Simple Functional Encryption Schemes
for Inner Products

Michel Abdalla[1], Florian Bourse[1], Angelo De Caro[1], and David Pointcheval[1]

Département d'Informatique, École normale supérieure
45 Rue d'Ulm, 75230 Paris Cedex 05, France
{michel.abdalla,florian.bourse,angelo.decaro,david.pointcheval}@ens.fr

Abstract. Functional encryption is a new paradigm in public-key encryption that allows users to finely control the amount of information that is revealed by a ciphertext to a given receiver. Recent papers have focused their attention on constructing schemes for general functionalities at expense of efficiency. Our goal, in this paper, is to construct functional encryption schemes for less general functionalities which are still expressive enough for practical scenarios. We propose a functional encryption scheme for the *inner-product* functionality, meaning that decrypting an encrypted vector $\mathbf{x}$ with a key for a vector $\mathbf{y}$ will reveal only $\langle \mathbf{x}, \mathbf{y} \rangle$ and nothing else, whose security is based on the DDH assumption. Despite the simplicity of this functionality, it is still useful in many contexts like descriptive statistics. In addition, we generalize our approach and present a generic scheme that can be instantiated, in addition, under the LWE assumption and offers various trade-offs in terms of expressiveness and efficiency.

Keywords. Functional Encryption, Inner-Product, Generic Constructions.

1 Introduction

Functional Encryption. Whereas, in traditional public-key encryption, decryption is an all-or-nothing affair (i.e., a receiver is either able to recover the entire message using its key, or nothing), in *functional encryption* (FE), it is possible to finely control the amount of information that is revealed by a ciphertext to a given receiver. For example, decrypting an encrypted data set with a key for computing the mean will reveal only the mean computed over the data set and nothing else. Somewhat more precisely, in a functional encryption scheme for functionality F, each secret key (generated by a master authority having a *master secret key*) is associated with value k in some *key space* K; Anyone can encrypt via the public parameters; When a ciphertext Ct_x that encrypts x, in some *message space* X, is decrypted using a secret key Sk_k for value k, the result is $F(k, x)$. A notable subclass of functional encryption is that of *predicate encryption* (PE) which are defined for functionalities whose message space

X consists of two subspaces I and M called respectively *index space* and *payload space*. In this case, the functionality F is defined in terms of a predicate $P : K \times I \rightarrow \{0, 1\}$ as follows: $F(k, (\mathsf{ind}; \mathsf{m})) = \mathsf{m}$ if $P(k, \mathsf{ind}) = 1$, and $\perp$ otherwise, where $k \in K$, $\mathsf{ind} \in I$ and $\mathsf{m} \in M$. Those schemes are also called *predicate encryption with private-index*. Examples of those schemes are Anonymous Identity-Based Encryption (AIBE) [BF01, Gen06], Hidden Vector Encryption [BW07] and Orthogonality [KSW08, LOS$^+$10, OT12], among the others. On the other hand, when the index ind is easily readable from the ciphertext those schemes are called *predicate encryption with public-index* (PIPE). Examples of PIPE schemes are Identity-Based Encryption (IBE) [Sha84, BF01, Coc01], Attribute-Based Encryption (ABE) [SW05, GPSW06], Functional Encryption for Regular Languages [Wat12].

The standard notion of security for functional encryption is that of *indistinguishability-based security* (IND). Informally, it requires that an adversary cannot tell apart which of two messages x_0, x_1 has been encrypted having oracle access to the key generation algorithm under the constraint that, for each k for which the adversary has seen a secret key, it holds that $F(k, x_0) = F(k, x_1)$. This models the idea that an individual's messages are still secure even if an arbitrary number of other users of the system collude against that user. Boneh, Sahai, and Waters [BSW11] and O'Neill [O'N10] showed that the IND definition is weak in the sense that a trivially insecure scheme implementing a certain functionality can be proved IND-secure anyway. The authors, then, initiate the study of *simulation-based* (SIM) notions of security for FE, which asks that the "view" of the adversary can be simulated by simulator given neither ciphertexts nor keys but only the corresponding outputs of the functionality on the underlying plaintexts, and shows that SIM-security is not always achievable.

In a recent series of outstanding results, [GGH$^+$13, BCP14, Wat14, GGHZ14] proposed IND-secure FE schemes for general circuits whose security is based either on indistinguishable obfuscation and its variants or polynomial hardness of simple assumptions on multilinear maps. Those schemes are far from being practical and this led us to investigate the possibility of having functional encryption schemes for *functionalities of practical interest* which are still expressive enough for practical scenarios. In doing so, we seek for schemes that offer *simplicity*, in terms of understanding of how the schemes work, and *adaptability* in terms of the possibility of choosing the instantiations and the parameters that better fit the constraints and needs of a specific scenario the user is interested in.

This Work. In this paper, we focus on the *inner-product* functionality, which has several practical applications. For example, in descriptive statistics, the discipline of quantitatively describing the main features of a collection of information, the *weighted mean* is a useful tool. Here are a few examples:

Slugging average in baseball. A batter's slugging average, also called slugging percentage, is computed by: $\mathsf{SLG} = (1 * \mathsf{SI} + 2 * \mathsf{DO} + 3 * \mathsf{TR} + 4 * \mathsf{HR})/\mathsf{AB}$, where SLG is the slugging percentage, SI is the number of singles, DO the number of doubles, TR the number of triples, HR the number of home runs,

and AB is total number of at-bats. Here, each single has a *weight* of 1, each double has a *weight* of 2, etc. The average counts home runs four times as important as singles, and so on. An at-bat without a hit has a *weight* of zero.

Course grades. A teacher might say that the *test average* is 60% of the grade, *quiz average* is 30% of the grade, and a *project* is 10% of the grade. Suppose Alice got 90 and 78 on the tests; 100, 100 and 85 on the quizzes; and an 81 on the project. Then, Alice's test average is $(90 + 78)/2 = 84$, quiz average is $(100 + 100 + 85)/3 = 95$, and her course grade would then be: $.60 \cdot 84 + .30 \cdot 95 + .10 \cdot 81 = 87$.

Our goal then is to design a simple and efficient functional encryption scheme for inner products that can be used, for instance, to compute a weighted mean and to protect the privacy of Alice's grades, in the example involving course grades. In fact, we can imagine that Alice's grades, represented as a vector $\mathbf{x} = (x_1, \ldots, x_\ell)$ in some finite field, says $\mathbb{Z}_p$ for prime p, are encrypted in a ciphertext $\mathsf{Ct}_\mathbf{x}$ and the teacher has a secret key $\mathsf{Sk}_\mathbf{y}$ for the vector of weights $\mathbf{y} = (y_1, \ldots, y_\ell)$. Then Alice's course grade can be computed as the *inner-product* of $\mathbf{x}$ and $\mathbf{y}$, written as $\langle \mathbf{x}, \mathbf{y} \rangle = \sum_{i \in [\ell]} x_i \cdot y_i$. We would like to stress here that, unlike the inner-product predicate schemes in [KSW08,LOS$^+$10,OT12], our goal is to output the actual value of the inner product.

A very simple scheme can be constructed to compute the above functionality whose security can be based on the DDH assumption. Informally, it is like this:

$$mpk = \big(\mathbb{G}, (h_i = g^{s_i})_{i \in [\ell]}\big)$$
$$\mathsf{Ct}_\mathbf{x} = \big(\mathsf{ct}_0 = g^r, (\mathsf{ct}_i = h_i^r \cdot g^{x_i})_{i \in [\ell]}\big)$$
$$\mathsf{Sk}_\mathbf{y} = \langle \mathbf{s}, \mathbf{y} \rangle = \sum_{i \in [\ell]} s_i \cdot y_i,$$

where $msk = \mathbf{s} = (s_1, ..., s_\ell)$ is the master secret key used to generate secret keys $\mathsf{Sk}_\mathbf{y}$. Then, decryption is done by computing the discrete log of $(\prod_{i \in [\ell]} \mathsf{ct}_i^{y_i})/\mathsf{ct}_0^{\mathsf{Sk}_\mathbf{y}}$. Please refer to Section 3 for more details.

Despite its simplicity, this DDH-based scheme can be proved secure, in a selective security model[1], against any adversary that issues an unbounded, but polynomially related to the security parameter, number of secret key queries. The adversary will not learn anything more than what it is implied by the linear combination of their keys.

An astute reader could now ask what happens if an adversary possesses secret keys $\mathsf{Sk}_{\mathbf{y}_i}$, for $i \in [q]$, such that the $\mathbf{y}_i$'s form a basis for $\mathbb{Z}_p^\ell$. Clearly, this adversary can then recover completely $\mathbf{x}$ from the ciphertext and wins the security game. But notice that this has nothing to do with the specific implementation of the functionality, it is something inherent to the functionality itself. This happens also for other functionalities: Consider the case of the *circuit* functionality, where secret keys correspond to Boolean circuits over ℓ Boolean variables and one-bit

---

[1] In the selective model, the adversary is asked to commit to its challenge before seeing the public parameters.

output, and ciphertexts to vectors in $\{0, 1\}^{\ell}$. Then, an adversary having secret keys for circuits C_i, where C_i extract the i-th bit of the input, can recover completely the input from a ciphertext no matter how the scheme, supporting this circuit functionality, is implemented. This subtle aspect will appear in the security proof, if the adversary asks such a set of secret keys then the simulator will not be able to answer all those queries. Notice that this is reasonable, because it is like requiring security when the adversary possesses the master secret key.

