Pr[\mathcal{A}(\mathcal{PG}, [\mathbf{A}]_s, [\mathbf{u}]_s) = 1]|$$

where probabilities are over $\mathcal{PG} \leftarrow_{\mathbb{R}} \text{PGGen}(1^\lambda)$, $\mathbf{A} \leftarrow_{\mathbb{R}} \mathcal{D}_k$, $\mathbf{w} \leftarrow_{\mathbb{R}} \mathbb{Z}_p^k$, $\mathbf{u} \leftarrow_{\mathbb{R}} \mathbb{Z}_p^{k+1}$.

Next, we recall a result on the uniform distribution over full-rank matrices:

Definition 10 (Uniform distribution). Let $\ell, k \in \mathbb{N}$, with $\ell > k$. We denote by $\mathcal{U}_{\ell,k}$ the uniform distribution over all full-rank $\ell \times k$ matrices over \mathbb{Z}_p .

Among all possible matrix distributions \mathcal{D}_k , the uniform matrix distribution $\mathcal{U}_{\ell,k}$ is the hardest possible instance, so in particular $k\text{-Lin} \Rightarrow \mathcal{U}_k\text{-MDDH}$, as stated in Lemma 1.

Lemma 1 ($\mathcal{D}_k\text{-MDDH} \Rightarrow \mathcal{U}_{\ell,k}\text{-MDDH}$, [11]). Let $\ell, k \in \mathbb{N}$ and \mathcal{D}_k a matrix distribution. For any PPT adversary \mathcal{A} , there exists a PPT \mathcal{B} such that

$$\text{Adv}_{\mathbb{G}_s, \mathcal{A}}^{\mathcal{U}_{\ell,k}\text{-mddh}}(\lambda) \leq \text{Adv}_{\mathbb{G}_s, \mathcal{B}}^{\mathcal{D}_k\text{-mddh}}(\lambda).$$

3 From Single to Multi-Input FE for Inner Product

In this section, we give a generic construction of MIFE for inner product from any single-input FE (Setup, Enc, KeyGen, Dec) for the same functionality. More precisely, we show two transformations: the first one addresses FE schemes that compute the inner product functionality over a finite ring \mathbb{Z}_L for some integer L , while the second transformation addresses FE schemes for bounded-norm inner product. The two transformations are almost the same, and the only difference is that in the case of bounded-norm inner product, we require additional structural properties on the single-input FE. Yet we stress that these properties are satisfied by all existing constructions. Both our constructions rely on a simple MIFE scheme that is one-AD-IND secure unconditionally. In particular, our constructions show how to use single-input FE in order to bootstrap the information-theoretic MIFE from one-time to many-time security.

$\text{Setup}^{\text{ot}}(1^\lambda, \mathcal{F}_{L,n}^m):$ For all $i \in [n]$, $\mathbf{u}_i \leftarrow_{\mathbb{R}} \mathbb{Z}_L^m$ Return $\mathbf{u} = \{\mathbf{u}_i\}_{i \in [n]}$	$\text{KeyGen}^{\text{ot}}(\mathbf{u}, \mathbf{y}_1 \ \dots \ \mathbf{y}_n):$ Return $z := \sum_{i \in [n]} \langle \mathbf{u}_i, \mathbf{y}_i \rangle \text{ mod } L$
$\text{Enc}^{\text{ot}}(\mathbf{u}, i, \mathbf{x}_i):$ Return $\mathbf{x}_i + \mathbf{u}_i \text{ mod } L$	$\text{Dec}^{\text{ot}}(z, \text{ct}_1, \dots, \text{ct}_n):$ Return $\sum_{i=1}^n \langle \text{ct}_i, \mathbf{y}_i \rangle - z \text{ mod } L$

Fig. 1. Private-key, information theoretically secure, multi-input FE scheme $\mathcal{MIFE}^{\text{ot}} = (\text{Setup}^{\text{ot}}, \text{Enc}^{\text{ot}}, \text{KeyGen}^{\text{ot}}, \text{Dec}^{\text{ot}})$ for the class $\mathcal{F}_{L,n}^m$.

3.1 Information-Theoretic MIFE with One-Time Security

Here we present the multi-input scheme $\mathcal{MIFE}^{\text{ot}}$ for the class $\mathcal{F}_{L,n}^m$, and we prove its one-AD-IND security. The scheme is described in Fig. 1.

Theorem 1. The MIFE described in Fig. 1 is one-AD-IND secure. Namely, for any adversary \mathcal{A} , $\text{Adv}_{\mathcal{MIFE}, \mathcal{A}}^{\text{one-AD-IND}}(\lambda) = 0$.

Proof overview. The proof of Theorem 1 has two main steps. First, we use the fact that any adaptive distinguisher against $\mathcal{MIFE}^{\text{ot}}$ with advantage ε can be transformed into a selective distinguisher with advantage $\varepsilon/|X|^2$ by randomly guessing the two challenge input vectors, where $|X|$ is the size of the input space ($|X| = L^{nm}$ in our case). Then, in a second step, we show that any selective distinguisher against $\mathcal{MIFE}^{\text{ot}}$ has advantage 0 since $\mathcal{MIFE}^{\text{ot}}$ behaves as the FE equivalent of the one-time pad. Hence, it follows that any adaptive distinguisher must also have advantage 0.

Proof. Let \mathcal{A} be an adversary against the one-AD-IND security of the MIFE. First, we use a complexity leveraging argument to build an adversary \mathcal{B} such that:

$$\text{Adv}_{\mathcal{MIFE}, \mathcal{A}}^{\text{one-AD-IND}}(\lambda) \leq L^{-2nm} \cdot \text{Adv}_{\mathcal{MIFE}, \mathcal{B}}^{\text{one-SEL-IND}}(\lambda).$$

The adversary \mathcal{B} simply guesses the challenge $\{\mathbf{x}_i^b\}_{i \in [n], b \in \{0,1\}}$ in advance, then simulates \mathcal{A} 's experiment using its own selective experiment. When \mathcal{B} receives \mathcal{A} 's challenge, it checks if the guess was successful (call E that event): if it was, it continues simulating \mathcal{A} 's experiment, otherwise, it returns 0. When the guess is successful, \mathcal{B} perfectly simulate \mathcal{A} 's view. Since event E happens with probability exactly L^{-2nm} , and is independent of the adversary \mathcal{A} 's view, we obtain $\text{Adv}_{\mathcal{MIFE}, \mathcal{A}}^{\text{one-AD-IND}}(\lambda) \leq L^{-2nm} \cdot \text{Adv}_{\mathcal{MIFE}, \mathcal{B}}^{\text{one-SEL-IND}}(\lambda)$.

It remains to prove that the MIFE presented in Fig. 1 satisfies perfect one-SEL-IND security, namely, for any adversary \mathcal{B} , $\text{Adv}_{\mathcal{MIFE}, \mathcal{B}}^{\text{one-SEL-IND}}(\lambda) = 0$. To do so, we introduce hybrid games $H_\beta(1^\lambda, \mathcal{B})$ described in Fig. 2. We prove that for all $\beta \in \{0, 1\}$, $H_\beta(1^\lambda, \mathcal{B})$ is identical to the experiment $\text{one-SEL-IND}_{\mathcal{B}}^{\mathcal{MIFE}}(1^\lambda, \mathcal{B})$. This can be seen using the fact that for all $\{\mathbf{x}_i^\beta \in \mathbb{Z}^m\}_{i \in [n]}$, the following distributions are identical: $\{\mathbf{u}_i \bmod L\}_{i \in [n]}$ and $\{\mathbf{u}_i - \mathbf{x}_i^\beta \bmod L\}_{i \in [n]}$, with $\mathbf{u}_i \leftarrow_{\text{R}} \mathbb{Z}_L^m$. Recall that here $i \in [n]$ is an index for input slots. Note that the independence of the \mathbf{x}_i^β from the \mathbf{u}_i is only true in the selective security game. Finally, we show that \mathcal{B} 's view in $H_\beta(1^\lambda, \mathcal{B})$ is independent of β . Indeed, the only information about β that leaks in this experiment is $\sum_i \langle \mathbf{x}_i^\beta, \mathbf{y}_i \rangle$, which is independent of β by definition of the security game. \square

<p>HYB$_{\beta}(1^\lambda, \mathcal{B})$:</p> <p>$\{\mathbf{x}_i^b\}_{i \in [n], b \in \{0,1\}} \leftarrow \mathcal{B}(1^\lambda, \mathcal{F}_{L,n}^m)$</p> <p>For all $i \in [n]$,</p> <p style="padding-left: 20px;">$\mathbf{u}_i \leftarrow_{\text{R}} \mathbb{Z}_L^m$; $\text{ct}_i \leftarrow \mathbf{u}_i$</p> <p>$\alpha \leftarrow \mathcal{B}^{\mathcal{O}_H(\cdot)}(\{\text{ct}_i\}_{i \in [n]})$</p> <p>Output α</p>	<p>$\mathcal{O}_H(\mathbf{y})$:</p> <p>For all $i \in [n]$,</p> <p style="padding-left: 20px;">$z := \sum_{i \in [n]} \langle \mathbf{u}_i, \mathbf{y}_i \rangle - \langle \mathbf{x}_i^\beta, \mathbf{y}_i \rangle \bmod L$</p> <p>Return z</p>
---	---

Fig. 2. Experiments for the proof of Theorem 1.

Remark 2 (one-SEL-SIM security). As a result of independent interest, in the full version of the paper [3] we show that the MIFE presented in Fig. 1 satisfies perfect one-SEL-SIM security, which implies perfect one-SEL-IND (which itself implies perfect one-AD-IND security via complexity leveraging, as shown in the proof above).

Remark 3 (Linear homomorphism). We use the fact that Enc^{ot} is linearly homomorphic, that is, for all input slots $i \in [n]$, $\mathbf{x}_i, \mathbf{x}'_i \in \mathbb{Z}_p^m$, $\mathbf{u} \leftarrow \text{Setup}^{\text{ot}}(1^\lambda, \mathcal{F}_{L,n}^m)$, $\text{Enc}^{\text{ot}}(\mathbf{u}, i, \mathbf{x}_i) + \mathbf{x}'_i \bmod L = \text{Enc}^{\text{ot}}(\mathbf{u}, i, \mathbf{x}_i + \mathbf{x}'_i)$. This property will be used when using the one-time scheme $\mathcal{MIFE}^{\text{ot}}$ from Fig. 1 as a building block to obtain a full-fledged many-AD-IND MIFE.

3.2 Our Transformation for Inner Product over \mathbb{Z}_L

We present our multi-input scheme \mathcal{MIFE} for the class $\mathcal{F}_{L,n}^m$ in Fig. 3. The construction relies on the one-time scheme $\mathcal{MIFE}^{\text{ot}}$ of Fig. 1, and any single-input FE for the class $\mathcal{F}_{L,1}^m$.

