Non-zero Inner Product Encryption with Short Ciphertexts and Private Keys

Jie Chen[1,2(✉)], Benoît Libert[1(✉)], and Somindu C. Ramanna[1(✉)]

[1] Laboratoire LIP, École Normale Supérieure de Lyon, Lyon, France
{benoit.libert,somindu.ramanna}@ens-lyon.fr
[2] East China Normal University, Shanghai, China
s080001@e.ntu.edu.sg

Abstract. We describe two constructions of non-zero inner product encryption (NIPE) systems in the public index setting, both having ciphertexts and secret keys of constant size. Both schemes are obtained by tweaking the Boneh-Gentry-Waters broadcast encryption system (Crypto 2005) and are proved selectively secure under previously considered assumptions in groups with a bilinear map. Our first realization builds on prime-order bilinear groups and is proved secure under the Decisional Bilinear Diffie-Hellman Exponent assumption, which is parameterized by the length n of vectors over which the inner product is defined. By moving to composite order bilinear groups, we are able to obtain security under static subgroup decision assumptions following the Déjà Q framework of Chase and Meiklejohn (Eurocrypt 2014) and its extension by Wee (TCC 2016). Our schemes are the first NIPE systems to achieve such parameters, even in the selective security setting. Moreover, they are the first proposals to feature optimally short private keys, which only consist of *one* group element. Our prime-order-group realization is also the first one with a deterministic key generation mechanism.

Keywords: Functional encryption · Non-zero inner products · (Identity-based) revocation

1 Introduction

Attribute-based encryption (ABE) [20,35] allows fine-grained access control to encrypted data. In an ABE system, a ciphertext has an associated attribute $\boldsymbol{x}$ and a secret key for a user associated to some attribute $\boldsymbol{y}$ can successfully decrypt iff some relation R on $\boldsymbol{x}, \boldsymbol{y}$ holds true i.e., $R(\boldsymbol{x}, \boldsymbol{y}) = 1$. An ABE scheme is said to be secure if a collusion attack by a group of users does not compromise the security of a ciphertext they are not allowed to decrypt. In this work, we consider attributes belonging to some inner product space V and the relation is given by $R(\boldsymbol{x}, \boldsymbol{y}) = 1$ iff $\langle \boldsymbol{x}, \boldsymbol{y} \rangle \neq 0$, for $\boldsymbol{x}, \boldsymbol{y} \in V$. Such an ABE (referred to as non-zero inner product encryption scheme or NIPE) is known to imply identity-based revocation, an important cryptographic primitive in its own right.

V. Zikas and R. De Prisco (Eds.): SCN 2016, LNCS 9841, pp. 23–41, 2016.
DOI: 10.1007/978-3-319-44618-9_2

Identity-based revocation (IBR) allows a sender to encrypt and broadcast a message to a number of identities, given a set of revoked users $\mathcal{R}$, so that only secret keys associated with identities outside of $\mathcal{R}$ can decrypt the message. NIPE systems are known to imply IBR – the attribute associated with the ciphertext (of length n) is nothing but the vector of coefficients of the polynomial $p_{\mathcal{R}}(Z) = \prod_{\mathsf{id}_i \in \mathcal{R}}(Z - \mathsf{id}_i)$ where $|\mathcal{R}| \leq n$ and the secret key for an identity id corresponds to the vector $(1, \mathsf{id}, \ldots, \mathsf{id}^n)$. The inner product is non-zero if and only if $p_{\mathcal{R}}(\mathsf{id}) \neq 0$ or equivalently $\mathsf{id} \notin \mathcal{R}$, in which case decryption succeeds.

In this paper, our main goal is to design NIPE (and thus revocation) schemes that simultaneously provide short ciphertexts and private keys. We will also seek to prove security under well-studied hardness assumptions.

Our Contribution. We first present a NIPE system employing prime-order bilinear groups where ciphertexts *and* secret keys *both* have constant[1] size. Our scheme is the first one where both sizes can be constant. Indeed, all earlier realizations [4,5,34] providing $O(1)$-size ciphertexts (resp. $O(1)$-size private keys) indeed required $O(n)$ group elements in private keys (resp. in ciphertexts), where n denotes the dimension of the inner product space which is fixed at setup time. Even in the selective model [4,5], all previous constructions thus had linear complexities in the size of ciphertexts or private keys.

The scheme is also the first NIPE realization to feature optimally short private keys – which only consist of one group element – via a deterministic private key extraction algorithm. In particular, our NIPE scheme implies the first (identity-based) revocation system that simultaneously provides $O(1)$-size ciphertexts and private keys. It thus performs in the same way as the Boneh-Gentry-Waters (BGW) broadcast encryption [12] system and relies on the same assumption. Like earlier NIPE proposals, our scheme requires $O(n)$ group elements in the public parameters. In the revocation setting, this translates into a linear public key size in the maximal number of revoked users per ciphertext, which is on par with solutions [29,38] based on the Naor-Pinkas technique [29].

The security of our scheme is proved against selective adversaries under the n-Decisional Bilinear Diffie-Hellman (n-DBDHE) assumption, the strength of which depends on the dimension n of handled vectors. While relying on such a parameterized assumption is certainly a caveat [17], our scheme can be modified so as to dispense with variable-size assumptions.

Our second contribution is a NIPE system based on composite order pairing groups with security under constant-size subgroup decision assumptions. The proof follows the Déjà Q framework of [16,40]. Even in the restrictive selective model of security, our scheme is the first one to achieve constant size ciphertexts and keys under static assumptions.

[1] One may object saying the linear-length vector $\boldsymbol{x}$ still has to be appended to the ciphertext. Nevertheless, in many applications the description of $\boldsymbol{x}$ can be very short. For example, in an ordinary (i.e., non-identity-based) broadcast encryption scheme for n users, $\boldsymbol{x}$ is uniquely determined by the n-bit word that specifies which users are in the revoked set. In this case, our ciphertexts reduce the communication overhead from $O(n\lambda)$ to $O(n + \lambda)$ bits if λ is the security parameter.

In the context of revocation, not only do we provide the first identity-based revocation systems with constant-size ciphertexts and keys, but we also give a solution based on fairly well-studied subgroup assumptions in composite order groups. It remains a challenging open problem (at least without using a complexity leveraging argument [8] entailing an exponential security loss) to achieve similar efficiency tradeoffs while proving security against adaptive adversaries.

Outline of the Constructions and Proofs. We begin with the first construction based on an asymmetric prime-order pairing $e : \mathbb{G} \times \hat{\mathbb{G}} \to \mathbb{G}_T$ with group order p. The public key consists of $g^{\alpha^i}, \hat{g}^{\alpha^i}$ for $i \in [1, 2n] \setminus \{n+1\}$ along with g^γ where g and α, γ are sampled at random from $\mathbb{G}$ and $\mathbb{Z}_p$, respectively. In addition the element $e(g, \hat{g})^{\alpha^{n+1}}$ is provided. A ciphertext for an attribute vector $\boldsymbol{x} \in \mathbb{Z}_p^n$ and message m consists of $(m \cdot e(g, \hat{g})^{\alpha^{n+1} s}, g^s, (v \cdot g^{\sum_{i=1}^n \alpha^i x_i})^s)$. Secret key associated with a vector $\boldsymbol{y}$ is computed deterministically as $\hat{g}^{\gamma \sum_{i=1}^n \alpha^{n-i+1} y_i}$. The structure is reminiscent of the Boneh-Gentry-Waters broadcast encryption scheme [12]. The proof of security is a reduction from the hardness of the n-DBDHE problem – an instance consists of $g^{\alpha^i}, \hat{g}^{\alpha^i}$ for $i \in [1, 2n] \setminus \{n+1\}, g^s \in \mathbb{G}, T \in \mathbb{G}_T$ and asks to decide whether $T = e(g, \hat{g})^{\alpha^{n+1} s}$ or $T \xleftarrow{R} \mathbb{G}_T$. The attacker declares a target vector $\boldsymbol{x}^*$ which is used to program $\gamma = \sum_{i=1}^n \alpha^i x_i^*$. For any $\boldsymbol{y} \in \mathbb{Z}_p^n$ with $\langle \boldsymbol{x}^*, \boldsymbol{y} \rangle = 0$, secret key $d_{\boldsymbol{y}}$ can be simulated using the elements provided in the instance because for $d_{\boldsymbol{y}}$, the coefficient of α^{n+1} in the exponent of $\hat{g}$ would be $\langle \boldsymbol{x}^*, \boldsymbol{y} \rangle = 0$. The attacker then provides two messages m_0, m_1 to which the challenger responds with the ciphertext $(m_\beta \cdot T, g^s, (v \cdot g^{\sum_{i=1}^n \alpha^i x_i})^s)$ for a randomly chosen bit β. An adversary's ability to determine whether the message encrypted in the challenge ciphertext is real or random can be leveraged to solve the given instance of the decision problem.