One drawback of the above scheme is that restrictions must be put in place in order to guarantee that the final computed value has small magnitude and the discrete log can be computed efficiently. To overcome this limitation, let us first describe some interesting characteristics of the above scheme. To start, please notice that a ciphertext for a vector $\mathbf{x}$ consists of ElGamal ciphertexts [ElG84] $(g^r, (h_i^r \cdot g^{x_i})_i)$ under public keys $h_i = g^{s_i}$, sharing the same randomness r. Then, a secret key for a vector $\mathbf{y}$ consists of a linear combination of the underlying ElGamal secrets s_i. Now notice that, by the ElGamal scheme's homomorphic properties, it holds that

$$\prod_{i \in [\ell]} \mathsf{ct}_i^{y_i} = \prod_{i \in [\ell]} h_i^{r \cdot y_i} \cdot g^{x_i \cdot y_i} = h^r \cdot g^{\langle \mathbf{x}, \mathbf{y} \rangle}$$

where h is an ElGamal public key corresponding to secret key $\langle \mathbf{s}, \mathbf{y} \rangle$. The above observations point out that, by possibly combining public-key encryption schemes secure under randomness reuse [BBS03], and having specific syntactical, non-security-related, properties, we can generalize the above construction with the aim of (1) having a scheme whose security can be based on different assumptions and that can provide different trade-offs in terms of efficiency and expressiveness, (2) have a generic proof of security that reduces security to that of the underlying public-key encryption scheme. We present our generalization in Section 4.

Related Work. One of the first example of investigation on reductions between various primitives has been given by [Rot11] who shows a simple reduction between any semantically secure private-key encryption scheme which possesses a simple homomorphic, namely that the product of two ciphertexts Ct_1 and Ct_2, encrypting plaintexts m_1 and m_2, yields a new ciphertext $\mathsf{Ct} = \mathsf{Ct}_1 \cdot \mathsf{Ct}_2$ which decrypts to $\mathsf{m}_1 + \mathsf{m}_2 \mod 2$. Goldwasser *et al.* [GLW12] investigated the construction of PIPE schemes for specific predicates. In particular, they show how public-key encryption schemes which possess a linear homomorphic property over their keys as well as hash proof system features with certain algebraic structure can be used to construct an efficient identity-based encryption (IBE) scheme that is secure against bounded collusions (BC-IBE). This weaker security notion restricts the adversary to issue only a bounded number of secret keys during the security game. In more details, they rely on a public-key encryption scheme whose secret keys and public keys are elements of respective groups (with possibly different operations, which we denote by $+$ and $\cdot$), and there exists an homomorphism μ such that $\mu(\mathsf{sk}_0 + \mathsf{sk}_1) = \mu(\mathsf{sk}_0) \cdot \mu(\mathsf{sk}_1)$, where $\mu(\mathsf{sk}_0)$ and $\mu(\mathsf{sk}_1)$

are valid public keys for which sk_0 and sk_1 yield correct decryption, respectively. Then, to obtain a BC-IBE, the construction by Goldwasser *et al.* generates multiple public-key/secret-key pairs $(\mathsf{pk}_1, \mathsf{sk}_1), \ldots, (\mathsf{pk}_\ell, \mathsf{sk}_\ell)$, letting the public-parameters and the master secret key of the scheme be $\mathsf{params} = (\mathsf{pk}_1, \ldots, \mathsf{pk}_\ell)$ and $msk = (\mathsf{sk}_1, \ldots, \mathsf{sk}_\ell)$, respectively. Then, an efficient map ϕ associates every identity ID with a vector $[id_1, \ldots, id_\ell]$, and a message m is encrypted for an identity ID as the ciphertext $\mathsf{Ct} = \mathsf{Encrypt}(\mathsf{pk}_{\mathsf{ID}}, m)$, where $\mathsf{pk}_{\mathsf{ID}} = \prod_{i=1}^{n} \mathsf{pk}_i^{\mathsf{ID}_i}$. Then, μ guarantees that Ct can be decrypted using $\mathsf{sk}_{\mathsf{ID}} = \sum_{i=1}^{n} \mathsf{ID}_i \cdot \mathsf{sk}_i$, since by the homomorphism it holds that $\mathsf{pk}_{\mathsf{ID}} = \mu(\mathsf{sk}_{\mathsf{ID}})$. The map ϕ is subject to a combinatorial requirement that disallows computing $\mathsf{sk}_{\mathsf{ID}}$ given $\mathsf{sk}_{\mathsf{ID}'}$ for t different $\mathsf{ID}' \neq \mathsf{ID}$. Later, Tessaro and Wilson [TW14] presented generic constructions of BC-IBE which rely on encryption schemes that solely satisfy the standard security notion of semantic security in addition to some syntactical, non-security-related, properties. We will use Tessaro and Wilson to present a generalization of our results.

As already mentioned, in a recent series of outstanding results, [GGH+13, BCP14, Wat14, GGHZ14] proposed IND-secure FE scheme for general circuits whose security is based either on indistinguishable obfuscation and its variants or polynomial hardness of simple assumptions on multilinear maps. Clearly, those schemes can be used to implement the inner-product functionality, but this would defeat our main goal which is to construct somewhat efficient functional encryption schemes for less general functionalities which are still expressive enough for practical scenarios.

In another line of research, Katz, Sahai, and Waters [KSW08] proposed a functional encryption scheme for a functionality called *orthogonality* (in the literature it is also know as inner-product encryption), meaning that decrypting an encrypted vector $\mathbf{x}$ with a key for a vector $\mathbf{y}$ will reveal only if $\langle \mathbf{x}, \mathbf{y} \rangle$ is equals to zero (meaning that the two vectors are orthogonal) or not, and nothing else. This functionality is of little help in our case because what we need is for the decryptor to be able to recover the value $\langle \mathbf{x}, \mathbf{y} \rangle$.

Parameter Sizes For Our Constructions. In Table 1, we present the size of the parameters and ciphertexts for our concrete construction. Each column refers to an instantiation of our general scheme and is indexed by its underlying assumption. Each row describes the size of an element of the scheme. The master public key size does not include public parameters that can be re-used such that $(g, \mathbf{A})$. Message space is the number of different messages that can be decrypted using this instantiation. The DDH instantiation can give short ciphertexts and keys using elliptic curves, but requires the computation of a discrete logarithm in the decryption, which makes it usable for small message space only. LWE enables short ciphertexts and keys together with a message space that is independent of the security parameter, which allows shorter ciphertexts for a smaller message space.

	DDH	LWE
mpk	$\ell \log p$	$(k+1)m\ell \log q$
msk	$\ell \log p$	$(k+1)n\ell \log q$
$\mathsf{Ct_x}$	$(\ell+1)\log p$	$(n+(k+1)\ell)\log q$
$\mathsf{Sk_y}$	$\log p$	$n \log q$
message space	bounded by computation	$\ell r^{2(k+1)}$
$\lvert\{\langle \mathbf{x}, \mathbf{y}\rangle\}\rvert$	of discrete logarithms	computations $\mod p$

Table 1. Parameter Sizes. Here, ℓ is the length of the vectors encrypted in the ciphertexts and encoded in the secret keys. In DDH, p is the size of a group element. In LWE, q is the size of a group element, p is the order of the message space, and k is the size of a message in base r.

2   Basic Tools

In this section, we recall some of the definitions and basic tools that will be used in the remaining sections, such as the syntax of code-based games, functional encryption and the assumptions.

2.1   Notation and Conventions

Let $\mathbb{N}$ denote the set of natural numbers. If $n \in \mathbb{N}$, then $\{0,1\}^n$ denotes the set of n-bit strings, and $\{0,1\}^*$ is the set of all bit strings. The empty string is denoted ε. More generally, if S is a set, then S^n is the set of n-tuples of elements of S, $S^{\leq n}$ is the set of tuples of length at most n. If x is a string then $|x|$ denotes its length, and if S is a set then $|S|$ denotes its size. If S is finite, then $x \xleftarrow{R} S$ denotes the assignment to x of an element chosen uniformly at random from S. If $\mathcal{A}$ is an algorithm, then $y \leftarrow \mathcal{A}(x)$ denotes the assignment to y of the output of $\mathcal{A}$ on input x, and if $\mathcal{A}$ is randomized, then $y \xleftarrow{R} \mathcal{A}(x)$ denotes that the output of an execution of $\mathcal{A}(x)$ with fresh coins is assigned to y. Unless otherwise indicated, an algorithm may be randomized. "PT" stands for polynomial time and "PTA" for polynomial-time algorithm or adversary. We denote by $\lambda \in \mathbb{N}$ the security parameter. A function $\nu : \mathbb{N} \to [0,1]$ is said to be *negligible* if for every $c \in \mathbb{N}$ there exists a $\lambda_c \in \mathbb{N}$ such that $\nu(\lambda) \leq \lambda^{-c}$ for all $\lambda > \lambda_c$, and it is said to be *overwhelming* if the function $|1 - \nu(\lambda)|$ is negligible.