<p><u>Setup'(1^λ, $\mathcal{F}_{L,n}^m$):</u> $\mathbf{u} \leftarrow \text{Setup}^{\text{ot}}(1^\lambda, \mathcal{F}_{L,n}^m)$, for all $i \in [n]$, $(\text{mpk}_i, \text{msk}_i) \leftarrow \text{Setup}(1^\lambda, \mathcal{F}_{L,1}^m)$ $(\text{mpk}, \text{msk}) := (\{\text{mpk}_i\}_{i \in [n]}, (\{\text{msk}_i\}_{i \in [n]}, \mathbf{u}))$ Return (mpk, msk)</p>
<p><u>Enc'(msk, i, \mathbf{x}_i):</u> $\mathbf{w}_i := \text{Enc}^{\text{ot}}(\mathbf{u}, i, \mathbf{x}_i)$ Return $\text{Enc}(\text{msk}_i, \mathbf{w}_i)$</p>
<p><u>KeyGen'(msk, $\mathbf{y}_1 \parallel \dots \parallel \mathbf{y}_n$):</u> For all $i \in [n]$, $\text{sk}_i \leftarrow \text{KeyGen}(\text{msk}_i, \mathbf{y}_i)$, $z := \text{KeyGen}^{\text{ot}}(\mathbf{u}, \mathbf{y}_1 \parallel \dots \parallel \mathbf{y}_n)$ $\text{sk}_{\mathbf{y}_1 \parallel \dots \parallel \mathbf{y}_n} := (\{\text{sk}_i\}_{i \in [n]}, z)$ Return $\text{sk}_{\mathbf{y}_1 \parallel \dots \parallel \mathbf{y}_n}$</p>
<p><u>Dec'(($\{\text{sk}_i\}_{i \in [n]}$), z, $\text{ct}_1, \dots, \text{ct}_n$):</u> For all $i \in [n]$, $D_i \leftarrow \text{Dec}(\text{sk}_i, \text{ct}_i)$ Return $\sum_{i \in [n]} D_i - z \bmod L$</p>

Fig. 3. Private-key multi-input FE scheme $\mathcal{MIFE} := (\text{Setup}', \text{Enc}', \text{KeyGen}', \text{Dec}')$ for the class $\mathcal{F}_{L,n}^m$ from a public-key single-input FE $\mathcal{FE} := (\text{Setup}, \text{Enc}, \text{KeyGen}, \text{Dec})$ for the class $\mathcal{F}_{L,1}^m$, and one-time multi-input FE $\mathcal{MIFE}^{\text{ot}} = (\text{Setup}^{\text{ot}}, \text{Enc}^{\text{ot}}, \text{KeyGen}^{\text{ot}}, \text{Dec}^{\text{ot}})$ for the class $\mathcal{F}_{L,n}^m$.

The correctness of \mathcal{MIFE} follows from the correctness properties of the single-input scheme \mathcal{FE} and the multi-input scheme $\mathcal{MIFE}^{\text{ot}}$. Indeed, correctness of the former implies that, for all input slots $i \in [n]$, $D_i = \langle \mathbf{w}_i, \mathbf{y}_i \rangle$

mod L , while correctness of $\mathcal{MIF}\mathcal{E}^{\text{ot}}$ implies that $\sum_{i \in [n]} D_i - z = \text{Dec}^{\text{ot}}(z, \mathbf{w}_1, \dots, \mathbf{w}_n) = \sum_{i \in [n]} \langle \mathbf{x}_i, \mathbf{y}_i \rangle \text{ mod } L$.

For the security we state the following theorem:

Theorem 2. *If the single-input FE, $\mathcal{F}\mathcal{E}$ is many-AD-IND-secure, and the multi-input scheme $\mathcal{MIF}\mathcal{E}^{\text{ot}}$ is one-AD-IND-secure, then the multi-input FE, $\mathcal{MIF}\mathcal{E}$, described in Fig. 3, is many-AD-IND-secure.*

Since the proof of the above theorem is almost the same as the one for the case of bounded-norm inner product, we only provide an overview here, and defer to the proof of Theorem 3 for further details.

Proof overview. Here, for any input slot $i \in [n]$, we denote by $(\mathbf{x}_i^{j,0}, \mathbf{x}_i^{j,1})$ the j 'th query to $\text{Enc}(i, \cdot, \cdot)$, for any $j \in [Q_i]$, where Q_i is the total number of queries to $\text{Enc}(i, \cdot, \cdot)$.

The proof is in two main steps. First, we switch encryptions of $\mathbf{x}_1^{1,0}, \dots, \mathbf{x}_n^{1,0}$ to those of $\mathbf{x}_1^{1,1}, \dots, \mathbf{x}_n^{1,1}$, using the one-AD-IND security of $\mathcal{MIF}\mathcal{E}^{\text{ot}}$. For the remaining ciphertexts, we switch from an encryption of $\mathbf{x}_i^{j,0} = (\mathbf{x}_i^{j,0} - \mathbf{x}_i^{1,0}) + \mathbf{x}_i^{1,0}$ to that of $(\mathbf{x}_i^{j,0} - \mathbf{x}_i^{1,0}) + \mathbf{x}_i^{1,1}$. In this step we use the fact that one can compute an encryption of $\text{Enc}^{\text{ot}}(\mathbf{u}, i, (\mathbf{x}_i^{j,0} - \mathbf{x}_i^{1,0}) + \mathbf{x}_i^{1,0})$ from an encryption $\text{Enc}^{\text{ot}}(\mathbf{u}, i, \mathbf{x}_i^{1,0})$, because the encryption algorithm Enc^{ot} of $\mathcal{MIF}\mathcal{E}^{\text{ot}}$ is linearly homomorphic (see Remark 3). Finally, we apply a hybrid argument across the slots to switch from encryptions of

$$(\mathbf{x}_i^{2,0} - \mathbf{x}_i^{1,0}) + \mathbf{x}_i^{1,1}, \dots, (\mathbf{x}_i^{Q_i,0} - \mathbf{x}_i^{1,0}) + \mathbf{x}_i^{1,1}$$

to those of

$$(\mathbf{x}_i^{2,1} - \mathbf{x}_i^{1,1}) + \mathbf{x}_i^{1,1}, \dots, (\mathbf{x}_i^{Q_i,1} - \mathbf{x}_i^{1,1}) + \mathbf{x}_i^{1,1},$$

using the many-AD-IND security of $\mathcal{F}\mathcal{E}$.

Instantiations. The construction in Fig. 3 can be instantiated using the single-input FE schemes of Agrawal et al. [5] that are many-AD-IND-secure and allow for computing inner products over a finite ring. Specifically, we obtain:

- A MIFE for inner product over \mathbb{Z}_p for a prime p , based on the LWE assumption. This is obtained by using the LWE-based scheme of Agrawal et al. [5, Sect. 4.2].
- A MIFE for inner product over \mathbb{Z}_N where N is an RSA modulus, based on the Composite Residuosity assumption. This is obtained by using the Paillier-based scheme of Agrawal et al. [5, Sect. 5.2].

We note that since both these schemes in [5] have a stateful key generation, our MIFE inherits this stateful property. Stateless MIFE instantiations are obtained from the transformation in the next section.

3.3 Our Transformation for Inner Product over \mathbb{Z}

Here we present our transformation for the case of bounded-norm inner product. In particular, in Fig. 4 we present a multi-input scheme \mathcal{MIFE} for the class $\mathcal{F}_n^{m,X,Y}$ from the one-time scheme $\mathcal{MIFE}^{\text{ot}}$ of Fig. 1, and a (single-input) scheme \mathcal{FE} for the class $\mathcal{F}_1^{m,3X,Y}$.⁷ For our transformation to work, we require \mathcal{FE} to satisfy two properties. The first one, that we call *two-step decryption*, intuitively says that the \mathcal{FE} decryption algorithm works in two steps: the first step uses the secret key to output an encoding of the result, while the second step returns the actual result $\langle \mathbf{x}, \mathbf{y} \rangle$ provided that the bounds $\|\mathbf{x}\|_\infty < X, \|\mathbf{y}\|_\infty < Y$ hold. The second property informally says that the \mathcal{FE} encryption algorithm is additively homomorphic.

We note that the two-step property also says that the encryption algorithm accepts inputs \mathbf{x} such that $\|\mathbf{x}\|_\infty > X$, yet correctness is guaranteed as long as the encrypted inputs are within the bound at the moment of invoking the second step of decryption.

Two-step decryption is formally defined as follows.

Property 1 (Two-step decryption). An FE scheme $\mathcal{FE} = (\text{Setup}, \text{Enc}, \text{KeyGen}, \text{Dec})$ satisfies *two-step decryption* if it admits PPT algorithms $\text{Setup}^*, \text{Dec}_1, \text{Dec}_2$ and an encoding function \mathcal{E} such that:

1. For all $\lambda, m, n, X, Y \in \mathbb{N}$, $\text{Setup}^*(1^\lambda, \mathcal{F}_1^{m,X,Y}, 1^n)$ outputs (msk, mpk) where mpk includes a bound $B \in \mathbb{N}$, and the description of a group \mathbb{G} (with group law \circ) of order $L > n \cdot m \cdot X \cdot Y$, which defines the encoding function $\mathcal{E} : \mathbb{Z}_L \times \mathbb{Z} \rightarrow \mathbb{G}$.
2. For all $(\text{msk}, \text{mpk}) \leftarrow \text{Setup}^*(1^\lambda, \mathcal{F}_1^{m,X,Y}, 1^n)$, $\mathbf{x} \in \mathbb{Z}^m$, $\text{ct} \leftarrow \text{Enc}(\text{msk}, \mathbf{x})$, $\mathbf{y} \in \mathbb{Z}^m$, and $\text{sk} \leftarrow \text{KeyGen}(\text{msk}, \mathbf{y})$, we have

$$\text{Dec}_1(\text{ct}, \text{sk}) = \mathcal{E}(\langle \mathbf{x}, \mathbf{y} \rangle \bmod L, \text{noise}),$$

for some $\text{noise} \in \mathbb{N}$ that depends on ct and sk . Furthermore, it holds that for all $\mathbf{x}, \mathbf{y} \in \mathbb{Z}^m$, $\Pr[\text{noise} < B] = 1 - \text{negl}(\lambda)$, where the probability is taken over the random coins of $\text{Setup}^*, \text{Enc}$ and KeyGen . Note that there is no restriction on the norm of $\langle \mathbf{x}, \mathbf{y} \rangle$ here, and that we are assuming that Enc accepts inputs \mathbf{x} whose norm may be larger than the bound.

3. Given any $\gamma \in \mathbb{Z}_L$, and mpk , one can efficiently compute $\mathcal{E}(\gamma, 0)$.
4. The encoding \mathcal{E} is linear, that is: for all $\gamma, \gamma' \in \mathbb{Z}_L, \text{noise}, \text{noise}' \in \mathbb{Z}$, we have

$$\mathcal{E}(\gamma, \text{noise}) \circ \mathcal{E}(\gamma', \text{noise}') = \mathcal{E}(\gamma + \gamma' \bmod L, \text{noise} + \text{noise}').$$

5. For all $\gamma < n \cdot m \cdot X \cdot Y$, and $\text{noise} < n \cdot B$, $\text{Dec}_2(\mathcal{E}(\gamma, \text{noise})) = \gamma$.

The second property is as follows.

Property 2 (Linear encryption). For any FE scheme $\mathcal{FE} = (\text{Setup}, \text{Enc}, \text{KeyGen}, \text{Dec})$ satisfying the two-step property, we define the following additional property.

⁷ The reason why we need $3X$ instead of X is due to maintain a correct distribution of the inputs in the security proof.

There exists a deterministic algorithm Add that takes as input a ciphertext and a message, such that for all $\mathbf{x}, \mathbf{x}' \in \mathbb{Z}^m$, the following are identically distributed:

$$\text{Add}(\text{Enc}(\text{msk}, \mathbf{x}), \mathbf{x}'), \text{ and } \text{Enc}(\text{msk}, (\mathbf{x} + \mathbf{x}' \bmod L)).$$

Note that the value $L \in \mathbb{N}$ is defined as part of the output of the algorithm Setup^* (see the two-step property above). We later use a single input FE with this property as a building block for a multi-input FE (see Fig. 4); this property however is only used in the security proof of our transformation.