We then consider a variant in the setting of a composite-order symmetric pairing $e : \mathbb{G} \times \mathbb{G} \to \mathbb{G}_T$ of common group order $N = p_1 p_2 p_3$, similar to Wee's composite-order variant [40] of the broadcast encryption in [12]. (Let $\mathbb{G}_q$ denote the subgroup of $\mathbb{G}$ of order q where q would be of the form $p_1^{e_1} p_2^{e_2} p_3^{e_3}$ for $e_1, e_2, e_3 \in \{0, 1\}$). The public key is composed of $v = g^\gamma, (g^{\alpha^i})_{i=1}^n$, $U_j = u^{\alpha^j}$, $j \in [1, 2n] \setminus \{n+1\}$ for some $g, u \xleftarrow{R} \mathbb{G}$ and $\alpha, \gamma \in \mathbb{Z}_N$ along with a pairwise-independent hash function $\mathsf{H} : \mathbb{G}_T \to \{0,1\}^\lambda$. Decryption key for a vector $\boldsymbol{y}$ is defined as $u^{\gamma \sum_{i=1}^n \alpha^{n-i+1} y_i}$ and the ciphertext for attribute $\boldsymbol{x}$ and message M is defined as $(M \oplus \mathsf{H}(e(g, u)^{\alpha^{n+1} s}), g^s, (v \cdot g^{\sum_{i=1}^n \alpha^i x_i})^s)$. In addition, the parameters U_j and secret keys are randomized with $\mathbb{G}_{p_3}$-components. The security is reduced to two standard subgroup decision assumptions, denoted $(p_1 \to p_1 p_2)$ and $(p_1 p_3 \to p_1 p_2 p_3)$, where $(q_1 \to q_2)$ subgroup decision problem asks to distinguish between random elements of $\mathbb{G}_{q_1}$ from random elements of $\mathbb{G}_{q_2}$. The reduction gradually adds $\mathbb{G}_{p_2}$-components to the challenge ciphertext as well as elements $(U_j)_{j=1}^{2n}$ so that at the end, each U_j has in its exponent a pseudorandom function $RF : [1, 2n] \to \mathbb{Z}_{p_2}$ evaluated at j. The element $v = g^\gamma$ is programmed based on the challenge attribute $\boldsymbol{x}^*$ in a manner similar to the reduction in the prime-order case. Additionally, this ensures that the challenge ciphertext components are independent of $\alpha \bmod p_2$. Given this and the fact that keys are

generated only for vectors $\boldsymbol{y}$ with $\langle \boldsymbol{x}^*, \boldsymbol{y} \rangle = 0$, α^{n+1} does not appear in the exponent of u in any of the keys. On the other hand, the message is masked by the hash of an element of $\mathbb{G}_T$ determined by $RF(n+1)$. Since all information provided to the attacker is independent of $RF(n+1)$, we use the left over hash lemma to argue that the mask on the message is uniformly distributed and hence statistically hides the message from the attacker.

Related Work. The inner product functionality was first considered by Katz et al. [22] in the design of predicate encryption systems (i.e., ABE schemes in the private index setting). Their construction [22] initiated a large body of work [2,24,30–34,36] which considered hierarchical extensions [30,33], additional properties in the secret-key setting [36] and adaptively secure realizations [24,31–34].

In the public-index setting, inner products also proved useful [4] to build adaptively secure identity-based broadcast encryption (IBBE) and revocation schemes with short ciphertexts under simple assumptions. The first construction of non-zero IPE appeared in [4] with security in the *co-selective* model under the Decision Linear [9] and Decisional Bilinear Diffie-Hellman assumptions. Co-selective security requires an adversary to commit to the attributes corresponding to private key queries before seeing the public parameters of the scheme, as opposed to target attribute set in the selective model. It is slightly stronger than the selective model but weaker than the adaptive model. The scheme has constant-size ciphertexts whereas its public parameters and keys are of size linear in n. More efficient realizations (but with asymptotically similar parameters) were put forth by Attrapadung *et al.* [5] and Yamada *et al.* [41] under the n-DBDHE assumption. While some of the NIPE constructions of [5,41] have exactly the same ciphertext length (resp. private key length) as our scheme, they require $O(n)$-size private keys (resp. $O(n)$-size ciphertexts). We thus prove security under the same assumption as [5,41] with only one group element per private key and 3 group elements per ciphertext.

The first adaptively secure NIPE scheme was proposed in [34] with $O(n)$ group elements in the public parameters and either $O(1)$-size ciphertexts or $O(1)$-size keys with a security reduction to the Decision Linear assumption. A more efficient construction was provided in [15] via an instantiation of predicate encodings [39] in prime-order groups. On the other hand, either ciphertexts or secret keys had size linear in n. Previously known constructions did not consider simultaneously achieving constant size ciphertexts and secret keys.

More recently, Abdalla *et al.* [1] suggested a different inner product functionality which evaluates linear functions of encrypted data (i.e., their inner product with a vector associated with the private key), instead of only testing if they evaluate to 0 as in [22,24,31–34]. Under simple assumptions, they obtained practical solutions based on the standard Decision Diffie-Hellman and Learning-With-Errors assumptions. Their results were extended to handle adaptive adversaries [3] and function-privacy in the secret-key setting [6].

In the context of IBBE scheme, Delerablée [18] suggested a selectively secure construction with constant-size ciphertexts and private keys based on strong q-type assumptions. Her construction actually remains the most efficient IBBE

in the literature to date. The IBR system implied by our first NIPE construction can be seen as the revocation analogue of Delerablée's IBBE as it simultaneously provides $O(1)$-size ciphertexts and keys (the public parameters also have linear length in the maximal number of receivers per ciphertext in [18]). Unlike our IBR system, however, [18] is not known to have a counterpart based on simple assumptions in composite order groups. In the identity-based revocation setting, the constructions of Lewko et al. [23] feature constant-size private keys and public parameters, but their ciphertext size is linear in the number of revoked users. While their first construction has very short private keys and public parameters (made of 3 and 4 group elements, respectively), its underlying complexity assumption is very *ad hoc* and even stronger than n-DBDHE.

The Déjà Q framework, introduced by Chase and Meiklejohn [16], allows reducing well-studied fixed-size assumptions, such as the Subgroup Decision assumption [11] to some families of parameterized assumptions in composite-order groups. As a result, some well-known constructions such as Dodis-Yampolskiy PRF [19] and Boneh-Boyen signatures [7], when instantiated in composite order groups, could be shown secure under subgroup decision assumptions. Wee [40] further advanced the framework to cover certain encryption primitives as well, in addition to removing the restriction to work with asymmetric composite order groups. The primitives include adaptively secure identity-based encryption and selectively secure broadcast encryption. Recently, Libert *et al.* [26] applied Wee's framework to obtain functional commitment schemes for linear functions and accumulators from simple assumptions.

2 Background

2.1 Bilinear Maps and Complexity Assumptions

Assumptions in prime order groups. Let $(\mathbb{G}, \hat{\mathbb{G}}, \mathbb{G}_T)$ be groups of prime order p with a bilinear map $e : \mathbb{G} \times \hat{\mathbb{G}} \to \mathbb{G}_T$. We rely on a parameterized assumption introduced by Boneh et al. [12]. While this assumption was defined using symmetric pairings [10,12], we consider a natural extension to asymmetric pairings, which will enable our most efficient construction.