Let $\mathbf{e}_i$ be the vector with all 0 but one 1 in the i-th position.

Code-Based Games. We use the code-based game-playing [BR06] to define the security notions. In such games, there exist procedures for initialization (**initialize**) and finalization (**Finalize**) and procedures to respond to adversary oracle queries. A game G is executed with an adversary $\mathcal{A}$ as follows. First, **initialize** executes and its outputs are the inputs to $\mathcal{A}$. Then $\mathcal{A}$ executes, its oracle queries being answered by the corresponding procedures of G. When $\mathcal{A}$ terminates, its output becomes the input to the **Finalize** procedure. The output of the latter, denoted $G(\mathcal{A})$, is called the output of the game, and "$G(\mathcal{A}) = y$" denotes the

event that the output takes a value y. Boolean flags are assumed initialized to false. Games G_i, G_j are *identical until* bad if their code differs only in statements that follow the setting of bad to true.

2.2 Public-Key Encryption

Definition 1 (Public-Key Encryption Scheme). *A* public-key encryption *(PKE) scheme $\mathcal{E}$ is a tuple $\mathcal{E} = (\mathsf{Setup}, \mathsf{Encrypt}, \mathsf{Decrypt})$ of 3 algorithms:*

1. $\mathsf{Setup}(1^\lambda)$ *outputs* public *and* secret *keys* $(\mathsf{pk}, \mathsf{sk})$ *for security parameter λ;*

2. $\mathsf{Encrypt}(\mathsf{pk}, m)$, *on input public key pk and message m in the allowed message space, outputs* ciphertext Ct;

3. $\mathsf{Decrypt}(\mathsf{sk}, \mathsf{Ct})$ *on input secret key sk and ciphertext Ct, outputs messages m'.*

In addition we make the following correctness *requirement: for all $(\mathsf{pk}, \mathsf{sk}) \leftarrow \mathsf{Setup}(1^\lambda)$, all messages m and ciphertexts $\mathsf{Ct} \leftarrow \mathsf{Encrypt}(\mathsf{pk}, m)$, we have that $\mathsf{Decrypt}(\mathsf{sk}, \mathsf{Ct}) = m$ except with negligible probability.*

We often also allow public-key encryption schemes to additionally depend on explicit public parameters params (randomly generated in an initial phase and shared across multiple instances of the PKE scheme) on which all of Setup, $\mathsf{Encrypt}$, and $\mathsf{Decrypt}$ are allowed to depend. Examples include the description of a group $\mathbb{G}$ with its generator g. We will often omit them in the descriptions of generic constructions from PKE schemes.

Indistinguishability-Based Security. We define *security against chosen-plaintext attacks* (IND-CPA *security*, for short) for a PKE scheme $\mathcal{E} = (\mathsf{Setup}, \mathsf{Encrypt}, \mathsf{Decrypt})$ via the security game depicted on Figure 1. Then, we say that $\mathcal{E}$ is secure against chosen-plaintext attacks (IND-CPA secure, for short) if

$$\left| \Pr[\mathsf{Exp}^{\text{ind-cpa-0}}_{\mathcal{E}, \lambda}(\mathcal{A}) = 1] - \Pr[\mathsf{Exp}^{\text{ind-cpa-1}}_{\mathcal{E}, \lambda}(\mathcal{A}) = 1] \right| = \mathsf{negl}(\lambda).$$

We also define *selective security against chosen-plaintext attacks* (s-IND-CPA *security*, for short) when the challenge messages m_0^* and m_1^* have to be chosen before hand. Actually, in this case, the procedures **initialize** and **LR** can be merged into an **initialize** procedure that outputs both the public key pk and the challenge ciphertext Ct^*.

2.3 Functional Encryption

Following Boneh *et al.* [BSW11], we start by defining the notion of functionality and then that of functional encryption scheme $\mathcal{FE}$ for functionality F.

Definition 2 (Functionality). *A functionality F defined over (K, X) is a function $F : K \times X \to \Sigma \cup \{\bot\}$ where K is the* key space, *X is the* message space *and Σ is the* output space *and $\bot$ is a special string not contained in Σ. Notice that the functionality is undefined for when either the key is not in the key space or the message is not in the message space.*

Game $\mathsf{Exp}^{\text{ind-cpa-}b}_{\mathcal{E},\lambda}(\mathcal{A})$		**Game** $\mathsf{Exp}^{\text{s-ind-cpa-}b}_{\mathcal{E},\lambda}(\mathcal{A})$
proc Initialize(λ)	**proc Finalize**(b')	**proc Initialize**$(\lambda, \mathsf{m}_0^*, \mathsf{m}_1^*)$
$(\mathsf{pk}, \mathsf{sk}) \overset{R}{\leftarrow} \mathsf{Setup}(1^\lambda)$	Return $(b' = b)$	$(\mathsf{pk}, \mathsf{sk}) \overset{R}{\leftarrow} \mathsf{Setup}(1^\lambda)$
Return pk		Return pk
proc LR$(\mathsf{m}_0^*, \mathsf{m}_1^*)$		**proc LR**$()$
$\mathsf{Ct}^* \overset{R}{\leftarrow} \mathsf{Encrypt}(mpk, \mathsf{m}_b^*)$		$\mathsf{Ct}^* \overset{R}{\leftarrow} \mathsf{Encrypt}(\mathsf{pk}, \mathsf{m}_b^*)$
Return Ct^*		Return Ct^*

Fig. 1. Games $\mathsf{Exp}^{\text{ind-cpa-}b}_{\mathcal{E},\lambda}(\mathcal{A})$ and $\mathsf{Exp}^{\text{s-ind-cpa-}b}_{\mathcal{E},\lambda}(\mathcal{A})$ define IND-CPA and s-IND-CPA security (respectively) of $\mathcal{E}$. The procedure **Finalize** is common to both games, which differ in their **initialize** and **LR** procedures.

Definition 3 (Functional Encryption Scheme). *A functional encryption (FE) scheme $\mathcal{FE}$ for functionality F is a tuple $\mathcal{FE} = (\mathsf{Setup}, \mathsf{KeyDer}, \mathsf{Encrypt}, \mathsf{Decrypt})$ of 4 algorithms:*

1. $\mathsf{Setup}(1^\lambda)$ *outputs* public *and* master secret *keys* (mpk, msk) *for security parameter λ;*
2. $\mathsf{KeyDer}(msk, k)$, *on input a master secret key msk and key $k \in K$ outputs secret key sk_k;*
3. $\mathsf{Encrypt}(mpk, x)$, *on input public key mpk and message $x \in X$ outputs ciphertext Ct;*
4. $\mathsf{Decrypt}(mpk, \mathsf{Ct}, \mathsf{sk}_k)$ *outputs $y \in \Sigma \cup \{\bot\}$.*

We make the following correctness *requirement: for all $(mpk, msk) \leftarrow \mathsf{Setup}(1^\lambda)$, all $k \in K$ and $m \in M$, for $\mathsf{sk}_k \leftarrow \mathsf{KeyDer}(msk, k)$ and $\mathsf{Ct} \leftarrow \mathsf{Encrypt}(mpk, m)$, we have that $\mathsf{Decrypt}(mpk, \mathsf{Ct}, \mathsf{sk}_k) = F(k, m)$ whenever $F(k, m) \neq \bot^2$, except with negligible probability.*

Indistinguishability-Based Security. For a functional encryption scheme $\mathcal{FE} = (\mathsf{Setup}, \mathsf{KeyDer}, \mathsf{Encrypt}, \mathsf{Decrypt})$ for functionality F, defined over (K, X), we define *security against chosen-plaintext attacks* (IND-CPA *security*, for short) via the security game depicted on Figure 2. Then, we say that $\mathcal{FE}$ is secure against chosen-plaintext attacks (IND-CPA secure, for short) if

$$\left| \Pr[\mathsf{Exp}^{\text{ind-cpa-0}}_{\mathcal{FE},\lambda}(\mathcal{A}) = 1] - \Pr[\mathsf{Exp}^{\text{ind-cpa-1}}_{\mathcal{FE},\lambda}(\mathcal{A}) = 1] \right| = \mathsf{negl}(\lambda).$$

We also define *selective security against chosen-plaintext attacks* (s-IND-CPA *security*, for short) when the challenge messages m_0^* and m_1^* have to be chosen before hand.

---

[2] See [BO12, ABN10] for a discussion about this condition.