Instantiations. It is not hard to check that these two properties are satisfied by known functional encryption schemes for (bounded-norm) inner product. In particular, in the full version of the paper we show that this is satisfied by the many-AD-IND secure FE schemes of Agrawal et al. [5].⁸ This allows us to obtain MIFE schemes for bounded-norm inner product based on a variety of assumptions such as plain DDH, Decisional Composite Residuosity, and LWE. In addition to obtaining the first schemes without the need of pairing groups, we also obtain schemes where decryption works efficiently even for large outputs. This stands in contrast to the previous result [4], where decryption requires to extract discrete logarithms.

Correctness. The correctness of the scheme \mathcal{MIFE} follows from (i) the correctness and Property 1 (two-step decryption) of the single-input scheme, and (ii) from the correctness of $\mathcal{MIFE}^{\text{ot}}$ and the linear property of its decryption algorithm Dec^{ot} .

More precisely, consider any vector $\mathbf{x} := (\mathbf{x}_1 \parallel \dots \parallel \mathbf{x}_n) \in (\mathbb{Z}^m)^n$, $\mathbf{y} \in \mathbb{Z}^{mn}$, such that $\|\mathbf{x}\|_\infty < X$, $\|\mathbf{y}\|_\infty < Y$, and let $(\text{mpk}, \text{msk}) \leftarrow \text{Setup}'(1^\lambda, \mathcal{F}_n^{m, X, Y})$, $\text{sk}_\mathbf{y} \leftarrow \text{KeyGen}'(\text{msk}, \mathbf{y})$, and $\text{ct}_i \leftarrow \text{Enc}'(\text{msk}, i, \mathbf{x}_i)$ for all $i \in [n]$.

By (2) of Property 1, the decryption algorithm $\text{Dec}'(\text{sk}_\mathbf{y}, \text{ct}_1, \dots, \text{ct}_n)$ computes $\mathcal{E}(\langle \mathbf{w}_i, \mathbf{y}_i \rangle \bmod L, \text{noise}_i) \leftarrow \text{Dec}_1(\text{sk}_i, \text{ct}_i)$ where for all $i \in [n]$, $\text{noise}_i < B$, with probability $1 - \text{negl}(\lambda)$.

By (4) of Property 1 (linearity of \mathcal{E}), and the correctness of $\mathcal{MIFE}^{\text{ot}}$ we have:

$$\begin{aligned} & \mathcal{E}(\langle \mathbf{w}_1, \mathbf{y}_1 \rangle \bmod L, \text{noise}_1) \circ \dots \circ \mathcal{E}(\langle \mathbf{w}_n, \mathbf{y}_n \rangle \bmod L, \text{noise}_n) \circ \mathcal{E}(-z, 0) \\ &= \mathcal{E} \left(\text{Dec}^{\text{ot}}(z, \mathbf{w}_1, \dots, \mathbf{w}_n), \sum_{i \in [n]} \text{noise}_i \right) = \mathcal{E} \left(\langle \mathbf{x}, \mathbf{y} \rangle \bmod L, \sum_{i \in [n]} \text{noise}_i \right). \end{aligned}$$

Since $\langle \mathbf{x}, \mathbf{y} \rangle < n \cdot m \cdot X \cdot Y < L$ and $\sum_{i \in [n]} \text{noise}_i < n \cdot B$, we have

$$\text{Dec}_2(\mathcal{E}(\langle \mathbf{x}, \mathbf{y} \rangle \bmod L, \sum_{i \in [n]} \text{noise}_i)) = \langle \mathbf{x}, \mathbf{y} \rangle,$$

by (5) of Property 1.

⁸ While in [5] the FE schemes are proven only one-AD-IND secure (i.e., for adversaries making a single encryption query), note that these are *public-key* schemes and thus many-AD-IND security can be obtained via a standard hybrid argument from one-AD-IND security.

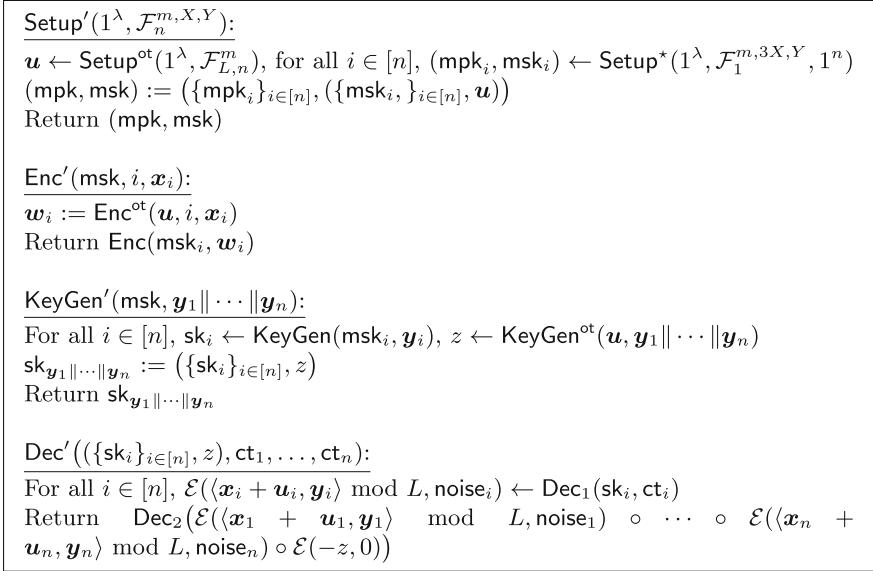


Fig. 4. Private-key multi-input FE scheme $\mathcal{MIFE} = (\text{Setup}', \text{Enc}', \text{KeyGen}', \text{Dec}')$ for the class $\mathcal{F}_n^{m,X,Y}$ from public-key single-input FE scheme $\mathcal{FE} = (\text{Setup}, \text{Enc}, \text{KeyGen}, \text{Dec})$ for the class $\mathcal{F}_1^{m,X,Y}$ and one-time multi-input FE $\mathcal{MIFE}^{\text{ot}} = (\text{Setup}^{\text{ot}}, \text{Enc}^{\text{ot}}, \text{KeyGen}^{\text{ot}}, \text{Dec}^{\text{ot}})$.

Proof of Security. In the following theorem we show that our construction is a many-AD-IND-secure MIFE, assuming that the underlying single-input FE scheme is many-AD-IND-secure, and the scheme $\mathcal{MIFE}^{\text{ot}}$ is one-AD-IND secure.

Theorem 3. *Assume that the single-input FE is many-AD-IND-secure and the multi-input FE $\mathcal{MIFE}^{\text{ot}}$ is one-AD-IND-secure. Then the multi-input FE \mathcal{MIFE} in Fig. 4 is many-AD-IND-secure. Namely, for any PPT adversary \mathcal{A} , there exist PPT adversaries \mathcal{B} and \mathcal{B}' such that*

$$\text{Adv}_{\mathcal{MIFE}, \mathcal{A}}^{\text{many-AD-IND}}(\lambda) \leq \text{Adv}_{\mathcal{MIFE}^{\text{ot}}, \mathcal{B}}^{\text{one-AD-IND}}(\lambda) + n \cdot \text{Adv}_{\mathcal{FE}, \mathcal{B}'}^{\text{many-AD-IND}}(\lambda).$$

Proof of Theorem 3. The proof proceeds by a sequence of games where G_0 is the **many-AD-IND** $_0^{\mathcal{MIFE}}(1^\lambda, \mathcal{A})$ experiment. A formal description of all the experiments used in this proof is given in Fig. 6, and a high-level summary is provided in Fig. 5. For any game G_i , we denote by $\text{Adv}_i(\mathcal{A})$ the advantage of \mathcal{A} in G_i , that is, $\Pr[G_i(1^\lambda, \mathcal{A}) = 1]$, where the probability is taken over the random coins of G_i and \mathcal{A} . In what follows we adopt the same notation from [4] for queried plaintexts, namely $(\mathbf{x}_i^{j,0}, \mathbf{x}_i^{j,1})$ denotes the j -th encryption query on the i -th slot.

Game G_1 : Here we change the way the challenge ciphertexts are created. In particular, for all slots and all queries simultaneously, we switch from $\text{Enc}'(\text{msk}, i, \mathbf{x}_i^{j,0} - \mathbf{x}_i^{1,0} + \mathbf{x}_i^{1,0})$ to $\text{Enc}'(\text{msk}, i, \mathbf{x}_i^{j,0} - \mathbf{x}_i^{1,0} + \mathbf{x}_i^{1,1})$.

Game	ct_i^j	justification/remark
G_0	$Enc'(msk', i, \mathbf{x}_i^{j,0} - \mathbf{x}_i^{1,0} + \mathbf{x}_i^{1,0})$	
G_1	$Enc'(msk', i, \mathbf{x}_i^{j,0} - \mathbf{x}_i^{1,0} + \mathbf{x}_i^{1,1})$	one-AD-IND of \mathcal{MIFE}^{ot} , Lemma 2
$G_{1..l}$	$Enc'(msk', i, \mathbf{x}_i^{j,1} - \mathbf{x}_i^{1,1} + \mathbf{x}_i^{1,1})$, for $i \leq l$ $Enc'(msk', i, \mathbf{x}_i^{j,0} - \mathbf{x}_i^{1,0} + \mathbf{x}_i^{1,1})$, for $i > l$	many-AD-IND of \mathcal{FE} , Lemma 3
G_2	$Enc'(msk', i, \mathbf{x}_i^{j,1})$	$G_2 = G_{1..n}$

Fig. 5. An overview of the games used in the proof of Theorem 3.

G_1 can be proved indistinguishable from G_0 by relying on the one-time security of the multi-input scheme. More formally,

Lemma 2. *There exists a PPT adversary \mathcal{B}_1 against the one-AD-IND security of \mathcal{MIFE}^{ot} scheme such that*

$$|\mathbf{Adv}_0(\mathcal{A}) - \mathbf{Adv}_1(\mathcal{A})| \leq \mathbf{Adv}_{\mathcal{MIFE}^{ot}, \mathcal{B}_1}^{one-AD-IND}(\lambda).$$

Proof. Here we replace encryptions of $\mathbf{x}_i^{j,0} - \mathbf{x}_i^{1,0} + \mathbf{x}_i^{1,0}$ with encryptions of $\mathbf{x}_i^{j,0} - \mathbf{x}_i^{1,0} + \mathbf{x}_i^{1,1}$ in all slots simultaneously. Recall that here, j is the index of the encryption query while i is the index for the slot. The argument relies on the one-AD-IND security of the multi-input scheme \mathcal{MIFE}^{ot} and on the fact that ciphertexts produced by the latter can be used as plaintext for the underlying single input FE scheme \mathcal{FE} that we are using as additional basic building block.

More in details, we build the adversary \mathcal{B}_1 so that it simulates G_β to \mathcal{A} when interacting with experiment **one-AD-IND** $_{\beta}^{\mathcal{MIFE}}$.

Initially \mathcal{B}_1 does not receive anything, since the one-AD-IND information-theoretically secure MIFE does not have any public key. For all $i \in [n]$ it runs $(mpk_i, msk_i) \leftarrow \text{Setup}^*(1^\lambda, \mathcal{F}_1^{m, 3X, Y}, 1^n)$, and hands the public parameters to \mathcal{A} . Also, whenever \mathcal{A} queries a secret key, \mathcal{B}_1 first queries its own oracle (on the same input) to get a corresponding key z . Next, for all $i \in [n]$, it sets $sk_i \leftarrow \text{KeyGen}(msk_i, \mathbf{y}_i)$ and gives back to \mathcal{A} the secret key $sk_{\mathbf{y}_1 || \dots || \mathbf{y}_n} := (\{sk_i\}_{i \in [n]}, z)$.