Definition 1. *Let $(\mathbb{G}, \hat{\mathbb{G}}, \mathbb{G}_T)$ be bilinear groups of prime order p. The* **n-Decision Bilinear Diffie-Hellman Exponent** *(n-DBDHE) problem is, given a tuple $(g, g^{\alpha}, g^{(\alpha^2)}, \ldots, g^{(\alpha^n)}, g^{(\alpha^{n+2})}, \ldots, g^{(\alpha^{2n})}, h, \hat{g}, \hat{g}^{\alpha}, \hat{g}^{(\alpha^2)}, \ldots, \hat{g}^{(\alpha^n)}, \hat{g}^{(\alpha^{n+2})}, T)$ where $g, h \xleftarrow{R} \mathbb{G}$, $\hat{g} \xleftarrow{R} \hat{\mathbb{G}}$, $\alpha \xleftarrow{R} \mathbb{Z}_p$ and $T \in_R \mathbb{G}_T$, to decide if $T = e(h, \hat{g})^{(\alpha^{n+1})}$ or if T is a random element of $\mathbb{G}_T$.*

Assumptions in composite order groups.We use groups $(\mathbb{G}, \mathbb{G}_T)$ of composite order $N = p_1 p_2 p_3$ endowed with an efficiently computable map (a.k.a. pairing) $e : \mathbb{G} \times \mathbb{G} \to \mathbb{G}_T$ such that: (1) $e(g^a, h^b) = e(g, h)^{ab}$ for any $(g, h) \in \mathbb{G} \times \mathbb{G}$ and $a, b \in \mathbb{Z}$; (2) if $e(g, h) = 1_{\mathbb{G}_T}$ for each $h \in \mathbb{G}$, then $g = 1_{\mathbb{G}}$. An important property of composite order groups is that pairing two elements of order p_i and p_j, with $i \neq j$, always gives the identity element $1_{\mathbb{G}_T}$.

In the following, for each $i \in \{1,2,3\}$, we denote by $\mathbb{G}_{p_i}$ the subgroup of order p_i. For all distinct $i,j \in \{1,2,3\}$, we call $\mathbb{G}_{p_ip_j}$ the subgroup of order p_ip_j. In this setting, we rely on the following assumptions introduced in [25].

Assumption 1. Given a description of $(\mathbb{G}, \mathbb{G}_T, e)$ as well as $g \xleftarrow{R} \mathbb{G}_{p_1}, g_3 \xleftarrow{R} \mathbb{G}_{p_3}$ and $T \in \mathbb{G}$, it is infeasible to efficiently decide if $T \in \mathbb{G}_{p_1p_2}$ or $T \in \mathbb{G}_{p_1}$.

Assumption 2. Let $g, X_1 \xleftarrow{R} \mathbb{G}_{p_1}, X_2, Y_2 \xleftarrow{R} \mathbb{G}_{p_2}, g_3, Y_3 \xleftarrow{R} \mathbb{G}_{p_3}$. Given a description of $(\mathbb{G}, \mathbb{G}_T, e)$, a set of group elements (g, X_1X_2, g_3, Y_2Y_3) and T, it is hard to decide if $T \in_R \mathbb{G}_{p_1p_3}$ or $T \in_R \mathbb{G}$.

These assumptions are non-interactive and falsifiable [28]. Moreover, in both of them, the number of input elements is constant (*i.e.*, independent of the number of adversarial queries).

2.2 Non-zero Inner Product Encryption (IPE)

Definition 2 (NIPE). *Let V denote an inner product space of dimension n and $\mathcal{M}$ denote the message space. A non-zero inner product encryption (NIPE) scheme for inner products over V, is defined by four probabilistic algorithms – Setup, Encrypt, KeyGen and Decrypt.*

Setup(λ, n)**:** *Takes as input a security parameter λ and the dimension of V. It outputs the public parameters mpk and the master secret msk.*

KeyGen$(msk, \boldsymbol{y})$**:** *On input a vector $\boldsymbol{y} \in V$ and the master secret msk; this algorithm outputs a secret key $d_{\boldsymbol{y}}$ for $\boldsymbol{y}$.*

Encrypt$(mpk, m, \boldsymbol{x})$**:** *Takes as input a message m and an attribute vector $\boldsymbol{x} \in V$ and outputs a ciphertext $\mathcal{C}$.*

Decrypt$(mpk, \mathcal{C}, d_{\boldsymbol{y}})$**:** *If $\langle \boldsymbol{x}, \boldsymbol{y} \rangle \neq 0$, this algorithm returns the message m and $\perp$ otherwise.*

Correctness. A NIPE scheme satisfies the correctness condition if for all vectors $\boldsymbol{x}, \boldsymbol{y} \in V$ with $\langle \boldsymbol{x}, \boldsymbol{y} \rangle \neq 0$ and for any message $m \in \mathcal{M}$, any keys $(\mathsf{mpk}, \mathsf{msk}) \leftarrow \mathsf{Setup}(\lambda, n)$, $d_{\boldsymbol{y}} \leftarrow \mathsf{KeyGen}(\mathsf{msk}, \boldsymbol{y})$ and any ciphertext $\mathcal{C} \leftarrow \mathsf{Encrypt}(\mathsf{mpk}, m, \boldsymbol{x})$, then $\Pr[m = \mathsf{Decrypt}(\mathsf{mpk}, \mathcal{C}, d_{\boldsymbol{y}})] = 1$.

Definition 3 (Selective Security). *Selective security of a non-zero inner product encryption scheme is formalized in terms of the following game between an adversary $\mathcal{A}$ and a challenger.*

Initialization: *The adversary $\mathcal{A}$ declares a challenge vector $\boldsymbol{x}^\star$.*

Setup: *The challenger runs the Setup algorithm of the NIPE and gives the public parameters to the adversary $\mathcal{A}$.*

Key Extraction Phase 1: *The adversary makes a number of key extraction queries adaptively. For a query on a vector $\boldsymbol{y}$ with the restriction that $\langle \boldsymbol{x}^\star, \boldsymbol{y} \rangle = 0$, the challenger responds with a key $d_{\boldsymbol{y}}$.*

Challenge: *The adversary $\mathcal{A}$ provides two equal-length messages M_0, M_1. The challenger chooses a bit β uniformly at random from $\{0,1\}$, encrypts M_β to $\boldsymbol{x}^\star$ and returns the resulting ciphertext $\mathcal{C}^\star$ to $\mathcal{A}$.*

Key Extraction Phase 2: *$\mathcal{A}$ makes more key extraction queries under the same restriction that it can only query keys for vectors $\boldsymbol{y}$ with $\langle \boldsymbol{x}^\star, \boldsymbol{y} \rangle = 0$.*

Guess: *$\mathcal{A}$ outputs a bit β'.*

If $\beta = \beta'$, then $\mathcal{A}$ wins the game. The advantage of $\mathcal{A}$ in winning the above game is defined as

$$Adv_{\mathrm{NIPE},\mathcal{A}}(\lambda) = \left| \Pr[\beta = \beta'] - \frac{1}{2} \right|.$$

The NIPE scheme is said to be secure if every PPT adversary has negligible advantage in winning the above game.

3 A Construction for Non-zero Inner Products with Constant-Size Ciphertexts and Private Keys

Our scheme builds on the Boneh-Gentry-Waters broadcast encryption [12] and inherits its efficiency. In particular, the public parameters are exactly those of the BGW construction. In order to adapt it in the context of non-zero inner product encryption, we extend earlier observations which leveraged the BGW technique in the design of accumulators [13] and vector commitments [21,27].

It was shown in [21] that a public key of the form

$$\{(g_i = g^{(\alpha^i)}, \hat{g}_i = \hat{g}^{(\alpha^i)})\}_{i \in [1,2n] \setminus \{n+1\}}$$

allows committing to a vector $\boldsymbol{x} = (x_1, \ldots, x_n)$ in such a way that the commitment string $C = g^\gamma \cdot \prod_{j=1}^n g_j^{x_j}$ makes it possible to convincingly reveal the partial information $z = \langle \boldsymbol{x}, \boldsymbol{y} \rangle$ about the committed message $\boldsymbol{x}$. Namely, a single group element

$$W_z = \prod_{i=1, i \neq j}^{n} (\hat{g}_{n+1-i}^{\gamma} \prod_{j=1}^{n} \hat{g}_{n+1+j-i}^{x_j})^{y_i} \in \hat{\mathbb{G}} \tag{1}$$

can serve as a witness that $z = \langle \boldsymbol{x}, \boldsymbol{y} \rangle$, for public $\boldsymbol{x} \in \mathbb{Z}_p^n$ and $z \in \mathbb{Z}_p$, and the verifier accepts (z, W_z) if and only if the following relation holds:

$$e(C, \prod_{j=1}^{n} \hat{g}_{n+1-j}^{y_j}) = e(g_1, \hat{g}_n)^z \cdot e(g, W_z) \tag{2}$$

The binding property of the commitment scheme relies on the fact that neither $g_{n+1} = g^{(\alpha^{n+1})}$ nor $\hat{g}_{n+1} = \hat{g}^{(\alpha^{n+1})}$ is publicly available.