Game $\mathsf{Exp}^{\mathrm{ind\text{-}cpa\text{-}}b}_{\mathcal{FE},\lambda}(\mathcal{A})$		**Game** $\mathsf{Exp}^{\mathrm{s\text{-}ind\text{-}cpa\text{-}}b}_{\mathcal{FE},\lambda}(\mathcal{A})$
proc Initialize(λ)	**proc LR**$(\mathsf{m}_0^*, \mathsf{m}_1^*)$	**proc Initialize**$(\lambda, \mathsf{m}_0^*, \mathsf{m}_1^*)$
$(mpk, msk) \xleftarrow{R} \mathsf{Setup}(1^\lambda)$	$\mathsf{Ct}^* \xleftarrow{R} \mathsf{Encrypt}(mpk, \mathsf{m}_b^*)$	$(mpk, msk) \xleftarrow{R} \mathsf{Setup}(1^\lambda)$
$V \leftarrow \emptyset$	Return Ct^*	$V \leftarrow \emptyset$
Return mpk		Return mpk
	proc Finalize(b')	
proc KeyDer(k)	**if** $\exists k \in V$ such that	
$V \leftarrow V \cup \{k\}$	$\quad F(k, \mathsf{m}_0^*) \neq F(k, \mathsf{m}_1^*)$	**proc LR**$()$
$\mathsf{sk}_k \xleftarrow{R} \mathsf{KeyDer}(msk, k)$	$\quad$ **then return** false	$\mathsf{Ct}^* \xleftarrow{R} \mathsf{Encrypt}(mpk, \mathsf{m}_b^*)$
Return sk_k	Return $(b' = b)$	Return Ct^*

Fig. 2. Games $\mathsf{Exp}^{\mathrm{ind\text{-}cpa\text{-}}b}_{\mathcal{FE},\lambda}(\mathcal{A})$ and $\mathsf{Exp}^{\mathrm{s\text{-}ind\text{-}cpa\text{-}}b}_{\mathcal{FE},\lambda}(\mathcal{A})$ define IND-CPA and s-IND-CPA security (respectively) of $\mathcal{FE}$. The procedures **KeyDer** and **Finalize** are common to both games, which differ in their **initialize** and **LR** procedures.

Inner-Product Functionality. In this paper we are mainly interested in the *inner-product functionality* over $\mathbb{Z}_p$ (IP, for short) defined in the following way. It is a family of functionalities with key space K_ℓ and message space X_ℓ both consisting of vectors in $\mathbb{Z}_p$ of length ℓ: for any $k \in K_\ell, x \in X_\ell$ the functionality $\mathsf{IP}_\ell(k, x) = \langle k, x \rangle \mod p$. When it is clear from the context we remove the reference to the length ℓ.

3   Inner-Product from DDH

In this section, we present our first functional encryption scheme for the inner-product functionality whose security can be based on the plain DDH assumption.

The Decisional Diffie-Hellman assumption. Let $\mathsf{GroupGen}$ be a probabilistic polynomial-time algorithm that takes as input a security parameter 1^λ, and outputs a triplet $(\mathbb{G}, p, g)$ where $\mathbb{G}$ is a group of order p that is generated by $g \in \mathbb{G}$, and p is an λ-bit prime number. Then, the *Decisional Diffie-Hellman (DDH) assumption* states that the tuples (g, g^a, g^b, g^{ab}) and (g, g^a, g^b, g^c) are computationally indistinguishable, where $(\mathbb{G}, p, g) \leftarrow \mathsf{GroupGen}(1^\lambda)$, and $a, b, c \in \mathbb{Z}_p$ are chosen independently and uniformly at random.

Construction 1 (DDH-IP Scheme). *We define our functional encryption scheme for the inner-product functionality* $\mathsf{IP} = (\mathsf{Setup}, \mathsf{KeyDer}, \mathsf{Encrypt}, \mathsf{Decrypt})$ *as follows:*
- $\mathsf{Setup}(1^\lambda, 1^\ell)$ *samples* $(\mathbb{G}, p, g) \leftarrow \mathsf{GroupGen}(1^\lambda)$ *and* $\mathbf{s} = (s_1, \ldots, s_\ell) \leftarrow \mathbb{Z}_p^\ell$, *and sets* $mpk = (h_i = g^{s_i})_{i \in [\ell]}$ *and* $msk = \mathbf{s}$. *The algorithm returns the pair* (mpk, msk);
- $\mathsf{Encrypt}(mpk, \mathbf{x})$ *on input master public key* mpk *and message* $\mathbf{x} = (x_1, \ldots, x_\ell) \in \mathbb{Z}_p^\ell$, *chooses a random* $r \leftarrow \mathbb{Z}_p$ *and computes* $\mathsf{ct}_0 = g^r$ *and, for each* $i \in [\ell]$, $\mathsf{ct}_i = h_i^r \cdot g^{x_i}$. *Then the algorithm returns the ciphertext* $\mathsf{Ct} = (\mathsf{ct}_0, (\mathsf{ct}_i)_{i \in [\ell]})$;

- KeyDer$(msk, \mathbf{y})$ *on input master secret key msk and vector* $\mathbf{y} = (y_1, \ldots, y_\ell) \in \mathbb{Z}_p^\ell$, *computes and outputs secret key* $\mathsf{sk}_\mathbf{y} = \langle \mathbf{y}, \mathbf{s} \rangle$;
- Decrypt$(mpk, \mathsf{Ct}, \mathsf{sk}_\mathbf{y})$ *on input master public key mpk, ciphertext* $\mathsf{Ct} = (\mathsf{ct}_0, (\mathsf{ct}_i)_{i \in [\ell]})$ *and secret key* $\mathsf{Sk}_\mathbf{y}$ *for vector* $\mathbf{y}$, *returns the discrete logarithm in basis* g *of*

$$\prod_{i \in [\ell]} \mathsf{ct}_i^{y_i} / \mathsf{ct}_0^{\mathsf{sk}_\mathbf{y}}.$$

Correctness. For all $(mpk, msk) \leftarrow \mathsf{Setup}(1^\lambda, 1^\ell)$, all $\mathbf{y} \in \mathbb{Z}_p^\ell$ and $\mathbf{x} \in \mathbb{Z}_p^\ell$, for $\mathsf{sk}_\mathbf{y} \leftarrow \mathsf{KeyDer}(msk, \mathbf{y})$ and $\mathsf{Ct} \leftarrow \mathsf{Encrypt}(mpk, \mathbf{x})$, we have that

$$\begin{aligned}
\mathsf{Decrypt}(mpk, \mathsf{Ct}, \mathsf{sk}) &= \frac{\prod_{i \in [\ell]} \mathsf{ct}_i^{y_i}}{\mathsf{ct}_0^{\mathsf{sk}_\mathbf{y}}} = \frac{\prod_{i \in [\ell]} (g^{s_i r + x_i})^{y_i}}{g^{r(\sum_{i \in [\ell]} y_i s_i)}} \\
&= g^{\sum_{i \in [\ell]} y_i s_i r + \sum_{i \in [\ell]} y_i x_i - r(\sum_{i \in [\ell]} y_i s_i)} \\
&= g^{\sum_{i \in [\ell]} y_i x_i} = g^{\langle \mathbf{x}, \mathbf{y} \rangle}.
\end{aligned}$$

The above scheme limits the expressiveness of the functionality that can be computed because in order to recover the final inner-product value a discrete logarithm computation must take place. In the next section, in order to overcome this limitation and to generalize to other settings we will present a generic scheme whose security can be based on the semantic security of the underlying public-key encryption scheme under randomness reuse. Before moving to the generic scheme and its proof of security, we sketch below the proof of security for the above IP scheme to offer to the reader useful intuitions that will be reused in the proof of security of our generic functional encryption scheme for the inner-product functionality.

Theorem 1. *Under the DDH assumption, the above* IP *scheme is* s-IND-CPA.

Proof Sketch. For the sake of contradiction, suppose that there exists an adversary $\mathcal{A}$ that breaks s-IND-CPA security of our IP scheme with non-negligible advantage. Then, we construct a simulator $\mathcal{B}$ that given in input an instance of the DDH assumption, (g, g^a, g^b, g^c) where c is either ab or uniformly random, breaks it by using $\mathcal{A}$.

If the challenge messages $\mathbf{x}_0$ and $\mathbf{x}_1$ are different, there exists a vector in the message space for which the key shouldn't be known by the adversary ($\mathbf{x}_1 - \mathbf{x}_0$ is one of them).

To generate the master public key, $\mathcal{B}$ first generates secret keys for a basis $(\mathbf{z}_i)$ of $(\mathbf{x}_1 - \mathbf{x}_0)^\perp$. Setting implicitly a as secret key for $\mathbf{x}_1 - \mathbf{x}_0$, $\mathcal{B}$ generates the master public key using g^a. Actually, once group elements are generated for the basis $(\mathbf{y}_i)$ completed with $(\mathbf{x}_1 - \mathbf{x}_0)$, one can find the public key, for the canonical basis.

To generate the challenge ciphertext, $\mathcal{B}$ chooses a random bit μ and using g^b and g^c, generates a ciphertext for message $\mathbf{x}_\mu$.

Finally, notice that by the constraints of the s-IND-CPA security game, $\mathcal{A}$ is allowed to ask only secret keys for vectors in the vector sub-space generated by the $\mathbf{z}_i$'s, and thus orthogonal to $\mathbf{x}_1 - \mathbf{x}_0$. For those vectors, $\mathcal{B}$ will be able to generate the corresponding secret keys.