When \mathcal{A} asks encryption queries, \mathcal{B}_1 proceeds as follows. For each slot i , when receiving the first query $(i, \mathbf{x}_i^{1,0}, \mathbf{x}_i^{1,1})$, it computes the challenge ciphertext, for slot i , by invoking its own encryption oracle on the same input. Calling $\mathbf{w}_i^1 := \text{Enc}^{ot}(\mathbf{u}, i, \mathbf{x}_i^{1,\beta})$ the received ciphertext, \mathcal{B}_1 computes $ct_i^1 = \text{Enc}(msk_i, \mathbf{w}_i^1) = \text{Enc}'(msk, i, \mathbf{x}_i^{1,\beta})$.

Subsequent queries, on slot i , are answered as follows. \mathcal{B}_1 produces ct_i^j (for $j > 1$) by encrypting $\mathbf{x}_i^{j,0} - \mathbf{x}_i^{1,0} + \mathbf{w}_i^1 \bmod L$, using msk_i . Note that Enc^{ot} is linearly homomorphic (see Remark 3), thus, $\mathbf{x}_i^{j,0} - \mathbf{x}_i^{1,0} + \mathbf{w}_i^1 \bmod L = \text{Enc}^{ot}(\mathbf{u}, i, \mathbf{x}_i^{1,\beta} + \mathbf{x}_i^{j,0} - \mathbf{x}_i^{1,0})$.

Finally, \mathcal{B}_1 outputs 1 iff \mathcal{A} outputs 1. One can see that \mathcal{B}_1 provides a perfect simulation to \mathcal{A} and thus:

$$|\mathbf{Adv}_0(\mathcal{A}) - \mathbf{Adv}_1(\mathcal{A})| \leq \mathbf{Adv}_{\mathcal{MIFE}, \mathcal{B}_1}^{one-AD-IND}(\lambda).$$

□

$G_0(1^\lambda, \mathcal{A}), \boxed{G_1(1^\lambda, \mathcal{A})}, \boxed{G_2(1^\lambda, \mathcal{A})}$: $(\text{mpk}, \text{msk}) \leftarrow \text{Setup}'(1^\lambda, \mathcal{F}_n^{m, X, Y})$ $\beta' \leftarrow \mathcal{A}^{\text{KeyGen}'(\text{msk}, \cdot), \text{EncO}'(\cdot, \cdot, \cdot)}(\text{mpk})$ return β' $\text{EncO}'(i, \mathbf{x}_i^{j,0}, \mathbf{x}_i^{j,1})$ $\text{ct}_i^j := \text{Enc}'(\text{msk}, i, \mathbf{x}_i^{j,0} - \mathbf{x}_i^{1,0} + \mathbf{x}_i^{1,1})$ $\boxed{\text{ct}_i^j := \text{Enc}'(\text{msk}, i, \mathbf{x}_i^{j,0} - \mathbf{x}_i^{1,0} + \mathbf{x}_i^{1,1})}$ $\boxed{\text{ct}_i^j := \text{Enc}'(\text{msk}, i, \mathbf{x}_i^{j,1} - \mathbf{x}_i^{1,1} + \mathbf{x}_i^{1,1})}$ return ct_i^j $\text{KeyGen}'(\text{msk}, \mathbf{y})$ return sk_y	$G_{1,\ell}(1^\lambda, \mathcal{A})$: $(\text{mpk}, \text{msk}) \leftarrow \text{Setup}'(1^\lambda, \mathcal{F}_n^{m, X, Y})$ $\beta' \leftarrow \mathcal{A}^{\text{KeyGen}'(\text{msk}, \cdot), \text{EncO}'(\cdot, \cdot, \cdot)}(\text{mpk})$ return β' $\text{EncO}'(i, \mathbf{x}_i^{j,0}, \mathbf{x}_i^{j,1})$ If $i \leq \ell$ return $\text{Enc}'(\text{msk}, i, \mathbf{x}_i^{j,1} - \mathbf{x}_i^{1,1} + \mathbf{x}_i^{1,1})$ If $i > \ell$ return $\text{Enc}'(\text{msk}, i, \mathbf{x}_i^{j,0} - \mathbf{x}_i^{1,0} + \mathbf{x}_i^{1,1})$ $\text{KeyGen}'(\text{msk}, \mathbf{y})$ return sk_y
---	--

Fig. 6. Experiments for the proof of Theorem 3.

Game G_2 : Here we change again the way the challenge ciphertexts are created. In particular, for all slots i and all queries j , we switch ct_i^j from $\text{Enc}'(\text{msk}, i, \mathbf{x}_i^{j,0} - \mathbf{x}_i^{1,0} + \mathbf{x}_i^{1,1})$ to $\text{Enc}'(\text{msk}, i, \mathbf{x}_i^{j,1} - \mathbf{x}_i^{1,1} + \mathbf{x}_i^{1,1})$.

G_2 can be proved indistinguishable from G_1 via an hybrid argument over the n slots, relying on the security of the underlying single-input scheme.

By looking at the games defined in Fig. 6, one can see that

$$|\mathbf{Adv}_1(\mathcal{A}) - \mathbf{Adv}_2(\mathcal{A})| = \sum_{\ell=1}^n |\mathbf{Adv}_{1,\ell-1}(\mathcal{A}) - \mathbf{Adv}_{1,\ell}(\mathcal{A})|$$

since G_1 corresponds to game $G_{1,0}$ and whereas G_2 is identical to game $G_{1,n}$.

Therefore, for every ℓ we bound the difference between each consecutive pair of games in the following lemma:

Lemma 3. *For every $\ell \in [n]$, there exists a PPT adversary $\mathcal{B}_{1,\ell}$ against the many-AD-IND security of the single-input scheme \mathcal{FE} such that*

$$|\mathbf{Adv}_{1,\ell-1}(\mathcal{A}) - \mathbf{Adv}_{1,\ell}(\mathcal{A})| \leq \mathbf{Adv}_{\mathcal{FE}, \mathcal{B}_{1,\ell}}^{\text{many-AD-IND}}(\lambda).$$

Proof. Here, we replace encryptions of $\mathbf{x}_i^{j,0} - \mathbf{x}_i^{1,0} + \mathbf{x}_i^{1,1}$ with encryptions of $\mathbf{x}_i^{j,1} - \mathbf{x}_i^{1,1} + \mathbf{x}_i^{1,1}$ in all slots. Let us recall that j is the index of the encryption query while i is the index for the slot. The argument relies on (1) the many-AD-IND

security of the underlying single input scheme $\mathcal{FE} := (\text{Setup}, \text{KeyGen}, \text{Enc}, \text{Dec})$, (2) the fact that Enc satisfies Property 2 (linear encryption), and (3) the restrictions imposed by the security game (see [4]). As for this latter point we notice that, indeed, the security experiment restriction in the case of the inner product functionality imposes that $\langle \mathbf{x}_i^{j,0} - \mathbf{x}_i^{1,0}, \mathbf{y}_i \rangle = \langle \mathbf{x}_i^{j,1} - \mathbf{x}_i^{1,1}, \mathbf{y}_i \rangle$, for all slots $i \in [n]$. In our scheme this becomes $\langle \mathbf{x}_i^{j,0} - \mathbf{x}_i^{1,0}, \mathbf{y}_i \rangle \bmod L = \langle \mathbf{x}_i^{j,1} - \mathbf{x}_i^{1,1}, \mathbf{y}_i \rangle \bmod L$, which in turn is equivalent to

$$\langle \mathbf{x}_i^{j,0} - \mathbf{x}_i^{1,0} + \mathbf{x}_i^{1,1} + \mathbf{u}_i, \mathbf{y}_i \rangle \bmod L = \langle \mathbf{x}_i^{j,1} - \mathbf{x}_i^{1,1} + \mathbf{x}_i^{1,1} + \mathbf{u}_i, \mathbf{y}_i \rangle \bmod L.$$

More formally, we build an adversary $\mathcal{B}_{1,\ell}$ that simulates $G_{1,\ell-1+\beta}$ to \mathcal{A} when interacting with the experiment $\mathbf{many-AD-IND}_{\beta}^{\mathcal{FE}}$.

$\mathcal{B}_{1,\ell}$ starts by receiving a public key for the scheme \mathcal{FE} , which is set to be the key mpk_{ℓ} for the ℓ -th instance of \mathcal{FE} . Next, it runs $\mathbf{u} \leftarrow \text{Setup}^{\text{ot}}$, and for all $i \neq \ell$, it runs Setup^* to get $(\text{mpk}_i, \text{msk}_i)$. It gives $(\text{mpk}_1, \dots, \text{mpk}_n)$ to \mathcal{A} .

$\mathcal{B}_{1,\ell}$ answers secret key queries $\mathbf{y} = \mathbf{y}_1 || \dots || \mathbf{y}_n$ by first running $\text{sk}_i \leftarrow \text{KeyGen}(\text{msk}_i, \mathbf{y}_i)$ for $i \neq \ell$. Also it invokes its own key generation oracle on \mathbf{y}_{ℓ} , to get sk_{ℓ} . Finally, it computes $\mathbf{z} \leftarrow \text{KeyGen}^{\text{ot}}(\mathbf{u}, \mathbf{y}_1 || \dots || \mathbf{y}_n)$ (recall that $\mathcal{B}_{1,\ell}$ knows \mathbf{u}). This key material is then sent to \mathcal{A} .

$\mathcal{B}_{1,\ell}$ answers encryption queries $(i, \mathbf{x}_i^{j,0}, \mathbf{x}_i^{j,1})$ to Enc' as follows.

If $i < \ell$, it computes $\text{Enc}(\text{msk}_i, \text{Enc}^{\text{ot}}(\mathbf{u}, i, \mathbf{x}_i^{j,1}))$.

If $i > \ell$, it computes $\text{Enc}(\text{msk}_i, \text{Enc}^{\text{ot}}(\mathbf{u}, i, \mathbf{x}_i^{j,0} - \mathbf{x}_i^{1,0} + \mathbf{x}_i^{1,1}))$.

If $i = \ell$, at the j -th encryption query on slot ℓ , $\mathcal{B}_{1,\ell}$ queries its own oracle on input $(\mathbf{x}_{\ell}^{j,0} - \mathbf{x}_{\ell}^{1,0} + \mathbf{x}_{\ell}^{1,1}, \mathbf{x}_{\ell}^{j,1} - \mathbf{x}_{\ell}^{1,1} + \mathbf{x}_{\ell}^{1,1})$ (note that these vectors have norm less than $3X$, and as such, are valid input to the encryption oracle), to get back $\text{ct}_{*}^j := \text{Enc}(\text{msk}_{\ell}, \mathbf{x}_{\ell}^{j,\beta} - \mathbf{x}_{\ell}^{1,\beta} + \mathbf{x}_{\ell}^{1,1})$ from the experiment $\mathbf{many-AD-IND}_{\beta}^{\mathcal{FE}}$.

Then, $\mathcal{B}_{1,\ell}$ computes $\text{ct}_{*}^j := \text{Add}(\text{ct}_{*}^j, \mathbf{u}_{\ell})$, and sends it to \mathcal{A} .