Our non-zero IPE scheme proceeds by randomizing both members of (2) – by raising them to a random power $s \in \mathbb{Z}_p$ – so that the randomized C can be

embedded in the ciphertext (together with g^s) while W_z serves as a decryption token. The decryption operation then computes $e(g_1, \hat{g}_n)^{s \cdot \langle \boldsymbol{x}, \boldsymbol{y} \rangle}$, which uncovers $e(g_1, \hat{g}_n)^s$ whenever $\langle \boldsymbol{x}, \boldsymbol{y} \rangle \neq 0$.

Our ciphertexts are of the form $\big(M \cdot e(g_1, \hat{g}_n)^s, g^s, (g^\gamma \cdot \prod_{j=1}^n g_j^{x_j})^s\big)$ and the challenge is thus to associate each vector $\boldsymbol{y} \in \mathbb{Z}_p$ with a short private key $d_{\boldsymbol{y}}$ so as to enable decryption. To achieve this, we observe that (1) can be re-written

$$W_z = (\prod_{i=1} \hat{g}_{n+1-i}^{y_i})^\gamma \cdot \prod_{i=1, i \neq j}^{n} \prod_{j=1}^{n} \hat{g}_{n+1+j-i}^{x_j y_i} \in \hat{\mathbb{G}},$$

where the second term is publicly computable as it does not depend on $\hat{g}_{n+1} = \hat{g}^{(\alpha^{n+1})}$. This implies that, if $\gamma \in \mathbb{Z}_p$ is the master secret key, the private key for a vector $\boldsymbol{y}$ can only consist of a single group element $d_{\boldsymbol{y}} = (\prod_{j=1}^n \hat{g}_{n+1-j}^{y_j})^\gamma \in \hat{\mathbb{G}}$.

Somewhat surprisingly, private keys are generated in a deterministic manner and, at first glance, their shape seems at odds with the collusion-resistance requirement: if $d_{\boldsymbol{y}_1}$ is a private key for $\boldsymbol{y}_1 \in \mathbb{Z}_p$ and $d_{\boldsymbol{y}_2}$ is a private key for $\boldsymbol{y}_2 \in \mathbb{Z}_p$, the product $d_{\boldsymbol{y}_1} \cdot d_{\boldsymbol{y}_2}$ is a valid private key for $\boldsymbol{y}_1 + \boldsymbol{y}_2$. However, this does not affect the functionality since any ciphertext that neither $d_{\boldsymbol{y}_1}$ nor $d_{\boldsymbol{y}_2}$ can decrypt must be labeled with a vector $\boldsymbol{x}$ such that $\langle \boldsymbol{x}, \boldsymbol{y}_1 \rangle = \langle \boldsymbol{x}, \boldsymbol{y}_2 \rangle = 0$, which implies $\langle \boldsymbol{x}, \boldsymbol{y}_1 + \boldsymbol{y}_2 \rangle = 0$. Said otherwise, combining several keys that cannot decrypt a given ciphertext only yields another key that remains unable to decrypt.

Setup(λ, n)**:** Choose bilinear groups $(\mathbb{G}, \hat{\mathbb{G}}, \mathbb{G}_T)$ of prime order $p > 2^\lambda$ and define the bilinear map e. Choose $g \xleftarrow{R} \mathbb{G}$, $\hat{g} \xleftarrow{R} \hat{\mathbb{G}}$, $\alpha, \gamma \xleftarrow{R} \mathbb{Z}_p$ at random in order to define $v = g^\gamma \in \mathbb{G}$ and

$$\begin{array}{ccc} g_1 = g^\alpha, & \cdots & g_n = g^{(\alpha^n)} \\ g_{n+2} = g^{(\alpha^{n+2})}, & \cdots & g_{2n} = g^{(\alpha^{2n})} \end{array}$$

and

$$\begin{array}{ccc} \hat{g}_1 = \hat{g}^\alpha, & \cdots & \hat{g}_n = \hat{g}^{(\alpha^n)} \\ \hat{g}_{n+2} = \hat{g}^{(\alpha^{n+2})}, & \cdots & \hat{g}_{2n} = \hat{g}^{(\alpha^{2n})} \end{array}$$

Define the master public key to consist of

$$\mathsf{mpk} := \Big((\mathbb{G}, \hat{\mathbb{G}}, \mathbb{G}_T, e),\ g,\ \hat{g},\ v,\ \{(g_j, \hat{g}_j)\}_{j \in [1,2n] \setminus \{n+1\}}\Big).$$

The master secret key is $\mathsf{msk} := \gamma$.

KeyGen$(\mathsf{msk}, \boldsymbol{y})$**:** To generate a key for the vector $\boldsymbol{y} = (y_1, \ldots, y_n) \in \mathbb{Z}_p^n$, compute and output $d_{\boldsymbol{y}} = \big(\prod_{i=1}^n \hat{g}_{n+1-i}^{y_i}\big)^\gamma \in \hat{\mathbb{G}}$.

Encrypt$(\mathsf{mpk}, \boldsymbol{x}, M)$**:** To encrypt $M \in \mathbb{G}_T$ under $\boldsymbol{x} = (x_1, \ldots, x_n) \in \mathbb{Z}_p^n$, choose $s \xleftarrow{R} \mathbb{Z}_p$ in order to compute and output

$$\mathcal{C} = (C_0, C_1, C_2) = \big(M \cdot e(g_1, \hat{g}_n)^s,\ g^s,\ (v \cdot \prod_{j=1}^n g_j^{x_j})^s\big).$$

Decrypt$(\mathsf{mpk}, \mathcal{C}, \boldsymbol{x}, d_{\boldsymbol{y}}, \boldsymbol{y})$**:** Given a ciphertext $\mathcal{C}$ labeled with $\boldsymbol{x} = (x_1, \ldots, x_n) \in \mathbb{Z}_p^n$ and a private key $d_{\boldsymbol{y}}$ associated with the vector $\boldsymbol{y} = (y_1, \ldots, y_n) \in \mathbb{Z}_p^n$, return $\perp$ if $\langle \boldsymbol{x}, \boldsymbol{y} \rangle = 0$. Otherwise, conduct the following steps.

1. Compute

$$\hat{A}_i = \prod_{j=1, j\neq i}^{n} \hat{g}_{n+1+j-i}^{x_j} \qquad \forall i \in \{1, \ldots, n\}. \tag{3}$$

2. Compute and output

$$M = C_0 \cdot \Big(\frac{e(C_1, d_{\boldsymbol{y}} \cdot \prod_{i=1}^{n} \hat{A}_i^{y_i})}{e(C_2, \prod_{i=1}^{n} \hat{g}_{n+1-i}^{y_i})} \Big)^{1/\langle \boldsymbol{x}, \boldsymbol{y} \rangle}. \tag{4}$$

The correctness of the scheme is easily verified by observing that

$$\begin{aligned} &\frac{e\big(g, (\prod_{i=1}^{n} \hat{g}_{n+1-i}^{y_i})^{\gamma} \cdot \prod_{i=1}^{n} \prod_{j=1, j\neq i}^{n} \hat{g}_{n+1-i+j}^{x_j y_i}\big)}{e\big(g^{\gamma} \cdot \prod_{j=1}^{n} g_j^{x_j}, \prod_{i=1}^{n} \hat{g}_{n+1-i}^{y_i}\big)} \\ &= \frac{e\big(g, (\prod_{i=1}^{n} \hat{g}_{n+1-i}^{y_i})^{\gamma} \cdot \prod_{i=1}^{n} \prod_{j=1, j\neq i}^{n} \hat{g}_{n+1-i+j}^{x_j y_i}\big)}{e(g^{\gamma} \cdot \prod_{j=1}^{n} \hat{g}_{n+1-i}^{y_i}) \cdot e\big(g, \prod_{i=1}^{n} \prod_{j=1}^{n} g_{n+1-i+j}^{x_j y_i}\big)} = e(g, \hat{g}_{n+1})^{-\sum_{i=1}^{n} x_i y_i}. \end{aligned} \tag{5}$$

By raising both members of (5) to the power $s \in \mathbb{Z}_p$ and using (3), we obtain the equality

$$e(C_1, d_{\boldsymbol{y}} \cdot \prod_{i=1}^{n} \hat{A}_i^{y_i}) / e(C_2, \prod_{i=1}^{n} \hat{g}_{n+1-i}^{y_i}) = e(g_1, \hat{g}_n)^{-s \cdot \langle \boldsymbol{x}, \boldsymbol{y} \rangle},$$

which explains why M can be computed as per (4) whenever $\langle \boldsymbol{x}, \boldsymbol{y} \rangle \neq 0$.