Now, if $c = ab$, then the challenge ciphertext is a well-distributed ciphertext of the message $\mathbf{x}_\mu$. On the other hand, if c is random, the challenge ciphertext encrypts message $u\mathbf{x}_0 + (1 - u)\mathbf{x}_1$ where u is uniformly distributed and in this case μ is information theoretically hidden: for all $\mathsf{sk_y}$ asked, since $\mathbf{y}$ is orthogonal to $\mathbf{x}_1 - \mathbf{x}_0$, and thus $\langle \mathbf{x}_0, \mathbf{y} \rangle = \langle \mathbf{x}_1, \mathbf{y} \rangle$, while the ciphertexts decrypt to the inner products

$$
\begin{aligned}
\langle u\mathbf{x}_0 + (1 - u)\mathbf{x}_1, \mathbf{y} \rangle &= u\langle \mathbf{x}_0, \mathbf{y} \rangle + (1 - u)\langle \mathbf{x}_1, \mathbf{y} \rangle \\
&= u\langle \mathbf{x}_\mu, \mathbf{y} \rangle + (1 - u)\langle \mathbf{x}_\mu, \mathbf{y} \rangle \\
&= \langle \mathbf{x}_\mu, \mathbf{y} \rangle \ .
\end{aligned}
$$

4   A Generic Inner-Product Encryption Scheme

In this section, we present a generic functional encryption scheme for the inner-product functionality based on any public-key encryption scheme that possesses some specific properties. The security of this scheme can be then based solely on the semantic security of the underlying public-key encryption scheme.

We start with a public-key scheme $\mathcal{E} = (\mathsf{Setup}, \mathsf{Encrypt}, \mathsf{Decrypt})$ whose secret keys are elements of a group $(G, +, 0_G)$, public keys are elements of group $(H, \cdot, 1_H)$, and messages are elements of $\mathbb{Z}_q$ for some q.

In addition, we rely on three other special properties of the cryptosystem. First, we require that it remains secure when we reuse randomness for encrypting under different keys [BBS03]. Then, we require some homomorphic operation, on the keys and on the ciphertexts. More specifically:

Randomness Reuse. We require the ciphertexts to consist of two parts ct_0 and ct_1. The first part ct_0 corresponds to some *commitment* $\mathsf{C}(r)$ of the randomness r used for the encryption. The second part ct_1 is the *encryption* $\mathsf{E}(\mathsf{pk}, x; r)$ in a group $(I, \cdot, 1_I)$ of the message x under public key pk and randomness r.

Then, we say that a PKE has *randomness reuse* (RR, for short) if $\mathsf{E}(\mathsf{pk}, x; r)$ is *efficiently computed* given the triple (x, pk, r), or the triple $(x, \mathsf{sk}, \mathsf{C}(r))$ where sk is a secret key corresponding to pk. (In [BBS03], this property is also called *reproducibility* and guarantees that it is secure to reuse randomness when encrypting under several independent public keys. The idea of randomness reuse was first considered by [Kur02]);

Linear Key Homomorphism. We say that a PKE has *linear key homomorphism* (LKH, for short) if for any two secret keys $\mathsf{sk}_1, \mathsf{sk}_2 \in G$ and any $y_1, y_2 \in \mathbb{Z}_q$, the component-wise G-linear combination formed by $y_1\mathsf{sk}_1 + y_2\mathsf{sk}_2$ can be computed efficiently only using public parameters, the secret keys sk_1 and sk_2 and the coefficients y_1 and y_2. And this combination $y_1\mathsf{sk}_1 + y_2\mathsf{sk}_2$

also functions as a secret key, with a corresponding public key that can be computed efficiently from y_1, y_2 and the public keys pk_1 and pk_2 of sk_1 and sk_2 respectively, fixing the same public parameters, for example $\mathsf{pk}_1^{y_1} \cdot \mathsf{pk}_2^{y_2}$.

Linear Ciphertext Homomorphism Under Shared Randomness. We say that a PKE has *linear ciphertext homomorphism under shared randomness* (LCH, for short) if $\mathsf{E}(\mathsf{pk}_1\mathsf{pk}_2, x_1 + x_2; r)$ can be efficiently computed from $\mathsf{E}(\mathsf{pk}_1, x_1; r), \mathsf{E}(\mathsf{pk}_2, x_2; r)$ and $\mathsf{C}(r)$, for example $\mathsf{E}(\mathsf{pk}_1\mathsf{pk}_2, x_1 + x_2; r) = \mathsf{E}(\mathsf{pk}_1, x_1; r) \cdot \mathsf{E}(\mathsf{pk}_2, x_2; r)$. (It doesn't actually have to be $\mathsf{pk}_1\mathsf{pk}_2$ but a public key corresponding to $\mathsf{sk}_1 + \mathsf{sk}_2$ if sk_1 and sk_2 are secret keys corresponding to pk_1 and pk_2).

4.1 Construction

Construction 2 (PKE-IP Scheme). *Let $\mathcal{E} = (\mathsf{Setup}, \mathsf{Encrypt}, \mathsf{Decrypt})$ be a PKE scheme with the properties defined above, we define our* functional encryption scheme for the inner-product functionality $\mathit{IP} = (\mathsf{Setup}, \mathsf{KeyDer}, \mathsf{Encrypt}, \mathsf{Decrypt})$ *as follows.*

- $\mathsf{Setup}(1^\lambda, 1^\ell)$ *calls $\mathcal{E}$'s key generation algorithm to generate ℓ independent $(\mathsf{sk}_1, \mathsf{pk}_1), \ldots, (\mathsf{sk}_\ell, \mathsf{pk}_\ell)$ pairs, sharing the same public parameters* params. *Then, the algorithm returns $mpk = (\mathsf{params}, \mathsf{pk}_1, \ldots, \mathsf{pk}_\ell)$ and $msk = (\mathsf{sk}_1, \ldots, \mathsf{sk}_\ell)$.*

- $\mathsf{KeyDer}(msk, \mathbf{y})$ *on input master secret key msk and a vector $\mathbf{y} = (y_1, \ldots, y_\ell) \in \mathbb{Z}_q^\ell$, computes $\mathsf{sk}_\mathbf{y}$ as an G-linear combination of $(\mathsf{sk}_1, \ldots, \mathsf{sk}_n)$ with coefficients $(y_1, \ldots, y_\ell)$, namely $\mathsf{sk}_\mathbf{y} = \sum_{i \in [\ell]} y_i \cdot \mathsf{sk}_i$.*

- $\mathsf{Encrypt}(mpk, \mathbf{x})$ *on input master public key mpk and message $\mathbf{x} = (x_1, \ldots, x_\ell) \in \mathbb{Z}_q^\ell$, chooses shared randomness r in the randomness space of $\mathcal{E}$, and computes $\mathsf{ct}_0 = \mathcal{E}.\mathsf{C}(r)$ and $\mathsf{ct}_i = \mathcal{E}.\mathsf{E}(\mathsf{pk}_i, x_i; r)$. Then the algorithm returns the ciphertext $\mathsf{Ct} = (\mathsf{ct}_0, (\mathsf{ct}_i)_{i \in [\ell]})$.*

- $\mathsf{Decrypt}(mpk, \mathsf{Ct}, \mathsf{sk}_\mathbf{y})$ *on input master public key mpk, ciphertext $\mathsf{Ct} = (\mathsf{ct}_0, (\mathsf{ct}_i)_{i \in [\ell]})$, and secret key $\mathsf{Sk}_\mathbf{y}$ for vector $\mathbf{y} = (y_1, \ldots, y_\ell)$, returns the output of $\mathcal{E}.\mathsf{Decrypt}(\mathsf{Sk}_\mathbf{y}, (\mathsf{ct}_0, \prod_{i \in [\ell]} \mathsf{ct}_i^{y_i}))$.*

Correctness. For all $(mpk, msk) \leftarrow \mathsf{Setup}(1^\lambda, 1^\ell)$, all $\mathbf{y} \in \mathbb{Z}_q^\ell$ and $\mathbf{x} \in \mathbb{Z}_q^\ell$, for $\mathsf{sk}_\mathbf{y} \leftarrow \mathsf{KeyDer}(msk, \mathbf{y})$ and $\mathsf{Ct} \leftarrow \mathsf{Encrypt}(mpk, \mathbf{x})$, we have that

$$
\begin{aligned}
\mathsf{Decrypt}(mpk, \mathsf{Ct}, \mathsf{sk}_\mathbf{y}) &= \mathcal{E}.\mathsf{Decrypt}(\mathsf{Sk}_\mathbf{y}, (\mathsf{ct}_0, \prod_{i \in [\ell]} \mathsf{ct}_i^{y_i})) \\
&= \mathcal{E}.\mathsf{Decrypt}(\mathsf{Sk}_\mathbf{y}, (\mathsf{ct}_0, \prod_{i \in [\ell]} \mathcal{E}.\mathsf{E}(\mathsf{pk}_i, x_i; r)^{y_i})) \\
&= \mathcal{E}.\mathsf{Decrypt}(\mathsf{Sk}_\mathbf{y}, (\mathsf{ct}_0, \mathcal{E}.\mathsf{Encrypt}(\prod_{i \in [\ell]} \mathsf{pk}_i^{y_i}, \sum_{i \in [\ell]} y_i x_i; r))) \\
&= \sum_{i \in [\ell]} y_i x_i \ .
\end{aligned}
$$

by the LCH property. Finally, note that the decryption is allowed because $(\mathsf{Sk_y}, \prod_{i \in [\ell]} \mathsf{pk}_i^{y_i})$ is a valid key pair, due to the LKH property.