Note that by Property 2 ct_{*}^j is identically distributed to $\text{Enc}(\text{msk}_{\ell}, \mathbf{x}_{\ell}^{j,\beta} - \mathbf{x}_{\ell}^{1,\beta} + \mathbf{x}_{\ell}^{1,1} + \mathbf{u}_{\ell} \bmod L)$, the latter being equal to $\text{Enc}(\text{msk}_{\ell}, \text{Enc}^{\text{ot}}(\mathbf{x}_{\ell}^{j,\beta} - \mathbf{x}_{\ell}^{1,\beta} + \mathbf{x}_{\ell}^{1,1}))$. Also, we remark that because $\mathcal{B}_{1,\ell}$ plays in the many-AD-IND security game, it can make several queries to its encryption oracle, which means that every ct_{*}^j obtained from the oracle is encrypted under fresh randomness r_j , i.e., $\text{ct}_{*}^j := \text{Enc}(\text{msk}_{\ell}, \mathbf{x}_{\ell}^{j,\beta} - \mathbf{x}_{\ell}^{1,\beta} + \mathbf{x}_{\ell}^{1,1}; r_j)$. Therefore, the simulated ciphertext ct_{*}^j uses randomness r_j which is independent of the randomness $r_{j'}$ used in $\text{ct}_{*}^{j'}$, for all $j \neq j'$. This means ct_{*}^j is distributed as in game $G_{1,\ell-1+\beta}$.

Finally, $\mathcal{B}_{1,\ell}$ outputs the same bit β' returned by \mathcal{A} . Thus:

$$|\text{Adv}_{1,\ell-1}(\mathcal{A}) - \text{Adv}_{1,\ell}(\mathcal{A})| \leq \text{Adv}_{\mathcal{FE}, \mathcal{B}_{1,\ell}}^{\mathbf{many-AD-IND}}(\lambda).$$

□

The proof of Theorem 3 follows by combining the bounds obtained in the previous lemmas. □

4 Function-Hiding Multi-Input FE for Inner Product

In this section, we give a function-hiding MIFE. We transform the MIFE for inner product proposed by Abdalla et al. in [4] into a function-hiding scheme, using a double layered encryption approach, similar to the one of Lin [13]. Namely, in Sect. 4.1, we give a generic construction that use any single-input FE on top of the MIFE from [4], which can prove selectively secure. Unlike the results in Sect. 3 that can be instantiated without pairings, for function-hiding we rely on pairing groups. Finally, in Sect. 4.2, we prove adaptive security, considering a specific instantiation of our construction.

Our construction. We present our function-hiding scheme \mathcal{MIFE} in Fig. 8. The construction relies on the multi-input scheme \mathcal{MIFE}' of Abdalla et al. [4] (recalled in Fig. 7), used together with any one-SEL-SIM secure, single-input FE for the functionality

$$\mathcal{F}_{\mathbb{G}_1, \mathbb{G}_2, \mathbb{G}_T}^\ell = \{f_{[\mathbf{y}]_1} : \mathbb{G}_2^\ell \rightarrow \mathbb{G}_T \text{ for } [\mathbf{y}]_1 \in \mathbb{G}_1^\ell\},$$

where

$$f_{[\mathbf{y}]_1}([\mathbf{x}]_2) := \langle \mathbf{x}, \mathbf{y} \rangle_T,$$

$\mathcal{PG} := (\mathbb{G}_1, \mathbb{G}_2, p, g_1, g_2)$ is a pairing group, and ℓ is the size of the ciphertext and secret keys in \mathcal{MIFE}' .

Concretely, we use the single-input FE from [5], generalized to the MDDH assumption, whose one-SEL-SIM security is proven in [4, 17], and whose description is recalled in the full version of the paper. Note that this single-input FE happens to be public-key, but this is not a property that we need for our overall MIFE.

Outline of the construction. Our starting point is the MIFE scheme for inner-products from [4], denoted by $\mathcal{MIFE}' := (\text{Setup}', \text{Enc}', \text{KeyGen}', \text{Dec}')$ and recalled in Fig. 7. This scheme is clearly not function-hiding, as the vector \mathbf{y} is given in the clear as part of functional secret key, in order to make decryption possible. In order to avoid the leakage of \mathbf{y} , we employ an approach similar to the one proposed in [13], which intuitively consists into adding a layer of encryption on top of the MIFE keys and ciphertexts; this is done by using a *single-input* inner product encryption scheme \mathcal{FE} . Slightly more in detail, using the \mathcal{FE} and \mathcal{MIFE}' schemes, we design our new function-hiding multi-input scheme \mathcal{MIFE} as follows.

We generate master keys $(\text{mpk}_i, \text{msk}_i) \leftarrow \mathcal{FE}.\text{Setup}(1^\lambda, \mathcal{F}_{\mathbb{G}_1, \mathbb{G}_2, \mathbb{G}_T}^\ell)$ for computing inner products on vectors of dimension ℓ , where ℓ is the size of the ciphertexts and secret keys of \mathcal{MIFE}' . To encrypt $\mathbf{x}_i \in \mathbb{Z}_p^m$ for each slot $i \in [n]$, we first compute $[\text{ct}_i^{\text{in}}]_1$ using \mathcal{MIFE}' , and then we compute $\text{ct}_i^{\text{out}} := \mathcal{FE}.\text{KeyGen}(\text{msk}_i, [\text{ct}_i^{\text{in}}]_1)$. To generate a key for $\mathbf{y} := (\mathbf{y}_1 \parallel \dots \parallel \mathbf{y}_n) \in \mathbb{Z}_p^{nm}$, we first compute the keys sk^{in} from \mathcal{MIFE}' , and then we would like to encrypt these keys using \mathcal{FE} in order to hide information about \mathbf{y} . A generic way to do it would be to set our secret key to be $\text{Enc}(\text{msk}_i, \text{sk}^{\text{in}})$, for all possible $i \in [n]$, so that we can

compute the inner product of $[\text{ct}_i^{\text{in}}]_1$ with sk^{in} for all $i \in [n]$. But that would yield keys of size $O(n^2m)$, since the key sk^{in} itself is of size $O(nm)$. We can do better, however. If we consider the specific $\mathcal{MIF}\mathcal{E}'$ scheme from [4], a secret key sk^{in} for \mathbf{y} consists of the components $(\|\text{sk}_1^{\text{in}}\| \dots \|\text{sk}_n^{\text{in}}\|_2, [z]_T)$, where each $\|\text{sk}_i^{\text{in}}\|_2$ only depends on \mathbf{y}_i and is of size $O(m)$, while $[z]_T \in \mathbb{G}_T$ does not depend on \mathbf{y} at all. Hence, we encrypt each vectors $\|\text{sk}_i^{\text{in}}\|_2$ to obtain $\text{sk}_i^{\text{out}} := \mathcal{F}\mathcal{E}.\text{Enc}(\text{mpk}_i, \|\text{sk}_i^{\text{in}}\|_2)$, which gives us a secret key $\text{sk}^{\text{out}} := (\{\text{sk}_i^{\text{out}}\}_{i \in [n]}, [z]_T)$ of total size $O(nm)$.

This way, decrypting the outer layer as $\mathcal{F}\mathcal{E}.\text{Dec}(\text{sk}_i^{\text{out}}, \text{ct}_i^{\text{out}})$ yields $[\langle \text{sk}_i^{\text{in}}, \text{ct}_i^{\text{in}} \rangle]_T$, which is what needs to be computed in the $\mathcal{MIF}\mathcal{E}'$ decryption algorithm Dec' . More precisely, correctness of $\mathcal{MIF}\mathcal{E}$ follows from the correctness of $\mathcal{MIF}\mathcal{E}'$, and the structural requirement of $\mathcal{F}\mathcal{E}.\text{Dec}$ that is used in the $\mathcal{MIF}\mathcal{E}'$ decryption algorithm, namely:

$$\begin{aligned} & \mathcal{MIF}\mathcal{E}.\text{Dec}(\{\text{sk}_i^{\text{out}}\}_{i \in [n]}, [z]_T, \{\text{ct}_i^{\text{out}}\}_{i \in [n]}) \\ &= \prod_{i=1}^n \mathcal{F}\mathcal{E}.\text{Dec}(\text{ct}_i^{\text{out}}, \text{sk}_i^{\text{out}}) / [z]_T = \prod_{i=1}^n [\langle \text{sk}_i^{\text{in}}, \text{ct}_i^{\text{in}} \rangle]_T / [z]_T \\ &= \mathcal{MIF}\mathcal{E}'.\text{Dec}(\{\|\text{sk}_i^{\text{in}}\|_2\}_{i \in [n]}, [z]_T, \{[\text{ct}_i^{\text{in}}]_1\}_{i \in [n]}). \end{aligned}$$

Definition 11 (one-SEL-SIM-secure FE). A single-input functional encryption $\mathcal{F}\mathcal{E}$ for the functionality $\mathcal{F}_{\mathbb{G}_1, \mathbb{G}_2, \mathbb{G}_T}^\ell$ is one-SEL-SIM-secure if there exist PPT simulator algorithms $(\widetilde{\text{Setup}}, \widetilde{\text{Enc}}, \widetilde{\text{KeyGen}})$ such that for every PPT (stateful) adversary \mathcal{A} and every $\lambda \in \mathbb{N}$, the following two distributions are computationally indistinguishable:

<p><u>Experiment</u> $\mathbf{REAL}_{\text{SEL}}^{\mathcal{MIF}\mathcal{E}}(1^\lambda, \mathcal{A})$:</p> <p>$\mathbf{x} \leftarrow \mathcal{A}(1^\lambda, \mathcal{F}_{\mathbb{G}_1, \mathbb{G}_2, \mathbb{G}_T}^\ell)$ $(\text{mpk}, \text{msk}) \leftarrow \widetilde{\text{Setup}}(1^\lambda, \mathcal{F}_{\mathbb{G}_1, \mathbb{G}_2, \mathbb{G}_T}^\ell)$ $\text{ct} \leftarrow \text{Enc}(\text{msk}, \mathbf{x})$ $\alpha \leftarrow \mathcal{A}^{\text{KeyGen}(\text{msk}, \cdot)}(\text{mpk}, \text{ct})$ Output: α</p>	<p><u>Experiment</u> $\mathbf{IDEAL}_{\text{SEL}}^{\mathcal{MIF}\mathcal{E}}(1^\lambda, \mathcal{A})$:</p> <p>$\mathbf{x} \leftarrow \mathcal{A}(1^\lambda, \mathcal{F}_{\mathbb{G}_1, \mathbb{G}_2, \mathbb{G}_T}^\ell)$ $(\widetilde{\text{mpk}}, \widetilde{\text{msk}}) \leftarrow \widetilde{\text{Setup}}(1^\lambda, \mathcal{F}_{\mathbb{G}_1, \mathbb{G}_2, \mathbb{G}_T}^\ell)$ $\text{ct} \leftarrow \widetilde{\text{Enc}}(\widetilde{\text{msk}})$ $\alpha \leftarrow \mathcal{A}^{\mathcal{O}(\cdot)}(\widetilde{\text{mpk}}, \text{ct})$ Output: α</p>
---	--

The oracle $\mathcal{O}(\cdot)$ in the ideal experiment above is given access to another oracle that, given $[\mathbf{y}]_1 \in \mathcal{F}_{\mathbb{G}_1, \mathbb{G}_2, \mathbb{G}_T}^\ell$, returns $[\langle \mathbf{x}, \mathbf{y} \rangle]_1$, and then $\mathcal{O}(\cdot)$ returns $\widetilde{\text{KeyGen}}(\widetilde{\text{msk}}, [\mathbf{y}]_1, [\langle \mathbf{x}, \mathbf{y} \rangle]_1)$.