From an efficiency point of view, the receiver has to compute a product of only two pairings (which is faster than two individual pairing evaluations) while the encryption and decryption algorithms both require at most $O(n)$ exponentiations. Indeed, the value $d_{\boldsymbol{y}} \cdot \prod_{i=1}^{n} \hat{A}_i^{y_i}$ is computable via a multi-exponentiation involving $2n - 1$ base elements (rather than n^2 in a naive computation).

Theorem 1. *The scheme is selectively secure under the n-DBDHE assumption.*

Proof. Towards a contradiction, let $\mathcal{A}$ be a PPT adversary with non-negligible advantage ε in the selective security game. We build a reduction algorithm that takes as input $((\mathbb{G}, \hat{\mathbb{G}}, \mathbb{G}_T, e), g, h, \{(g_i, \hat{g}_i) = (g^{(\alpha^i)}, \hat{g}^{(\alpha^i)})\}_{i \in [1,2n] \setminus \{n+1\}}, T)$ and uses $\mathcal{A}$ to decide if $T = e(h, \hat{g})^{(\alpha^{n+1})}$ or $T \in_R \mathbb{G}_T$.

The adversary $\mathcal{A}$ first chooses a target vector $\boldsymbol{x}^{\star} = (x_1^{\star}, \ldots, x_n^{\star}) \in \mathbb{Z}_p^n$. To construct the master public key mpk, $\mathcal{B}$ chooses $\tilde{\gamma} \xleftarrow{R} \mathbb{Z}_p$ and computes

$$v = g^{\tilde{\gamma}} \cdot \prod_{j=1}^{n} g_j^{-x_j^{\star}} \in \mathbb{G},$$

which implicitly defines the master secret key msk to be $\gamma = \tilde{\gamma} - \sum_{j=1}^{n} x_j \cdot \alpha^j$. The adversary $\mathcal{A}$ is run on input of

$$\mathsf{mpk} := \Big(g,\ \hat{g},\ v,\ \{(g_i, \hat{g}_i) = (g^{(\alpha^i)}, \hat{g}^{(\alpha^i)})\}_{i \in [1,2n] \setminus \{n+1\}}\Big).$$

Observe that mpk is distributed as in the real scheme as v is uniformly distributed over $\mathbb{G}$. At any time, $\mathcal{A}$ can request a private key $d_{\boldsymbol{y}}$ for any vector $\boldsymbol{y} \in \mathbb{Z}_p^N$ such that $\langle \boldsymbol{x}, \boldsymbol{y} \rangle = 0$. To generate the private key $d_{\boldsymbol{y}} = \big(\prod_{i=1}^{n} \hat{g}_{n+1-i}^{y_i}\big)^{\gamma} \in \hat{\mathbb{G}}$, algorithm $\mathcal{B}$ can exploit the fact that, in the product,

$$\Big(\sum_{i=1}^{n} y_i \cdot \alpha^{n+1-i}\Big) \cdot \Big(\sum_{j=1}^{n} x_j^{\star} \cdot \alpha^j\Big) = \sum_{i=1}^{n} \sum_{j=1}^{n} x_j^{\star} \cdot y_i \cdot \alpha^{n+1-i+j},$$

the coefficient of α^{n+1} is exactly $\langle \boldsymbol{x}^{\star}, \boldsymbol{y} \rangle$, which must be zero in any legal private key query $\boldsymbol{y} \in \mathbb{Z}_p^n$. Specifically, $\mathcal{B}$ can compute

$$d_{\boldsymbol{y}} = \Big(\prod_{i=1}^{n} \hat{g}_{n+1-i}^{y_i}\Big)^{\tilde{\gamma}} \Big/ \prod_{i=1}^{n} \prod_{j=1, j \neq i}^{n} \hat{g}_{n+1-i+j}^{x_j^{\star} \cdot y_i}. \tag{6}$$

For any vector $\boldsymbol{y} \in \mathbb{Z}_p^n$ such that $\langle \boldsymbol{x}^{\star}, \boldsymbol{y} \rangle = 0$, $\mathcal{B}$ can thus compute the private key $d_{\boldsymbol{y}}$ as per (6).

In the challenge phase, $\mathcal{A}$ chooses messages $M_0, M_1 \in \mathbb{G}_T$ and expects to receive an encryption of one of these. At this point, $\mathcal{B}$ flips a fair coin $\beta \xleftarrow{R} \{0,1\}$ and computes

$$\mathcal{C} = (C_0, C_1, C_2) = \big(M_{\beta} \cdot T,\ h,\ h^{\tilde{\gamma}}\big),$$

which is returned as a challenge to $\mathcal{B}$. It is easy to see that, if $T = e(h, \hat{g})^{(\alpha^{n+1})}$, then $\mathcal{C}$ is a valid encryption of M_{β} for the vector $\boldsymbol{x}^{\star} = (x_1^{\star}, \ldots, x_n^{\star})$ and the encryption exponent $s = \log_g(h)$. In contrast, if $T \in_R \mathbb{G}_T$, the ciphertext carries no information about $\beta \in \{0,1\}$.

When $\mathcal{A}$ halts, it outputs a bit $\beta' \in \{0,1\}$. If $\beta' = \beta$, the reduction $\mathcal{B}$ outputs 1 (meaning that $T = e(h, \hat{g})^{(\alpha^{n+1})}$). Otherwise, it outputs 0. □

4 NIPE from Constant-Size Subgroup Assumptions

In this section, we present a non-zero inner-product encryption (NIPE) scheme based on composite order pairings $e : \mathbb{G} \times \mathbb{G} \to \mathbb{G}_T$ of common group order $N = p_1 p_2 p_3$, with security under the subgroup decision assumptions. For inner products over length-n vectors in $\mathbb{Z}_N$, the public parameter size is linear in n while ciphertexts and keys have constant size (independent of n). The resulting scheme is the first to achieve such parameters with selective security under constant size assumptions.

Similar to the prime-order case, it seems possible to derive this construction from a functional commitment scheme for linear functions [26] by randomizing

commitments and the verification equation. However, the transformation is not generic. A commitment C to $\boldsymbol{x} \in \mathbb{Z}_N^n$ in [26] is computed as $C = g^{\gamma} \cdot g^{\sum_{i=1}^{n} \alpha^i \cdot x_i}$. Elements $\left(g^{\gamma}, \{g^{\alpha^i}\}_{i=1}^{n}\right)$ are made available in the public parameters along with elements $U_j = u^{\alpha^j} \cdot R_{3,j}$ for $j \in [1, 2n] \setminus \{n+1\}$ with $R_{3,j}$ being randomly distributed in $\mathbb{G}_{p_3}$. The U_j's allow creating a short witness W_z for the statement $z = \langle \boldsymbol{x}, \boldsymbol{y} \rangle$ (for some $\boldsymbol{y} \in \mathbb{Z}_N^n$) using the secret random exponent γ.

$$W_z = \prod_{i=1}^{n} W_i^{y_i}, \quad \textit{where} \quad W_i = U_{n-i+1}^{\gamma} \prod_{j=1, j \neq i}^{n} U_{n+1+j-i}.$$

Consolidating all the terms that depend on γ into $W_{z,1}$, write $W_z = W_{z,1} \cdot W_{z,2}$. More precisely, we have

$$W_{z,1} = \prod_{i=1}^{n} U_{n-i+1}^{\gamma} \qquad \text{and} \qquad W_{z,2} = \prod_{i=1}^{n} \left(\prod_{j=1, j \neq i}^{n} U_{n+1+j-i} \right)^{y_i}.$$

Observe that the computation of $W_{z,2}$ is solely based on information available in the public parameters and $W_{z,1}$ is independent of $\boldsymbol{x}$. One can verify the validity of the witness W_z by simply checking whether the following equation holds.

$$e(C, \prod_{i=1}^{n} U_i^{y_i}) = e(g^{\alpha}, U_n)^z \cdot e(g, W_z).$$

Randomizing both sides of the above equation with $s \in \mathbb{Z}_N$ in the exponent leads us to the non-zero IPE. Namely, a ciphertext for a vector $\boldsymbol{x}$ and a message $M \in \{0,1\}^{\lambda}$ would consist of C^s, g^s and $M \oplus \mathsf{H}\big(e(g^{\alpha}, U_n)^s\big)$, where $\mathsf{H} : \mathbb{G}_T \to \{0,1\}^{\lambda}$ is a pairwise-independent hash function. The decryption key for a vector $\boldsymbol{y}$ is nothing but $W_{z,1}$. For a valid key, the fact that $z = \langle \boldsymbol{x}, \boldsymbol{y} \rangle \neq 0$ enables us to recover the blinding factor on the message from $e(g^{\alpha}, U_n)^{zs}$.