4.2   Proof of Security

In this section, we prove the following theorem:

Theorem 2. *If the underlying PKE $\mathcal{E}$ is* s-IND-CPA, *even under randomness-reuse, and satisfies the LKH and LCH properties, then the* IP *scheme of Section 4.1 is* s-IND-CPA.

Proof. This proof follows the intuition provided in the proof sketch of Theorem 1. To prove the security of our scheme we will show that the s-IND-CPA game is indistinguishable from a game where the challenge ciphertext encrypts a random combination of the challenge messages whose coefficients sum up to one. Thus, the challenge ciphertext decrypts to the expected values and information theoretically hides the challenge bit.

   More specifically, given an adversary $\mathcal{A}$ that breaks the s-IND-CPA security of our IP scheme with non-negligible probability ε, we construct an adversary $\mathcal{B}$ that breaks the s-IND-CPA security of the underlying PKE scheme $\mathcal{E}$ with comparable probability.

   $\mathcal{B}$ starts by picking a random element a in the full message space of the underlying PKE $\mathcal{E}$, and sends challenge messages 0 and a to the challenger $\mathcal{C}$ of PKE security game. $\mathcal{C}$ answers by sending an encryption $\mathsf{Ct} = (\mathsf{ct}_0, \mathsf{ct}_1)$ of either 0 or a and public key pk.

   $\mathcal{B}$ then invokes $\mathcal{A}$ on input the security parameter and gets two different challenge messages in output, namely $(\mathbf{x}_i = (x_{i,1}, \ldots, x_{i,\ell}))_{i \in \{0,1\}}$ both in $\mathbb{Z}_q^\ell$.

   Recall that, by the constraints of security game, the adversary can only issue secret key queries for vectors $\mathbf{y}$ such that $\langle \mathbf{x}_0, \mathbf{y} \rangle = \langle \mathbf{x}_1, \mathbf{y} \rangle$. Thus, we have that $\langle \mathbf{y}, \mathbf{x}_1 - \mathbf{x}_0 \rangle = 0$ meaning that $\mathbf{y}$ is in the vector space defined by $(\mathbf{x}_1 - \mathbf{x}_0)^\perp$.

   Then, $\mathcal{B}$ generates the view for $\mathcal{A}$ in the following way:

Public Key. To generate master public key mpk, $\mathcal{B}$ does the following. First, $\mathcal{B}$ finds a basis $(\mathbf{z}_1, \mathbf{z}_2, \ldots, \mathbf{z}_{\ell-1})$ of $(\mathbf{x}_1 - \mathbf{x}_0)^\perp$. Then we can write the canonical vectors in this basis: for $i \in [\ell], j \in [\ell-1]$, there exist $\lambda_{i,j} \in \mathbb{Z}_q$ and $\alpha_i \in \mathbb{Z}_q$ such that:

$$\mathbf{e}_i = \alpha_i(\mathbf{x}_1 - \mathbf{x}_0) + \sum_{j \in [\ell-1]} \lambda_{i,j}\mathbf{z}_j. \tag{1}$$

Then, for $j \in [\ell-1]$, $\mathcal{B}$ sets $(\mathsf{pk}_{\mathbf{z}_j}, \mathsf{sk}_{\mathbf{z}_j}) = \mathcal{E}.\mathsf{Setup}(1^\lambda)$, and for $i \in [\ell]$,

$$\gamma_i = \prod_{j \in [\ell-1]} \mathsf{pk}_{\mathbf{z}_j}^{\lambda_{i,j}} \quad \text{and} \quad \mathsf{pk}_i = \mathsf{pk}^{\alpha_i}\gamma_i.$$

Eventually, $\mathcal{B}$ invokes $\mathcal{A}$ on input $mpk = (\mathsf{pk}_i)_{i \in [\ell]}$.
Notice that, $\mathcal{B}$ is implicitly setting $\mathsf{sk}_i = \alpha_i\mathsf{sk} + \sum_{j \in [\ell-1]} \lambda_{i,j}\mathsf{sk}_{\mathbf{z}_j}$ because of the LKH property, where sk is the secret key corresponding to pk, which is unknown to $\mathcal{B}$.

Challenge Ciphertext. $\mathcal{B}$ computes the challenge ciphertext Ct^* as follows. $\mathcal{B}$ randomly picks $b \xleftarrow{R} \{0,1\}$, computes $\mathcal{E}.\mathsf{E}(\gamma_i,0;r)$ from ct_0 and $\sum_{j\in[\ell-1]} \lambda_{i,j}$ $\mathsf{sk}_{\mathbf{z}_j}$ and $\mathcal{E}.\mathsf{E}(1_H,x_{b,i};r)$ from secret key 0_G and ct_0, by randomness reuse. $\mathcal{B}$ then sets

$$\mathsf{ct}_0^* = \mathsf{ct}_0 \quad \text{and} \quad (\mathsf{ct}_i^* = \mathsf{ct}_1^{\alpha_i} \cdot \mathcal{E}.\mathsf{E}(\gamma_i,0;r) \cdot \mathcal{E}.\mathsf{E}(1_H,x_{b,i};r))_{i\in[\ell]} \ ,$$

Then the algorithm returns the challenge ciphertext $\mathsf{Ct}^* = (\mathsf{ct}_0^*, (\mathsf{ct}_i^*)_{i\in[\ell]})$.
Secret Keys. To generate a secret key for vector $\mathbf{y}$, $\mathcal{B}$ computes $\mathsf{sk}_{\mathbf{y}}$ as

$$\mathsf{sk}_{\mathbf{y}} = \sum_{j\in[\ell-1]} \left(\sum_{i\in[\ell]} y_i\lambda_{i,j} \right) \mathsf{sk}_{\mathbf{z}_j}$$

At the end of the simulation, if $\mathcal{A}$ correctly guesses b, then $\mathcal{B}$ returns 0 ($\mathcal{B}$ guesses that $\mathcal{C}$ encrypted 0), else $\mathcal{B}$ returns 1 ($\mathcal{B}$ guesses that $\mathcal{C}$ encrypted a). This concludes the description of adversary $\mathcal{B}$.

It remains to verify that $\mathcal{B}$ correctly simulates $\mathcal{A}$'s environment.

First see that the master public key is well distributed, because we are just applying a change of basis to a well distributed master public key. Now it holds that $\alpha_i = \frac{x_{1,i}-x_{0,i}}{\|\mathbf{x}_1-\mathbf{x}_0\|^2}$ because

$$\begin{aligned} x_{1,i} - x_{0,i} &= \langle \mathbf{x}_1 - \mathbf{x}_0, \mathbf{e}_i \rangle \\ &= \alpha_i\|\mathbf{x}_1-\mathbf{x}_0\|^2 + \sum_{j\in[\ell-1]} \lambda_{i,j}\langle \mathbf{x}_1-\mathbf{x}_0, \mathbf{z}_j \rangle \\ &= \alpha_i\|\mathbf{x}_1-\mathbf{x}_0\|^2 \ . \end{aligned}$$

Now recall that a vector $\mathbf{y}$ satisfying the security game constraints is such that $\langle \mathbf{y}, \mathbf{x}_0 \rangle = \langle \mathbf{y}, \mathbf{x}_1 \rangle$, so

$$\sum_{i\in[\ell]} y_i\alpha_i = \sum_{i\in[\ell]} y_i \frac{x_{1,i}-x_{0,i}}{\|\mathbf{x}_1-\mathbf{x}_0\|^2} = 0$$

which in turn implies that a secret key $\mathsf{sk}_{\mathbf{y}}$ for the vector $\mathbf{y}$ is distributed as

$$\begin{aligned} \mathsf{sk}_{\mathbf{y}} = \sum_{i\in[\ell]} y_i\mathsf{sk}_i &= \sum_{i\in[\ell]} y_i\alpha_i\mathsf{sk} + \sum_{i\in[\ell]}\sum_{j\in[\ell-1]} y_i\lambda_{i,j}\mathsf{sk}_{\mathbf{z}_j} \\ &= \sum_{j\in[\ell-1]} \left(\sum_{i\in[\ell]} y_i\lambda_{i,j} \right) \mathsf{sk}_{\mathbf{z}_j} \end{aligned}$$

On the other hand, if $\mathcal{A}$ asks for a secret key for some vector $\mathbf{y} \notin (\mathbf{x}_1 - \mathbf{x}_0)^{\perp}$, $\mathcal{B}$ would need to know sk in order to generate a correct secret key for $\mathbf{y}$.