For every stateful adversary \mathcal{A} , we define its advantage as

$$\begin{aligned} & \text{Adv}_{\mathcal{F}\mathcal{E}, \mathcal{A}}^{\text{one-SEL-SIM}}(\lambda) \\ &= \left| \Pr [\mathbf{REAL}_{\text{SEL}}^{\mathcal{F}\mathcal{E}}(1^\lambda, \mathcal{A}) = 1] - \Pr [\mathbf{IDEAL}_{\text{SEL}}^{\mathcal{F}\mathcal{E}}(1^\lambda, \mathcal{A}) = 1] \right|, \end{aligned}$$

and we require that for every PPT \mathcal{A} , there exists a negligible function negl such that for all $\lambda \in \mathbb{N}$, $\text{Adv}_{\mathcal{F}\mathcal{E}, \mathcal{A}}^{\text{one-SEL-SIM}}(\lambda) = \text{negl}(\lambda)$.

Multi-input scheme $\mathcal{MIFE}'[3]$	
$\text{Setup}'(1^\lambda, \mathcal{F}_n^{m,X,Y}):$ $\forall i \in [n]:$ $(\text{mpk}'_i, \text{msk}'_i) \leftarrow \mathcal{FE}'.\text{Setup}(1^\lambda, \mathcal{F}_1^{m+k,X,Y})$ $\forall i \in [n]: \mathbf{z}_i \leftarrow_{\mathbb{R}} \mathbb{Z}_q^k$ $\text{mpk}' := (\{\text{mpk}'_i\}_{i \in [n]})$ $\text{msk}' := (\{\text{msk}'_i, \mathbf{z}_i\}_{i \in [n]})$ return $(\text{mpk}', \text{msk}')$	$\text{Enc}'(\text{msk}, i, \mathbf{x}_i):$ $[\text{ct}'_i]_1 := \mathcal{FE}'.\text{Enc}(\text{mpk}'_i, \mathbf{x}_i \ \mathbf{z}_i)$ return $[\text{ct}'_i]_1$
$\text{Dec}'(\{([\text{sk}'_i]_2]_{i \in [n]}, [z]_T), \{[\text{ct}'_i]_1\}_{i \in [n]}\}):$ $\forall i \in [n]: [a_i]_T \leftarrow \mathcal{FE}'.\text{Dec}([\text{sk}'_i]_2, [\text{ct}'_i]_1)$ return the discrete log of $(\prod_{i=1}^n [a_i]_T) / [z]_T$	$\text{KeyGen}'(\text{msk}, \mathbf{y}_1 \ \dots \ \mathbf{y}_n):$ $\mathbf{r} \leftarrow_{\mathbb{R}} \mathbb{Z}_q^k$ $\forall i \in [n]:$ $\text{sk}'_i \leftarrow \mathcal{FE}'.\text{KeyGen}(\text{msk}'_i, \mathbf{y}_i \ \mathbf{r})$ $\mathbf{z} := \langle \mathbf{z}_1 + \dots + \mathbf{z}_n, \mathbf{r} \rangle$ $\text{sk}^{\text{in}} := (\{[\text{sk}'_i]_2\}_{i \in [n]}, [z]_T)$ return sk^{in}

Fig. 7. Multi-input, FE for $\mathcal{F}_n^{m,X,Y}$ from [4], whose many-SEL-IND relies on the \mathcal{D}_k -MDDH assumption. Here $\mathcal{FE}' := (\mathcal{FE}'.\text{Setup}, \mathcal{FE}'.\text{Enc}, \mathcal{FE}'.\text{KeyGen}, \mathcal{FE}'.\text{Dec})$ is a one-SEL-SIM secure, public-key, single-input FE for $\mathcal{F}_1^{m+k,X,Y}$, where k is the parameter used by the \mathcal{D}_k -MDDH assumption (concretely, $k = 1$ for SXDH, $k = 2$ for DLIN).

4.1 Proof of Selective Security

In the following theorem we state the selective security of our scheme \mathcal{MIFE} . Precisely, the theorem proves that our scheme is weakly function-hiding. We stress that this does not entail any limitation in the final result, as full-fledged function-hiding can be achieved in a generic way via a simple transformation, proposed in [14] (for single-input FE). The main idea is to work with slightly larger vectors where both input vectors \mathbf{x} and secret-key vectors \mathbf{y} are padded with zeros. In the full version of the paper we show how to do this transformation in the multi-input setting.

Theorem 4 (many-SEL-wFH-IND security). *Let \mathcal{MIFE}' be the many-SEL-IND secure multi-input FE from Fig. 7. Suppose the single-input $\mathcal{FE} := (\mathcal{FE}.\text{Setup}, \mathcal{FE}.\text{Enc}, \mathcal{FE}.\text{KeyGen}, \mathcal{FE}.\text{Dec})$ is one-SEL-SIM-secure. Then the multi-input scheme $\mathcal{MIFE} := (\text{Setup}, \text{Enc}, \text{KeyGen}, \text{Dec})$ in Fig. 8 is many-SEL-wFH-IND-secure.*

Proof Overview. The proof is done via a hybrid argument that consists of two main phases: we first switch the ciphertexts from encryptions of $\mathbf{x}_i^{j_i,0}$ to encryptions of $\mathbf{x}_i^{j_i,1}$ for all slots $i \in [n]$, and ciphertext queries $j_i \in [Q_i]$, where Q_i denotes the number of ciphertext query on the i 'th slot. This change is justified by the many-SEL-IND security of the underlying \mathcal{MIFE}' in a black box manner. In addition, this change relies on the weak-function-hiding property that imposes the constraints $\sum_{i=1}^n \langle \mathbf{x}_i^{j_i,0}, \mathbf{y}_i^{j_f,0} \rangle = \sum_{i=1}^n \langle \mathbf{x}_i^{j_i,1}, \mathbf{y}_i^{j_f,0} \rangle$, for all secret key queries $j_f \in [Q_f]$, where Q_f denotes the number of secret key queries, which thus disallow the adversary from trivially distinguishing the two games.

The second main change in the proof is to switch the decryption keys from keys corresponding to $\mathbf{y}_1^{j_1,0} \| \dots \| \mathbf{y}_n^{j_n,0}$ to keys corresponding to $\mathbf{y}_1^{j_1,1} \| \dots \| \mathbf{y}_n^{j_n,1}$ for

New function-hiding scheme \mathcal{MLFE}

Setup($1^\lambda, \mathcal{F}_n^{m,X,Y}$):
 $(\text{mpk}', \text{msk}') \leftarrow \text{Setup}'(1^\lambda, \mathcal{F}_n^{m,X,Y})$
 $\forall i \in [n] : (\text{mpk}_i, \text{msk}_i) \leftarrow \mathcal{FE}.\text{Setup}(1^\lambda, \mathcal{F}_1^{\ell,X,Y})$
 $\text{mpk} := (\{\text{mpk}_i\}_{i \in [n]}, \text{mpk}')$, $\text{msk} := (\{\text{msk}_i\}_{i \in [n]}, \text{msk}')$
 return (mpk, msk)

Enc($\text{msk}, i, \mathbf{x}_i$):
 $[\text{ct}_i^{\text{in}}]_1 := \text{Enc}'(\text{msk}', i, \mathbf{x}_i)$
 $\text{ct}_i^{\text{out}} := \mathcal{FE}.\text{KeyGen}(\text{msk}_i, [\text{ct}_i^{\text{in}}]_1)$
 return ct_i^{out}

KeyGen($\text{msk}, \mathbf{y}_1 \parallel \dots \parallel \mathbf{y}_n$):
 $(\{\text{sk}_i^{\text{in}}\}_2, [z]_T) \leftarrow \text{KeyGen}'(\text{msk}', \mathbf{y}_1 \parallel \dots \parallel \mathbf{y}_n)$
 $\forall i \in [n] : \text{sk}_i^{\text{out}} \leftarrow \mathcal{FE}.\text{Enc}(\text{msk}_i, [\text{sk}_i^{\text{in}}]_2)$
 $\text{sk}_{\mathbf{y}_1 \parallel \dots \parallel \mathbf{y}_n} := (\{\text{sk}_i^{\text{out}}\}_{i \in [n]}, [z]_T)$
 return $\text{sk}_{\mathbf{y}_1 \parallel \dots \parallel \mathbf{y}_n}$

Dec($(\{\text{sk}_i^{\text{out}}\}_{i \in [n]}, [z]_T), \{\text{ct}_i^{\text{out}}\}_{i \in [n]}$):
 $\forall i \in [n] : [a_i]_T \leftarrow \mathcal{FE}.\text{Dec}(\text{ct}_i^{\text{out}}, \text{sk}_i^{\text{out}})$
 return the discrete log of $(\prod_{i=1}^n [a_i]_T) / [z]_T$

Fig. 8. Many-SEL-wFH-IND secure, private-key, multi-input, FE for the class $\mathcal{F}_n^{m,X,Y}$. Here $\mathcal{FE} := (\mathcal{FE}.\text{Setup}, \mathcal{FE}.\text{Enc}, \mathcal{FE}.\text{KeyGen}, \mathcal{FE}.\text{Dec})$ is a one-SEL-SIM secure, single-input FE for $\mathcal{F}_1^{\ell,X,Y}$, where by ℓ we denote the output size of Enc' and KeyGen' , and $\mathcal{MLFE}' := (\text{Setup}', \text{Enc}', \text{KeyGen}', \text{Dec}')$ is the many-AD-IND secure, multi-input FE from Fig. 7.

every $j \in [Q_f]$. This in turn requires a hybrid argument over all decryption keys, changing one key at a time. To switch the ρ 'th key, we use the selective simulation security of the underlying \mathcal{FE} to embed the value $\langle \mathbf{x}_i^{j,1}, \mathbf{y}_i^{\rho,\beta} \rangle + \langle \mathbf{r}^\rho, \mathbf{z}_i \rangle$ in the ciphertexts ct_i^j , for all slots $i \in [n]$ and all $j \in [Q_i]$. Next, we use the \mathcal{D}_k -MDDH

Game	$[\text{ct}_i^{\text{in},k}]_1$	$[\text{sk}_i^{\text{in},j}]_2$	justification/remark
G_0	$\text{Enc}'(\text{msk}', i, \mathbf{x}_i^{k,0})$	$\text{KeyGen}'(\text{msk}', \mathbf{y}_1^{j,0} \parallel \dots \parallel \mathbf{y}_n^{j,0})$	many-SEL-wFH-IND ₀ security game
G_1	$\text{Enc}'(\text{msk}', i, \mathbf{x}_i^{k,1})$	$\text{KeyGen}'(\text{msk}', \mathbf{y}_1^{j,0} \parallel \dots \parallel \mathbf{y}_n^{j,0})$	many-SEL-IND of \mathcal{MLFE}'
$G_{1,\rho}$	$\text{Enc}'(\text{msk}', i, \mathbf{x}_i^{k,1})$	$\text{KeyGen}'(\text{msk}', \mathbf{y}_1^{j,1} \parallel \dots \parallel \mathbf{y}_n^{j,1})$, for $j < \rho$ $\text{KeyGen}'(\text{msk}', \mathbf{y}_1^{j,0} \parallel \dots \parallel \mathbf{y}_n^{j,0})$, for $j \geq \rho$	Lemma 5

Fig. 9. An overview of the games used in the proof of Theorem 4. By $[\text{ct}_i^{\text{in},k}]_1$ and $[\text{sk}_i^{\text{in},j}]_2$ we denote the k^{th} ciphertext and the j^{th} decryption key of the inner scheme \mathcal{MLFE}' .

assumption to argue that $\langle \mathbf{r}^\rho, \mathbf{z}_i \rangle_T$ is indistinguishable from a uniform random value and thus perfectly hides $\langle \mathbf{x}_i^{1,1}, \mathbf{y}_i^{\rho,\beta} \rangle$ for the first ciphertext of each slot: ct_i^1 . For all the other remaining $\langle \mathbf{x}_i^{j,1}, \mathbf{y}_i^{\rho,\beta} \rangle$, for $j \in [Q_i], j > 1$, we use the fact that $\langle \mathbf{x}_i^{j,1} - \mathbf{x}_i^{1,1}, \mathbf{y}_i^{\rho,0} \rangle = \langle \mathbf{x}_i^{j,1} - \mathbf{x}_i^{1,1}, \mathbf{y}_i^{\rho,1} \rangle$, as implied by the game's restrictions.