Setup(λ, n)**:** Takes as input n, the dimension of the inner product space. Choose bilinear groups $(\mathbb{G}, \mathbb{G}_T)$ of composite order $N = p_1 p_2 p_3$, where $p_i > 2^{l(\lambda)}$ for each $i \in \{1,2,3\}$, for a suitable polynomial $l : \mathbb{N} \to \mathbb{N}$. Define the bilinear map $e : \mathbb{G} \times \mathbb{G} \to \mathbb{G}_T$. We consider inner products defined over $\mathbb{Z}_N^n$. Choose $g, u \xleftarrow{R} \mathbb{G}_{p_1}$, $R_3 \xleftarrow{R} \mathbb{G}_{p_3}$ and $\alpha, \gamma \xleftarrow{R} \mathbb{Z}_N$ at random in order to define

$$G_1 = g^{\alpha}, \qquad G_2 = g^{(\alpha^2)}, \qquad \ldots \qquad , G_n = g^{(\alpha^n)}$$

and

$$U_1 = u^{\alpha} \cdot R_{3,1}, \qquad U_2 = u^{(\alpha^2)} \cdot R_{3,2}, \qquad \ldots \qquad , U_n = u^{(\alpha^n)} \cdot R_{3,n}$$
$$U_{n+2} = u^{(\alpha^{n+2})} \cdot R_{3,n+2}, \qquad \ldots \qquad , U_{2n} = u^{(\alpha^{2n})} \cdot R_{3,2n},$$

where $R_{3,j} \xleftarrow{R} \mathbb{G}_{p_3}$ for each $j \in [1, 2n] \setminus \{n+1\}$. Define the public parameters to consist of

$$\mathsf{mpk} := \Big((\mathbb{G}, \mathbb{G}_T, e), g,\ g^{\gamma},\ \{G_j\}_{j=1}^{n},\ \{U_j\}_{j\in[1,2n]\setminus\{n+1\}},\ \mathsf{H}\Big),$$

where $\mathsf{H} : \mathbb{G}_T \to \{0,1\}^{\lambda}$ is a pairwise-independent hash function. The master secret key is given by $\mathsf{msk} := (u, R_3, \gamma, \alpha)$.

Encrypt$(\mathsf{mpk}, M, \boldsymbol{x} = (x_1, \ldots, x_n))$**:** To encrypt $M \in \{0,1\}^{\lambda}$ under $\boldsymbol{x} \in \mathbb{Z}_N^n$, choose $s \xleftarrow{R} \mathbb{Z}_N$ and define the ciphertext $\mathcal{C}$ to consist of three components – one from $\mathbb{G}_T$ and two from $\mathbb{G}$ given by

$$C_0 = M \oplus \mathsf{H}(e(g,u)^{\alpha^{n+1}s}), \qquad C_1 = g^{s}, \qquad C_2 = g^{s\cdot(\gamma+\sum_{i=1}^{n}\alpha^{i}\cdot x_i)},$$

where C_0 and C_2 are computed as $M \oplus \mathsf{H}\big(e(G_1, U_n)^{s}\big)$ and $(g^{\gamma}\cdot\prod_{i=1}^{n} G_i^{x_i})^{s}$ respectively. The algorithm outputs $\mathcal{C} = (C_0, C_1, C_2)$.

KeyGen$(\mathsf{msk}, \boldsymbol{y})$**:** The secret key for $\boldsymbol{y} = (y_1, \ldots, y_n) \in \mathbb{Z}_N^n$ is given by

$$d_{\boldsymbol{y}} = \left(\prod_{i=1}^{n} u^{\alpha^{i}\cdot y_i}\right)^{\gamma} \cdot X_3,$$

where $X_3 \xleftarrow{R} \mathbb{G}_{p_3}$ is sampled using R_3.

Decrypt$(\mathcal{C}, \boldsymbol{x}, \boldsymbol{y}, d_{\boldsymbol{y}})$**:** Let $z = \langle \boldsymbol{x}, \boldsymbol{y}\rangle \bmod N$. If $z \neq 0$ the algorithm computes $A_i = \prod_{j=1, j\neq i}^{n} U_{n+1+j-i}^{x_j}$ for all $i \in [1,n]$, and recovers $M \in \{0,1\}^{\lambda}$ as

$$M = C_0 \oplus \mathsf{H}\left(\left(\frac{e(C_1, d_{\boldsymbol{y}}\cdot\prod_{i=1}^{n} A_i^{y_i})}{e(C_2, \prod_{i=1}^{n} U_{n-i+1}^{y_i})}\right)^{1/z}\right).$$

Correctness. Correctness follows from the observation that

$$\begin{aligned}
e(C_2, U_{n-i+1}) &= e\left(g^{s\cdot(\gamma+\sum_{i=1}^{n}\alpha^{i}x_i)}, u^{(\alpha^{n-i+1})}\cdot R_{3,n+2}\right)\\
&= e\left(g^{\gamma}\cdot\prod_{i=1}^{n} g^{\alpha^{i}\cdot x_i}, u^{(\alpha^{n-i+1})}\right)^{s}\\
&= e(g,u)^{\alpha^{n+1}\cdot s\cdot x_i}\cdot e\left(g, u_{n-i+1}^{\gamma}\cdot\prod_{j=1, j\neq i}^{n} u^{\alpha^{n+1+j-i}\cdot x_j}\right)^{s}\\
&= e(g,u)^{\alpha^{n+1}\cdot s\cdot x_i}\cdot e\left(g, u_{n-i+1}^{\gamma}\cdot A_i\right)^{s}.
\end{aligned}$$

Raising both sides of the above equality to y_i and taking a product over all $i \in [1,n]$ gives us

$$\begin{aligned}
e\left(C_2, \prod_{i=1}^{n} U_{n-i+1}^{y_i}\right) &= \prod_{i=1}^{n} e(g,u)^{\alpha^{n+1}\cdot s\cdot x_i\cdot y_i}\cdot\prod_{i=1}^{n} e\left(g, u^{(\alpha^{n-i+1})\cdot\gamma}\cdot A_i\right)^{s\cdot y_i}\\
&= e(g,u)^{\alpha^{n+1}\cdot s\cdot\langle\boldsymbol{x},\boldsymbol{y}\rangle}\cdot e\left(g^{s}, \prod_{i=1}^{n} u^{(\alpha^{n-i+1})\cdot\gamma\cdot y_i}\cdot A_i^{y_i}\right)\\
&= e(g,u)^{\alpha^{n+1}\cdot s\cdot z}\cdot e\left(C_1, d_{\boldsymbol{y}}\cdot\prod_{i=1}^{n} A_i^{y_i}\right),
\end{aligned}$$

as required. Note that in the last step, we replaced $\prod_{i=1}^{n} u^{(\alpha^{n-i+1})\cdot\gamma\cdot y_i}$ by $d_{\boldsymbol{y}}$ as the $\mathbb{G}_{p_3}$ component vanishes upon pairing.