Now, we have to analyze the following two cases, depending on which message was encrypted by $\mathcal{C}$ in the challenge ciphertext:

1. $\mathcal{C}$ encrypted 0. Then, the challenge ciphertext Ct^* for message $\mathbf{x}_b$ is distributed as

$$\mathsf{ct}_0^* = \mathsf{ct}_0$$

and

$$\begin{aligned}
\mathsf{ct}_i^* &= \mathcal{E}.\mathsf{E}(\mathsf{pk}, 0; r)^{\alpha_i} \cdot \mathcal{E}.\mathsf{E}(\gamma_i, 0; r) \cdot \mathcal{E}.\mathsf{E}(1_H, x_{b,i}; r) \\
&= \mathcal{E}.\mathsf{E}(\mathsf{pk}_i, x_{b,i}; r) ,
\end{aligned}$$

thanks to the LCH property, and then as in the real game.

Thus, in this case, $\mathcal{B}$ generates a view identical to that $\mathcal{A}$ would see in the real game. Hence, the advantage of $\mathcal{B}$ in this game is ε, the same advantage as $\mathcal{A}$ against s-IND-CPA of IP when 0 has been encrypted.

2. $\mathcal{C}$ encrypted a. First, in Equation 1, we have $\alpha_i = (x_{1,i} - x_{0,i})/\|\mathbf{x}_1 - \mathbf{x}_0\|^2$. Let us analyze the distribution of the challenge ciphertext in this case. We have $\mathsf{ct}_0^* = \mathsf{ct}_0$ and

$$\begin{aligned}
\mathsf{ct}_i^* &= \mathcal{E}.\mathsf{E}(\mathsf{pk}, a; r)^{\alpha_i} \cdot \mathcal{E}.\mathsf{E}(\gamma_i, 0; r) \cdot \mathcal{E}.\mathsf{E}(1_H, x_{b,i}; r) \\
&= \mathcal{E}.\mathsf{E}(\mathsf{pk}_i, x_{b,i} + \alpha_i a; r) \\
&= \mathcal{E}.\mathsf{E}(\mathsf{pk}_i, \hat{x}_i; r),
\end{aligned}$$

thanks to the LCH property, where $\hat{x}_i$ is defined as follows:

$$\begin{aligned}
\hat{x}_i = x_{b,i} + \alpha_i a &= \frac{a}{\|\mathbf{x}_1 - \mathbf{x}_0\|^2}(x_{1,i} - x_{0,i}) + x_{b,i} \\
&= \frac{a}{\|\mathbf{x}_1 - \mathbf{x}_0\|^2}(x_{1,i} - x_{0,i}) + x_{0,i} + b(x_{1,i} - x_{0,i}).
\end{aligned}$$

Let us set $u = a/\|\mathbf{x}_1 - \mathbf{x}_0\|^2 + b$, which is a random value in the full message space of $\mathcal{E}$, given that a is random in the same space, then $\hat{x}_i = ux_{1,i} + (1 - u)x_{0,i}$. Then, the challenge ciphertext is a valid ciphertext for the message $\hat{\mathbf{x}} = u\mathbf{x}_1 + (1 - u)\mathbf{x}_0$, which is a random linear combination of the vectors $\mathbf{x}_0$ and $\mathbf{x}_1$ whose coefficients sum up to one, as expected. Notice that b is information theoretically hidden because the distribution of u is independent from b. Hence, the advantage of $\mathcal{B}$ in this game is 0, when a random non-zero a has been encrypted.

Eventually, this shows that ε is bounded by the best advantage one can get against s-IND-CPA of $\mathcal{E}$. Hence, taking the maximal values, the best advantage one can get against s-IND-CPA of IP is bounded by the best advantage one can get against s-IND-CPA of $\mathcal{E}$.

5 Instantiations

5.1 Instantiation from DDH

The scheme of Section 3 can be obtained by plugging into our generalization the ElGamal encryption scheme [ElG85] which supports the properties that we require. Namely:

RR, LKH and LCH properties. The secret key space of this PKE is the group $(\mathbb{Z}_p, +, 0)$, and the public key space is the group $(\mathbb{G}, \times, 1)$. It is easy to see that $\mathsf{pk}_1^a \mathsf{pk}_2$ is the public key corresponding to the secret key $a\mathsf{sk}_1 + \mathsf{sk}_2$, and that

$$\mathsf{ct}_1 \cdot \mathsf{ct}_1' = \mathsf{pk}^r g^m \cdot \mathsf{pk}'^r g^{m'} = (\mathsf{pk} \cdot \mathsf{pk}')^r \cdot g^{m+m'}.$$

For RR, see that $\mathsf{ct}_0^{\mathsf{sk}} = \mathsf{pk}^r$.

5.2   Instantiation from LWE

The LWE Assumption. The learning with errors (LWE) problem was introduced by Regev [Reg05]. Let n, q be integer parameters. For any noise distribution χ on $\mathbb{Z}_q$, and vector $\mathbf{s} \in \mathbb{Z}_q^n$, the oracle $\mathsf{LWE}_{q,n,\chi}(\mathbf{s})$ samples a fresh random n-dimensional vector $\mathbf{a} \leftarrow \mathbb{Z}_q^n$, as well as noise $e \leftarrow \chi$, and returns $(\mathbf{a}, \langle \mathbf{a}, \mathbf{s} \rangle + e)$. The LWE assumption with noise χ states that for every PPT distinguisher $\mathcal{D}$,

$$\Pr[\mathbf{s} \leftarrow \mathbb{Z}_q^n \;:\; \mathcal{D}^{\mathsf{LWE}_{q,n,\chi}(\mathbf{s})} = 1] - \Pr[\mathbf{s} \leftarrow \mathbb{Z}_q^n \;:\; \mathcal{D}^{\mathsf{LWE}_{q,n,U}(\mathbf{s})} = 1] = \mathsf{negl}(n),$$

where U is the uniform distribution on $\mathbb{Z}_q$.

In other words, in addition to $\mathbf{a}$, the oracle $\mathsf{LWE}_{q,n,U}(\mathbf{s})$ simply returns uniform random samples, independent of $\mathbf{s}$. In general, the error distribution χ is chosen to be a discrete Gaussian on $\mathbb{Z}_q$.

Let $n \in \mathbb{Z}^+$ be a security parameter. Let $q = q(n)$, $m = m(n)$, and $1 < p < q$ be positive integers. Let $\alpha = \alpha(n)$ be positive real Gaussian parameters. Let $r = r(n) \geq 2$ be an integer and define $k = k(n) := \lfloor \log_r q \rfloor$.

Construction 3 (LWE-PKE Scheme). *We define our* public-key encryption scheme $\mathcal{E} = (\mathsf{Setup}, \mathsf{Encrypt}, \mathsf{Decrypt})$ *as follows.*

- $\mathsf{Setup}(1^n)$ *samples* $\mathbf{A} \leftarrow \mathbb{Z}_q^{m \times n}$. *Then, for* $\gamma = 0, \ldots, k$, *samples* $\mathbf{s}_\gamma \leftarrow \mathbb{Z}_q^n$ *and computes* $\mathbf{b}_\gamma = \mathbf{A}\mathbf{s}_\gamma + \mathbf{e}_\gamma \in \mathbb{Z}_q^m$, *where* $\mathbf{e}_\gamma \leftarrow \chi^m$. *Then, the algorithm sets* $\mathsf{pk} = (\mathbf{A}, (\mathbf{b}_\gamma)_\gamma)$ *and* $\mathsf{sk} = (\mathbf{s}_\gamma)_\gamma$, *and returns the pair* $(\mathsf{pk}, \mathsf{sk})$.
- $\mathsf{Encrypt}(\mathsf{pk}, x)$ *on input public key* pk, *and message* $x \in \mathbb{Z}_p$, *chooses random* $\mathbf{r} \leftarrow \{0, 1\}^m$ *and computes*

$$\mathsf{ct}_0 = \mathbf{A}^\top \mathbf{r} \in \mathbb{Z}_q^n$$

and, for $\gamma = 0, \ldots, k$,

$$\mathsf{ct}_\gamma = \langle \mathbf{r}, \mathbf{b}_\gamma \rangle + \left\lfloor \frac{q}{p} \right\rfloor r^\gamma x_\gamma \in \mathbb{Z}_q,$$

where $\sum_\gamma x_\gamma r^\gamma = x$. *Then the algorithm returns the ciphertext* $\mathsf{Ct} = (\mathsf{ct}_0, (\mathsf{ct}_\gamma)_\gamma)$.
- $\mathsf{Decrypt}(\mathsf{pk}, \mathsf{Ct}, \mathsf{sk})$ *on input public key* pk, *ciphertext* $\mathsf{Ct} = (\mathsf{ct}_0, (\mathsf{ct}_\gamma))$ *and secret key* sk, *returns the evaluation*

$$y = \left\lfloor \frac{p}{q} \cdot \left(\sum \mathsf{ct}_\gamma - \langle \mathsf{ct}_0, \mathsf{sk} \rangle \right) \right\rceil .$$

Semantic security. Notice that the above scheme is the Regev scheme where the message is r-decomposed to ensure that the error doesn't grow too much. Then, the proof of the semantic security of this encryption scheme can be found in [Reg05]. Essentially, the proof relies on two key properties:

- $\mathbf{A}^\top \mathbf{r} \in \mathbb{Z}_q^n$ is computationally indistinguishable from a random vector;
- Distinguishing a $\mathbf{As}_\gamma + \mathbf{e}_\gamma$ from a random vector is breaking the LWE assumption.