Proof of Theorem 4. We proceed via a series of Games $G_0, G_1, G_{1,\rho}$, for $\rho \in [Q_f + 1]$, described in Fig. 10. An overview is provided in Fig. 9. Let \mathcal{A} be a PPT adversary, and $\lambda \in \mathbb{N}$ be the security parameter. We denote by $\text{Adv}_{G_i}(\mathcal{A})$ the advantage of \mathcal{A} in game G_i .

- G_0 :** is the experiment **many-SEL-wFH-IND** $_0^{MIFE}$ (see Definition 6).
- G_1 :** we replace the inner encryption of $\mathbf{x}_i^{j,0}$ by encryptions of $\mathbf{x}_i^{j,1}$, for all $i \in [n], j \in [Q_i]$, using the many-SEL-IND security of $MIFE'$. This is possible due to the weak function-hiding constraint, which states in particular that $\sum_{i=1}^n \langle \mathbf{x}_i^{j_i,0}, \mathbf{y}_i^{j_f,0} \rangle = \sum_{i=1}^n \langle \mathbf{x}_i^{j_i,1}, \mathbf{y}_i^{j_f,0} \rangle$, for all indices $j_i \in [Q_i], j_f \in [Q_f]$.
- $G_{1,\rho}$:** for the first $\rho - 1$ queries to KeyGen , we replace inner secret key $\text{KeyGen}'(\text{msk}', \mathbf{y}_1^0 \parallel \dots \parallel \mathbf{y}_n^0)$, by $\text{KeyGen}'(\text{msk}', \mathbf{y}_1^1 \parallel \dots \parallel \mathbf{y}_n^1)$. Note that G_1 is the same as $G_{1,1}$, and G_{1,Q_f+1} is the same as **many-SEL-wFH-IND** $_1^{MIFE}$.

We prove $G_0 \approx_c G_1$ in Lemma 4, and $G_{1,\rho} \approx_c G_{1,\rho+1}$ for all $\rho \in [Q_f]$ in Lemma 5. □

$G_0, \boxed{G_1, \boxed{G_{1,\rho^1}}}$, for $\rho \in [Q_f + 1]$:

$\{\mathbf{x}_i^{j,\beta}\}_{i \in [n], j \in [Q_i], \beta \in \{0,1\}}, \{\mathbf{y}_i^{j,\beta}\}_{i \in [n], j \in [Q_f], \beta \in \{0,1\}} \leftarrow \mathcal{A}(1^\lambda, \mathcal{F}_n^{m,X,Y})$
 $(\text{mpk}', \text{msk}') \leftarrow \text{Setup}'(1^\lambda, \mathcal{F}_n^{m,X,Y})$
 $\forall i \in [n] : (\text{mpk}_i, \text{msk}_i) \leftarrow \mathcal{FE}.\text{Setup}(1^\lambda, \mathcal{F}_1^{\ell,X,Y})$
 $\text{mpk} := (\{\text{mpk}_i\}_{i \in [n]}, \text{mpk}')$, $\text{msk} := (\{\text{msk}_i\}_{i \in [n]}, \text{msk}')$
 $\forall i \in [n], j \in [Q_i]$:
 $\boxed{\text{ct}_i^{\text{in},j} := \text{Enc}'(\text{msk}', i, \mathbf{x}_i^{j,0})}$, $\boxed{\text{ct}_i^{\text{in},j} := \text{Enc}'(\text{msk}', i, \mathbf{x}_i^{j,1})}$
 $\text{ct}_i^{\text{out},j} := \mathcal{FE}.\text{KeyGen}(\text{msk}_i, \text{ct}_i^{\text{in},j})$
 $\alpha \leftarrow \mathcal{A}^{\text{KeyGen}(\text{msk}, \cdot)}(\text{mpk}, \{\text{ct}_i^{\text{out},j}\}_{i \in [n], j \in [Q_i]})$
Output: α .

KeyGen($\text{msk}, (\mathbf{y}_1^{j,\beta} \parallel \dots \parallel \mathbf{y}_n^{j,\beta})_{\beta \in \{0,1\}}$):

$(\{\text{sk}_i^{\text{in}}\}_2, [z]_T) \leftarrow \text{KeyGen}'(\text{msk}', \mathbf{y}_1^0 \parallel \dots \parallel \mathbf{y}_n^0)$
 If $j < \rho$: $(\{\text{sk}_i^{\text{in}}\}_2, [z]_T) \leftarrow \text{KeyGen}'(\text{msk}', \mathbf{y}_1^1 \parallel \dots \parallel \mathbf{y}_n^1)$

$\text{sk}_i^{\text{out}} \leftarrow \mathcal{FE}.\text{Enc}(\text{msk}_i, \{\text{sk}_i^{\text{in}}\}_2)$
 $\text{sk}_{\mathbf{y}_1 \parallel \dots \parallel \mathbf{y}_n} := (\{\text{sk}_i^{\text{out}}\}_{i \in [n]}, [z]_T)$
 return $\text{sk}_{\mathbf{y}_1 \parallel \dots \parallel \mathbf{y}_n}$

Fig. 10. Games for the proof of Theorem 4. In each procedure, the components inside a solid (dotted) frame are only present in the games marked by a solid (dotted) frame.

Lemma 4 (G_0 to G_1). *There exists a PPT adversary \mathcal{B}_1 such that*

$$\text{Adv}_{G_0}(\mathcal{A}) - \text{Adv}_{G_1}(\mathcal{A}) \leq \text{Adv}_{\text{MIFE}', \mathcal{B}_1}^{\text{many-SEL-IND}}(\lambda).$$

Proof. In order to show that we can switch $\mathbf{x}_i^{j,0}$ to $\mathbf{x}_i^{j,1}$, we rely on the security of the underlying MIFE' scheme. Intuitively, adding an additional layer of encryption on the decryption keys sk_i^{in} cannot invalidate the security of the underlying MIFE' .

More formally, we design an adversary \mathcal{B}_1 against the many-SEL-IND security of MIFE' . Adversary \mathcal{B}_1 draws public and secret keys for the outer encryption layer and then uses its own experiment to simulate either G_0 or G_1 . We describe adversary \mathcal{B}_1 in the full version of the paper and give a textual description here.

Simulation of master public key mpk . Since the game is selective, the adversary \mathcal{B}_1 first gets the challenges $\{\mathbf{x}_i^{j,b}\}_{i \in [n], j \in [Q_i], b \in \{0,1\}}$ from \mathcal{A} , and it sends them to its experiment $\text{many-SEL-IND}_{\beta}^{\text{MIFE}'}$. Then, \mathcal{B}_1 receives the public key mpk' of the MIFE' scheme. To construct the full public key, it draws $(\text{mpk}_i, \text{msk}_i) \leftarrow \mathcal{FE}.\text{Setup}(1^\kappa, \mathcal{F}_1^{\ell, X, Y})$, for all slots $i \in [n]$, independently. It then sets $\text{mpk} := \{\text{mpk}_i\}_{i \in [n]} \cup \{\text{mpk}'\}$ and returns mpk to adversary \mathcal{A} .

Simulation of the challenge ciphertexts. The adversary \mathcal{B}_1 receives $[\text{ct}_i^{\text{in},j}]_1$ from the encryption oracle of the experiment $\text{many-SEL-IND}_{\beta}^{\text{MIFE}'}$, for all $i \in [n]$. This corresponds to encryptions of either $\mathbf{x}_i^{j,\beta}$, for $\beta = 0$ or 1 . Since it knows msk_i , it computes $\text{ct}_i^{\text{out},j} := \mathcal{FE}.\text{KeyGen}(\text{msk}_i, [\text{ct}_i^{\text{in},j}]_1)$ for all $i \in [n]$ and returns $\{\text{ct}_i^{\text{out},j}\}_{i \in [n]}$ to \mathcal{A} .

Simulation of $\text{KeyGen}(\text{msk}, \cdot)$. On every secret key query $(\mathbf{y}_1^b \parallel \dots \parallel \mathbf{y}_n^b)_{b \in \{0,1\}}$, adversary \mathcal{B}_1 queries the KeyGen' oracle of the experiment $\text{many-SEL-IND}_{\beta}^{\text{MIFE}'}$ on $\mathbf{y}_1^0 \parallel \dots \parallel \mathbf{y}_n^0$. It obtains $\{[\text{sk}_i^{\text{in}}]_2\}_{i \in [n]}, [z]_T$. Finally, it computes $\text{sk}_i^{\text{out}} := \mathcal{FE}.\text{Enc}(\text{mpk}_i, [\text{sk}_i^{\text{in}}]_2)$ and returns $(\{\text{sk}_i^{\text{out}}\}_{i \in [n]}, [z]_T)$ to \mathcal{A} . □

Lemma 5 ($G_{1,\rho}$ to $G_{1,\rho+1}$). *For all $\rho \in [Q_f]$, there exist PPT adversaries \mathcal{B}_ρ and \mathcal{B}'_ρ such that*

$$\text{Adv}_{G_{1,\rho}}(\mathcal{A}) - \text{Adv}_{G_{1,\rho+1}}(\mathcal{A}) \leq 2n \cdot \text{Adv}_{\mathcal{FE}, \mathcal{B}_\rho}^{\text{one-SEL-SIM}}(\lambda) + 2 \cdot \text{Adv}_{G_1, \mathcal{B}'_\rho}^{\text{Dk-mddh}}(\lambda) + \frac{2k}{p}.$$

For lack of space the proof of Lemma 5 appears in the full version of the paper [3].

4.2 Adaptively-Secure Multi-Input Function-Hiding FE for Inner Product

In this section, we prove that if we instantiate the construction described in Fig. 8 (Sect. 4), with the many-AD-IND-secure, single-input FE from [5], we obtain an adaptively secure function-hiding MIFE. Specifically, we consider the generalized

version of single-input FE, as described in [4] (recalled in the full version of the paper [3]). For completeness, we present this new MIFE instantiation in Fig. 11. Proving adaptive security for our construction in a generic way would require the underlying \mathcal{FE} to achieve strong security notions, such as one-AD-SIM (which is not achieved by any known scheme). We overcome this issue, managing to prove adaptive security of our concrete MIFE in Fig. 8, using non-generic techniques inspired by [4].

Theorem 5 (many-AD-IND-wFH security). *If the \mathcal{D}_k -MDDH assumption holds in \mathbb{G}_1 and \mathbb{G}_2 , then the multi-input FE for $\mathcal{F}_n^{m,X,Y}$ described in Fig. 11 is many-AD-IND-wFH-secure.*

Proof overview. Similarly to the selective-security proof presented in Sect. 4.1, we prove weakly-function-hiding. This is sufficient, since it can be transformed generically into a fully function-hiding MIFE by using techniques from [14] (see in the full version of the paper [3] for more details).