Theorem 2. *The NIPE construction is selectively secure if Assumption 1 and Assumption 2 hold.*

Proof. The proof relies on a series of modifications to the distribution of public parameters. To define these alternative distributions, we use a family of functions

$$\{F_k : [1, 2n] \to \mathbb{Z}_{p_2}\}_{k=0}^{2n}$$

such that for all $j \in [1, 2n]$,

$$F_k(j) = \begin{cases} 0 & \text{if } k = 0 \\ \sum_{i=1}^{k} r_j \cdot \alpha_i^j \bmod p_2 & \text{if } k \in [1, 2n] \end{cases}$$

where $r_1, \ldots, r_{2n}, \alpha_1, \ldots, \alpha_{2n}$ are randomly distributed in $\mathbb{Z}_{p_2}$. The modified distributions are defined on the parameters $\{U_j\}_{j=1}^{2n}$.

Type k parameters ($0 \le k \le 2n$): are parameters where elements $\{U_i\}_{i\in[1,2n]}$ have a $\mathbb{G}_{p_2}$ component determined by the function $F_k(.)$: namely,

$$U_i = u^{(\alpha^i)} \cdot g_2^{F_k(i)} \cdot R_{3,i} \qquad \forall i \in [1, 2n].$$

The proof proceeds through a sequence of $2n + 4$ games denoted G_0, G_1, G_2, $G_{3,1}, \ldots, G_{3,2n}$, G_4 as defined below. Let $\mathsf{win}_\square$ denote the event that the adversary $\mathcal{A}$ wins in game $G_\square$.

Game G_0: is the real attack game (described in Sect. 2.2).

Game G_1: This game is similar to G_0 except for the following changes. At the beginning of the game, the challenger chooses $\tilde{\gamma} \xleftarrow{R} \mathbb{Z}_N$ and sets $\gamma = \tilde{\gamma} - \sum_{i=1}^{n} \alpha^i x_i^\star$ where $\boldsymbol{x}^\star = (x_1^\star, \ldots, x_n^\star)$ is the challenge vector. The public parameter g^γ is generated as $g^{\tilde{\gamma}} \cdot \prod_{i=1}^{n} G_i^{-x_i^\star}$. The challenge ciphertext is computed as:

$$C_1 \xleftarrow{R} \mathbb{G}_{p_1}, \qquad C_2 = C_1^{\tilde{\gamma}}, \qquad C_0 = M_\beta \oplus \mathsf{H}\big(e(C_1, U_{n+1})\big).$$

Since γ is known to the challenger, secret key queries can be answered by running the KeyGen algorithm. The change is only conceptual and hence $\Pr[\mathsf{win}_0] = \Pr[\mathsf{win}_1]$.

Game G_2: In this game, we start modifying the distribution of the challenge ciphertext. Namely, the challenger now picks C_1 uniformly at random in $\mathbb{G}_{p_1 p_2}$ instead of $\mathbb{G}_{p_1}$. The adversary's ability to distinguish between games G_1 and G_2 can be leveraged to break Assumption 1 as formalized in the following lemma.

Lemma 1. *If Assumption 1 holds, then* $|\Pr[\mathit{win}_1] - \Pr[\mathit{win}_2]|$ *is negligible.*

Game $G_{3,k}$ **for** $k = 1, \ldots, 2n$: We let game $G_{3,0}$ be identical to G_2 for notational convenience. In game $G_{3,k}$ the adversary is given Type k parameters. We argue that the adversary can detect this change with negligible probability if Assumption 2 holds.

Lemma 2. *If Assumption 2 holds, then* $|\Pr[\mathsf{win}_{3,k-1}] - \Pr[\mathsf{win}_{3,k}]|$ *is negligible for each* $k \in [1, 2n]$.

In game $G_{3,2n}$ the parameters U_j have their $\mathbb{G}_{p_2}$ components defined by $F_{2n}(j)$, which is a $2n$-wise independent function from $[1, 2n]$ to $\mathbb{Z}_{p_2}$. The adversary's view thus remains identical if we replace the function F_{2n} by a truly random function $RF : [1, 2n] \to \mathbb{Z}_{p_2}$ which allows defining the $\mathbb{G}_{p_2}$ component of U_j as $g_2^{RF(j)}$ for each $j \in [1, 2n]$.

Game G_4: This game is identical to game $G_{3,2n}$ with the difference that, in the challenge ciphertext, C_0 is chosen as a random string in $\{0,1\}^\lambda$. We argue that any legitimate adversary's view remains statistically close to that of game $G_{3,2n}$. To see this, we first note that the $\mathbb{G}_{p_2}$ components of the secret keys contain linear combinations of $RF(j)$ in the exponent excluding $RF(n+1)$. Indeed, recall that the adversary can only make private key queries on vectors $\boldsymbol{y}$ such that $\langle \boldsymbol{y}, \boldsymbol{x}^\star \rangle = 0$. Programming γ as $\gamma = \tilde{\gamma} - \sum_{i=1}^{n} \alpha^i \cdot x_i^\star$ requires the creation of a $\mathbb{G}_{p_1}$ component with the exponent

$$\left(\sum_{i=1}^{n} y_i \cdot \alpha^{n-i+1} \right) \cdot \left(\tilde{\gamma} - \sum_{i=1}^{n} \alpha^i \cdot x_i^\star \right),$$

in order to generate a secret key for $\boldsymbol{y}$. Note that the coefficient of α^{n+1} is $\langle \boldsymbol{y}, \boldsymbol{x}^\star \rangle$ which is 0 for all legal private key queries. Hence, the private key $d_{\boldsymbol{y}}$ can be computed without using U_{n+1}, ensuring that $RF(n+1)$ remains completely independent of any information revealed to $\mathcal{A}$. As a result, the distribution of

$$\mathsf{H}\big(e(C_1, U_{n+1})\big) = \mathsf{H}\big(e(C_1, u^{\alpha^{n+1}}) \cdot e(C_1, g_2^{RF(n+1)})\big)$$

is statistically uniform over $\{0,1\}^\lambda$ as long as C_1 as a non-trivial $\mathbb{G}_{p_2}$ component (which occurs with probability $1 - 1/p_2$). This follows from the fact that, if $e(C_1, g_2) \neq 1_{\mathbb{G}_T}$, the $\mathbb{G}_{p_2}$ component of $e(C_1, g_2^{RF(n+1)})$ has $\log(p_2)$ bits of min-entropy. Since $\mathsf{H} : \mathbb{G}_T \to \{0,1\}^\lambda$ is a pairwise-independent hash function, the Leftover Hash Lemma ensures that, conditionally on the adversary's view, the distribution of $\mathsf{H}\big(e(C_1, u^{\alpha^{n+1}}) \cdot e(C_1, g_2^{RF(n+1)})\big)$ is within distance $2^{-\lambda}$ from the uniform distribution over $\{0,1\}^\lambda$. This implies that $|\Pr[\mathsf{win}_{3,2n}] - \Pr[\mathsf{win}_4]| \leq 1/p_2 + 1/2^\lambda$, which is statistically negligible as claimed. Since $\beta \in \{0,1\}$ is perfectly hidden from the adversary in G_4, we have $\Pr[\mathsf{win}_4] = 1/2$.

Combining the above, we find

$$\mathsf{Adv}_{\mathrm{NIPE},\mathcal{A}}(\lambda) = |\Pr[\mathsf{win}_0] - \Pr[\mathsf{win}_4]| \leq \mathsf{Adv}^1_{\mathcal{G},\mathcal{B}}(\lambda) + 2n \cdot \mathsf{Adv}^2_{\mathcal{G},\mathcal{B}}(\lambda) + \frac{1}{p_2} + \frac{1}{2^\lambda}$$

which is negligible in the security parameter λ provided Assumption 1 and Assumption 2 both hold in $(\mathbb{G}, \mathbb{G}_T)$. □

Proof (of Lemma 1). Let (g, g_3, T) be an instance of Assumption 1. We show how $\mathcal{B}$ simulates the different stages of the security game.

Initialize: $\mathcal{A}$ commits to the challenge vector $\boldsymbol{x}^\star = (x_1^\star, \ldots, x_n^\star)$.