RR, LKH and LCH properties. The secret key space of this PKE is the group $((\mathbb{Z}_q^n)^{k+1}, +, \mathbf{0})$, and the public key space is the group $((\mathbb{Z}_q^m)^{k+1}, +, \mathbf{0})$.
It is easy to see that $a\mathsf{pk}_1 + \mathsf{pk}_2$ corresponds to the secret key $a\mathsf{sk}_1 + \mathsf{sk}_2$, and that

$$\mathsf{ct}_\gamma + \mathsf{ct}'_\gamma = \langle \mathbf{r}, \mathbf{b}_\gamma \rangle + \left\lfloor \frac{q}{p} \right\rfloor r^\gamma x_\gamma + \langle \mathbf{r}, \mathbf{b}'_\gamma \rangle + \left\lfloor \frac{q}{p} \right\rfloor r^\gamma x'_\gamma$$

$$= \langle \mathbf{r}, \mathbf{b}_\gamma + \mathbf{b}'_\gamma \rangle + \left\lfloor \frac{q}{p} \right\rfloor r^\gamma (x_\gamma + x'_\gamma).$$

For RR, see that $\langle \mathbf{A}^\top \mathbf{r}, \mathbf{s} \rangle = \langle \mathbf{r}, \mathbf{As} \rangle$.

Acknowledgments

References

[ABN10]   Michel Abdalla, Mihir Bellare, and Gregory Neven. Robust encryption. In Daniele Micciancio, editor, *TCC 2010*, volume 5978 of *LNCS*, pages 480–497. Springer, February 2010.

[BBS03]   Mihir Bellare, Alexandra Boldyreva, and Jessica Staddon. Randomness re-use in multi-recipient encryption schemeas. In Yvo Desmedt, editor, *PKC 2003*, volume 2567 of *LNCS*, pages 85–99. Springer, January 2003.

[BCP14]   Elette Boyle, Kai-Min Chung, and Rafael Pass. On extractability obfuscation. In Yehuda Lindell, editor, *TCC 2014*, volume 8349 of *LNCS*, pages 52–73. Springer, February 2014.

[BF01]   Dan Boneh and Matthew K. Franklin. Identity-based encryption from the Weil pairing. In Joe Kilian, editor, *CRYPTO 2001*, volume 2139 of *LNCS*, pages 213–229. Springer, August 2001.

[BO12]   Mihir Bellare and Adam O'Neill. Semantically-secure functional encryption: Possibility results, impossibility results and the quest for a general definition. Cryptology ePrint Archive, Report 2012/515, 2012. `http://eprint.iacr.org/2012/515`.

[BR06]   Mihir Bellare and Phillip Rogaway. The security of triple encryption and a framework for code-based game-playing proofs. In Serge Vaudenay, editor, *EUROCRYPT 2006*, volume 4004 of *LNCS*, pages 409–426. Springer, May / June 2006.

[BSW11]   Dan Boneh, Amit Sahai, and Brent Waters. Functional encryption: Definitions and challenges. In Yuval Ishai, editor, *TCC 2011*, volume 6597 of *LNCS*, pages 253–273. Springer, March 2011.

[BW07]   Dan Boneh and Brent Waters. Conjunctive, subset, and range queries on encrypted data. In Salil P. Vadhan, editor, *TCC 2007*, volume 4392 of *LNCS*, pages 535–554. Springer, February 2007.

[Coc01]   Clifford Cocks. An identity based encryption scheme based on quadratic residues. In Bahram Honary, editor, *8th IMA International Conference on Cryptography and Coding*, volume 2260 of *LNCS*, pages 360–363. Springer, December 2001.

[ElG84]   Taher ElGamal. A public key cryptosystem and a signature scheme based on discrete logarithms. In G. R. Blakley and David Chaum, editors, *CRYPTO'84*, volume 196 of *LNCS*, pages 10–18. Springer, August 1984.

[ElG85]   Taher ElGamal. A public key cryptosystem and a signature scheme based on discrete logarithms. *IEEE Transactions on Information Theory*, 31:469–472, 1985.

[Gen06]   Craig Gentry. Practical identity-based encryption without random oracles. In Serge Vaudenay, editor, *EUROCRYPT 2006*, volume 4004 of *LNCS*, pages 445–464. Springer, May / June 2006.

[GGH+13] Sanjam Garg, Craig Gentry, Shai Halevi, Mariana Raykova, Amit Sahai, and Brent Waters. Candidate indistinguishability obfuscation and functional encryption for all circuits. Cryptology ePrint Archive, Report 2013/451, 2013. http://eprint.iacr.org/2013/451.

[GGHZ14] Sanjam Garg, Craig Gentry, Shai Halevi, and Mark Zhandry. Fully secure attribute based encryption from multilinear maps. Cryptology ePrint Archive, Report 2014/622, 2014. http://eprint.iacr.org/2014/622.

[GLW12]   Shafi Goldwasser, Allison B. Lewko, and David A. Wilson. Bounded-collusion IBE from key homomorphism. In Ronald Cramer, editor, *TCC 2012*, volume 7194 of *LNCS*, pages 564–581. Springer, March 2012.

[GPSW06] Vipul Goyal, Omkant Pandey, Amit Sahai, and Brent Waters. Attribute-based encryption for fine-grained access control of encrypted data. In Ari Juels, Rebecca N. Wright, and Sabrina De Capitani di Vimercati, editors, *ACM CCS 06*, pages 89–98. ACM Press, October / November 2006. Available as Cryptology ePrint Archive Report 2006/309.

[KSW08]   Jonathan Katz, Amit Sahai, and Brent Waters. Predicate encryption supporting disjunctions, polynomial equations, and inner products. In Nigel P. Smart, editor, *EUROCRYPT 2008*, volume 4965 of *LNCS*, pages 146–162. Springer, April 2008.

[Kur02]   Kaoru Kurosawa. Multi-recipient public-key encryption with shortened ciphertext. In David Naccache and Pascal Paillier, editors, *PKC 2002*, volume 2274 of *LNCS*, pages 48–63. Springer, February 2002.

[LOS+10]   Allison B. Lewko, Tatsuaki Okamoto, Amit Sahai, Katsuyuki Takashima, and Brent Waters. Fully secure functional encryption: Attribute-based encryption and (hierarchical) inner product encryption. In Henri Gilbert, editor, *EUROCRYPT 2010*, volume 6110 of *LNCS*, pages 62–91. Springer, May 2010.

[O'N10]   Adam O'Neill. Definitional issues in functional encryption. Cryptology ePrint Archive, Report 2010/556, 2010. `http://eprint.iacr.org/2010/556`.

[OT12]   Tatsuaki Okamoto and Katsuyuki Takashima. Adaptively attribute-hiding (hierarchical) inner product encryption. In David Pointcheval and Thomas Johansson, editors, *EUROCRYPT 2012*, volume 7237 of *LNCS*, pages 591–608. Springer, April 2012.

[Reg05]   Oded Regev. On lattices, learning with errors, random linear codes, and cryptography. In Harold N. Gabow and Ronald Fagin, editors, *37th ACM STOC*, pages 84–93. ACM Press, May 2005.

[Rot11]   Ron Rothblum. Homomorphic encryption: From private-key to public-key. In Yuval Ishai, editor, *TCC 2011*, volume 6597 of *LNCS*, pages 219–234. Springer, March 2011.

[Sha84]   Adi Shamir. Identity-based cryptosystems and signature schemes. In G. R. Blakley and David Chaum, editors, *CRYPTO'84*, volume 196 of *LNCS*, pages 47–53. Springer, August 1984.

[SW05]   Amit Sahai and Brent R. Waters. Fuzzy identity-based encryption. In Ronald Cramer, editor, *EUROCRYPT 2005*, volume 3494 of *LNCS*, pages 457–473. Springer, May 2005.

[TW14]   Stefano Tessaro and David A. Wilson. Bounded-collusion identity-based encryption from semantically-secure public-key encryption: Generic constructions with short ciphertexts. In Hugo Krawczyk, editor, *PKC 2014*, volume 8383 of *LNCS*, pages 257–274. Springer, March 2014.

[Wat12]   Brent Waters. Functional encryption for regular languages. In Reihaneh Safavi-Naini and Ran Canetti, editors, *CRYPTO 2012*, volume 7417 of *LNCS*, pages 218–235. Springer, August 2012.

[Wat14]   Brent Waters. A punctured programming approach to adaptively secure functional encryption. Cryptology ePrint Archive, Report 2014/588, 2014. `http://eprint.iacr.org/2014/588`.