To prove weak function-hiding we proceed in two stages. First, we switch from $\text{Enc}(\text{msk}, i, \mathbf{x}_i^{j,0})$ to $\text{Enc}(\text{msk}, i, \mathbf{x}_i^{j,1})$ for all slots $i \in [n]$ and all queries $j \in [Q_i]$ simultaneously, using the many-AD-IND security of \mathcal{MIFE}' (the underlying

<p><u>Setup($1^\lambda, \mathcal{F}_n^{m,X,Y}$):</u></p> <p>$\mathcal{PG} \leftarrow \text{GGen}(1^\lambda), \mathbf{A}_1, \mathbf{B}_1, \dots, \mathbf{A}_n, \mathbf{B}_n \leftarrow_{\mathcal{R}} \mathcal{D}_k, \mathbf{U}_1, \dots, \mathbf{U}_n \leftarrow_{\mathcal{R}} \mathbb{Z}_p^{(k+m) \times (k+1)}$ $\mathbf{V}_1, \dots, \mathbf{V}_n \leftarrow_{\mathcal{R}} \mathbb{Z}_p^{(2k+m+1) \times (k+1)}, \mathbf{z}_1, \dots, \mathbf{z}_n \leftarrow_{\mathcal{R}} \mathbb{Z}_p^k$ $\text{mpk} := \mathcal{PG}, \text{msk} := \{\mathbf{A}_i, \mathbf{B}_i, \mathbf{U}_i, \mathbf{V}_i, \mathbf{z}_i\}_{i \in [n]}$ return (mpk, msk)</p> <p><u>Enc(msk, $i, \mathbf{x}_i \in \mathbb{Z}_p^m$):</u></p> <p>$\mathbf{s}_i \leftarrow_{\mathcal{R}} \mathbb{Z}_p^k, \mathbf{c}_i := \mathbf{A}_i \mathbf{s}_i, \mathbf{c}'_i := \begin{pmatrix} \mathbf{x}_i \\ \mathbf{z}_i \end{pmatrix} + \mathbf{U}_i \mathbf{A}_i \mathbf{s}_i, \mathbf{c}''_i := \begin{pmatrix} \mathbf{c}_i \\ \mathbf{c}'_i \end{pmatrix}^\top \mathbf{V}_i$ return ($[\mathbf{c}_i]_1, [\mathbf{c}'_i]_1, [\mathbf{c}''_i]_1$)</p> <p><u>KeyGen(msk, $\mathbf{y}_1 \parallel \dots \parallel \mathbf{y}_n \in (\mathbb{Z}_p^m)^n$):</u></p> <p>$\mathbf{r} \leftarrow_{\mathcal{R}} \mathbb{Z}_p^k, \mathbf{z} := \langle \mathbf{z}_1 + \dots + \mathbf{z}_n, \mathbf{r} \rangle$</p> <p>$\forall i \in [n]: \mathbf{t}_i \leftarrow_{\mathcal{R}} \mathbb{Z}_p^k, \mathbf{d}_i := \mathbf{B}_i \mathbf{t}_i, \mathbf{d}'_i := \begin{pmatrix} -\mathbf{U}_i^\top \begin{pmatrix} \mathbf{y}_i \\ \mathbf{r} \end{pmatrix} \\ \begin{pmatrix} \mathbf{y}_i \\ \mathbf{r} \end{pmatrix} \end{pmatrix} + \mathbf{V}_i \mathbf{B}_i \mathbf{t}_i$</p> <p>return ($\{[\mathbf{d}_i]_2, [\mathbf{d}'_i]_2\}_{i \in [n]}, [\mathbf{z}]_T$)</p> <p><u>Dec($(\{[\mathbf{d}_i]_2, [\mathbf{d}'_i]_2\}_{i \in [n]}, [\mathbf{z}]_T), \{[\mathbf{c}_i]_1, [\mathbf{c}'_i]_1, [\mathbf{c}''_i]_1\}_{i \in [n]}$):</u></p> <p>$out \leftarrow \left(\prod_i \left(e \left(\begin{bmatrix} \mathbf{c}_i \\ \mathbf{c}'_i \end{bmatrix}_1^\top, [\mathbf{d}_i]_2 \right) / e([\mathbf{c}''_i]_1, [\mathbf{d}_i]_2) \right) \right) / [\mathbf{z}]_T$ return discrete log of out</p>
--

Fig. 11. Many-AD-IND-wFH secure, multi-input FE scheme for the class $\mathcal{F}_n^{m,X,Y}$ (self-contained description).

MIFE from [4]). For completeness, we also give a concrete description of \mathcal{MIFE}' in the full version of the paper.

Secondly, we use a hybrid argument over all Q_f queried keys, switching them one by one from $\text{KeyGen}(\text{msk}, \mathbf{y}_1^0 \parallel \dots \parallel \mathbf{y}_n^0)$ to $\text{KeyGen}(\text{msk}, \mathbf{y}_1^1 \parallel \dots \parallel \mathbf{y}_n^1)$. To switch the ρ 'th key, we use the security of \mathcal{FE} in a non-generic way. Structurally, we do a proof similar to the selective one of the previous section. In order to apply complexity leveraging, we first do all the computational steps. Afterwards, only at some particular transition in the proof (transition from $H''_{\rho,0}$ to $H''_{\rho,1}$ the full version of the paper), we use complexity leveraging, and we simulate the selective proof arguments. This multiplies the security loss by an exponential factor. We can do so here because this particular transition is perfect: the exponential term is multiplied by a zero advantage.

Although this proof strategy shares similarities with the adaptive security proof of the MIFE in [4], our proof has some crucial differences: mainly, the role of the keys and ciphertexts in our proof is switched. Since the multi-input model is asymmetric with respect to the ciphertexts and decryption keys (only ciphertexts can be mixed-and-matched), this results in a different proof strategy.

For lack of space, the full proof appears in the full version of the paper.

Acknowledgments. Michel Abdalla was supported in part by SAFEcrypto (H2020 ICT-644729) and by the European Union's Horizon 2020 Research and Innovation Programme under grant agreement 780108 (FENITEC). Dario Fiore was partially supported by the Spanish Ministry of Economy under project references TIN2015-70713-R (DEDETIS), RTC-2016-4930-7 (DataMantium), and by the Madrid Regional Government under project N-Greens (ref. S2013/ICE-2731). Romain Gay was partially supported by a Google PhD Fellowship in Privacy and Security and by the ERC Project aSCEND (H2020 639554). Bogdan Ursu was partially supported by ANR-14-CE28-0003 (Project EnBiD) and by the ERC Project PREP-CRYPTO (H2020 724307).

References

1. Abdalla, M., Bourse, F., De Caro, A., Pointcheval, D.: Simple functional encryption schemes for inner products. In: Katz, J. (ed.) PKC 2015. LNCS, vol. 9020, pp. 733–751. Springer, Heidelberg (2015). https://doi.org/10.1007/978-3-662-46447-2_33
2. Abdalla, M., Bourse, F., De Caro, A., Pointcheval, D.: Better security for functional encryption for inner product evaluations. Cryptology ePrint Archive, Report 2016/011 (2016). <http://eprint.iacr.org/2016/011>
3. Abdalla, M., Catalano, D., Fiore, D., Gay, R., Ursu, B.: Multi-Input Functional Encryption for Inner Products: Function-Hiding Realizations and Constructions without Pairings. Cryptology ePrint Archive, Report 2017/972 (2017). <http://eprint.iacr.org/2017/972>
4. Abdalla, M., Gay, R., Raykova, M., Wee, H.: Multi-input inner-product functional encryption from pairings. In: Coron, J.-S., Nielsen, J.B. (eds.) EUROCRYPT 2017, Part I. LNCS, vol. 10210, pp. 601–626. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-56620-7_21

5. Agrawal, S., Libert, B., Stehlé, D.: Fully secure functional encryption for inner products, from standard assumptions. In: Robshaw, M., Katz, J. (eds.) CRYPTO 2016, Part III. LNCS, vol. 9816, pp. 333–362. Springer, Heidelberg (2016). https://doi.org/10.1007/978-3-662-53015-3_12
6. Ananth, P., Jain, A.: Indistinguishability obfuscation from compact functional encryption. In: Gennaro, R., Robshaw, M. (eds.) CRYPTO 2015, Part I. LNCS, vol. 9215, pp. 308–326. Springer, Heidelberg (2015). https://doi.org/10.1007/978-3-662-47989-6_15
7. Badrinarayanan, S., Gupta, D., Jain, A., Sahai, A.: Multi-input functional encryption for unbounded arity functions. In: Iwata, T., Cheon, J.H. (eds.) ASIACRYPT 2015, Part I. LNCS, vol. 9452, pp. 27–51. Springer, Heidelberg (2015). https://doi.org/10.1007/978-3-662-48797-6_2
8. Boneh, D., Sahai, A., Waters, B.: Functional encryption: definitions and challenges. In: Ishai, Y. (ed.) TCC 2011. LNCS, vol. 6597, pp. 253–273. Springer, Heidelberg (2011). https://doi.org/10.1007/978-3-642-19571-6_16
9. Brakerski, Z., Komargodski, I., Segev, G.: Multi-input functional encryption in the private-key setting: stronger security from weaker assumptions. In: Fischlin, M., Coron, J.-S. (eds.) EUROCRYPT 2016, Part II. LNCS, vol. 9666, pp. 852–880. Springer, Heidelberg (2016). https://doi.org/10.1007/978-3-662-49896-5_30
10. Datta, P., Okamoto, T., Tomida, J.: Full-hiding (unbounded) multi-input inner product functional encryption from the k -linear assumption. In: Abdalla, M., Dahab, R. (eds.) PKC 2018, Part II. LNCS, vol. 10770, pp. 245–277. Springer, Cham (2018). https://doi.org/10.1007/978-3-319-76581-5_9
11. Escala, A., Herold, G., Kiltz, E., Ràfols, C., Villar, J.: An algebraic framework for Diffie-Hellman assumptions. In: Canetti, R., Garay, J.A. (eds.) CRYPTO 2013, Part II. LNCS, vol. 8043, pp. 129–147. Springer, Heidelberg (2013). https://doi.org/10.1007/978-3-642-40084-1_8
12. Goldwasser, S., et al.: Multi-input functional encryption. In: Nguyen, P.Q., Oswald, E. (eds.) EUROCRYPT 2014. LNCS, vol. 8441, pp. 578–602. Springer, Heidelberg (2014). https://doi.org/10.1007/978-3-642-55220-5_32
13. Lin, H.: Indistinguishability obfuscation from SXDH on 5-linear maps and locality-5 PRGs. In: Katz, J., Shacham, H. (eds.) CRYPTO 2017, Part I. LNCS, vol. 10401, pp. 599–629. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-63688-7_20
14. Lin, H., Vaikuntanathan, V.: Indistinguishability obfuscation from DDH-like assumptions on constant-degree graded encodings. In: Dinur, I. (ed.) 57th FOCS, pp. 11–20. IEEE Computer Society Press, October 2016
15. O’Neill, A.: Definitional issues in functional encryption. Cryptology ePrint Archive, Report 2010/556 (2010). <http://eprint.iacr.org/2010/556>
16. Sahai, A., Waters, B.: Fuzzy identity-based encryption. In: Cramer, R. (ed.) EUROCRYPT 2005. LNCS, vol. 3494, pp. 457–473. Springer, Heidelberg (2005). https://doi.org/10.1007/11426639_27
17. Wee, H.: Attribute-hiding predicate encryption in bilinear groups, revisited. In: Kalai, Y., Reyzin, L. (eds.) TCC 2017, Part I. LNCS, vol. 10677, pp. 206–233. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-70500-2_8