Setup: Pick $u \xleftarrow{R} \mathbb{G}_{p_1}$, $\alpha \xleftarrow{R} \mathbb{Z}_N$ and compute $G_j = g^{\alpha^j}$ for $j = 1, \ldots, n$, $U_j = u^{\alpha^j} \cdot R_{3,j}$ for $j \in [1, 2n]$ where $R_{3,j}$'s are sampled from $\mathbb{G}_{p_3}$ using g_3. Choose $\tilde{\gamma} \xleftarrow{R} \mathbb{Z}_N$ and set $\gamma = \tilde{\gamma} - \sum_{i=1}^{n} \alpha^i \cdot x_i^\star$. The adversary is given the following public parameters

$$\mathsf{mpk} := \big(g,\ g^\gamma,\ \{G_j\}_{j=1}^n,\ \{U_j\}_{j \in [1,2n] \setminus \{n+1\}},\ \mathsf{H}\big).$$

Key Extraction: Upon a query on vector $\boldsymbol{y} \in \mathbb{Z}_N^n$, the adversary is given $d_{\boldsymbol{y}} = \left(u^{\sum_{i=1}^{n} \alpha^{n-i+1} \cdot y_i}\right)^\gamma \cdot X_3$, where $X_3 \xleftarrow{R} \mathbb{G}_{p_3}$.

Challenge: $\mathcal{A}$ provides two messages M_0, M_1. $\mathcal{B}$ picks $\beta \xleftarrow{R} \{0, 1\}$ and computes the ciphertext $\mathcal{C}^\star = (C_0, C_1, C_2)$, where,

$$C_1 = T, \qquad C_2 = T^{\tilde{\gamma}}, \qquad C_0 = M_\beta \oplus \mathsf{H}\big(e(C_1, U_{n+1})\big).$$

Guess: $\mathcal{A}$ returns a bit β'. $\mathcal{B}$ returns 1 if $\beta = \beta'$ and 0 otherwise.

If $T \xleftarrow{R} \mathbb{G}_{p_1}$, then $\mathcal{C}^\star$ is distributed as in G_1. Otherwise, $T \xleftarrow{R} \mathbb{G}_{p_1 p_2}$ and $\mathcal{B}$ simulates G_2. We have

$$\begin{aligned} |\Pr[\mathsf{win}_1] - \Pr[\mathsf{win}_2]| &= |\Pr[\beta = \beta' | T \xleftarrow{R} \mathbb{G}_{p_1}] - \Pr[\beta = \beta' | T \xleftarrow{R} \mathbb{G}_{p_1 p_2}]| \\ &= |\Pr[\mathcal{B} \text{ returns } 1 | T \xleftarrow{R} \mathbb{G}_{p_1}] - \Pr[\mathcal{B} \text{ returns } 1 | T \xleftarrow{R} \mathbb{G}_{p_1 p_2}]| \\ &= \mathsf{Adv}^1_{\mathcal{G},\mathcal{B}}(\lambda)\ , \end{aligned}$$

which is negligible under Assumption 1. □

Proof (of Lemma 2). Using $\mathcal{A}$ show how to construct an algorithm $\mathcal{B}$ that breaks Assumption 2. $\mathcal{B}$ receives an instance $(g, X_1 X_2, g_3, Y_2 Y_3, T)$ of the problem and simulates the game as follows. Suppose that $T = u \cdot g_2^{r_2} \cdot g_3^{r_3}$ where either $r_2 = 0$ or $r_2 \xleftarrow{R} \mathbb{Z}_{p_2}$.

Initialize: $\mathcal{A}$ commits to the challenge vector $\boldsymbol{x}^\star = (x_1^\star, \ldots, x_n^\star)$.

Setup: Pick $\alpha \xleftarrow{R} \mathbb{Z}_N, r'_1, \ldots, r'_{k-1} \xleftarrow{R} \mathbb{Z}_N$ and compute $G_j = g^{\alpha^j}$ for $j = 1, \ldots, n$ and

$$U_j = T^{\alpha^j} \cdot (Y_2 Y_3)^{\sum_{i=1}^{k-1} r'_i \cdot \alpha_i^j} \cdot R'_{3,j}$$

for $j \in [1, 2n]$ where $R'_{3,j} \xleftarrow{R} \mathbb{G}_{p_3}$. Choose $\tilde{\gamma} \xleftarrow{R} \mathbb{Z}_N$ and set $\gamma = \tilde{\gamma} - \sum_{i=1}^{n} \alpha^i x_i^\star$. The adversary is given the following public parameters

$$\mathsf{mpk} := \big(g,\ g^\gamma,\ \{G_j\}_{j=1}^n,\ \{U_j\}_{j \in [1,2n] \setminus \{n+1\}},\ \mathsf{H}\big).$$

Key Extraction: Upon a query on vector $\boldsymbol{y} \in \mathbb{Z}_N^n$, the adversary is given $d_{\boldsymbol{y}} = \left(\prod_{i=1}^n U_{n-i+1}^{y_i}\right)^{\gamma} \cdot X_3'$, where $X_3' \xleftarrow{R} \mathbb{G}_{p_3}$.

Challenge: $\mathcal{A}$ provides two messages M_0, M_1. $\mathcal{B}$ picks $\beta \xleftarrow{R} \{0,1\}$ and computes the ciphertext $\mathcal{C}^\star = (C_0, C_1, C_2)$, where,

$$C_1 = X_1 X_2, \qquad C_2 = (X_1 X_2)^{\tilde{\gamma}}, \qquad C_0 = M_\beta \oplus \mathsf{H}\big(e(C_1, U_{n+1})\big).$$

Guess: $\mathcal{A}$ returns a bit β'. $\mathcal{B}$ returns 1 if $\beta = \beta'$ and 0 otherwise.

If $r_2 = 0$, then the parameters have the Type $k-1$ distribution. Otherwise, $r_2 \xleftarrow{R} \mathbb{Z}_{p_2}$ and the parameters have the Type k distribution for reasons explained next. The $\mathbb{G}_{p_2}$-components of U_j (for $j \in [1, 2n]$) would be given by

$$g_2^{r_2 \cdot \alpha^j} \cdot Y_2^{\sum_{i=1}^{k-1} r_i \cdot \alpha_i^j}. \tag{7}$$

All the information provided to $\mathcal{A}$ is independent of $\alpha \bmod p_2$ (by the Chinese Remainder Theorem) and hence we can substitute $\alpha \bmod p_2$ with a uniformly random $\alpha_k \in \mathbb{Z}_{p_2}$. The $\mathbb{G}_{p_2}$ component of U_j in (7) can thus be replaced by

$$g_2^{\sum_{i=1}^{k} r_i \cdot \alpha_i^j}.$$

as required. Moreover, the $\mathbb{G}_{p_3}$ component of U_j is uniformly distributed since we randomize it by $R_{3,j}'$. We thus have

$$|\Pr[\mathsf{win}_{3,k-1}] - \Pr[\mathsf{win}_{3,k}]| \leq \mathsf{Adv}^2_{\mathcal{G},\mathcal{B}}(\lambda),$$

which is negligible under Assumption 2. □

Acknowledgements. The authors were funded by the "Programme Avenir Lyon Saint-Etienne de l'Université de Lyon" in the framework of the programme "Investissements d'Avenir" (ANR-11-IDEX-0007). Jie Chen was also supported in part by the National Natural Science Foundation of China (Grant No. 61472142).

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)
2. Agrawal, S., Freeman, D.M., Vaikuntanathan, V.: Functional encryption for inner product predicates from learning with errors. In: Lee, D.H., Wang, X. (eds.) ASIACRYPT 2011. LNCS, vol. 7073, pp. 21–40. Springer, Heidelberg (2011)
3. Agrawal, S., Libert, B., Stehlé, D.: Fully secure functional encryption for inner products, from standard assumptions. In: Robshaw, M., Katz, J. (eds.) CRYPTO 2016. LNCS, vol. 9816, pp. 333–362. Springer, Heidelberg (2016). doi:10.1007/978-3-662-53015-3_12. Cryptology ePrint Archive: Report 2015/608
4. Attrapadung, N., Libert, B.: Functional encryption for inner product: achieving constant-size ciphertexts with adaptive security or support for negation. In: Nguyen, P.Q., Pointcheval, D. (eds.) PKC 2010. LNCS, vol. 6056, pp. 384–402. Springer, Heidelberg (2010)

5. Attrapadung, N., Libert, B., de Panafieu, E.: Expressive key-policy attribute-based encryption with constant-size ciphertexts. In: Catalano, D., Fazio, N., Gennaro, R., Nicolosi, A. (eds.) PKC 2011. LNCS, vol. 6571, pp. 90–108. Springer, Heidelberg (2011)
6. Bishop, A., Jain, A., Kowalczyk, L.: Function-hiding inner product encryption. In: Iwata, T., et al. (eds.) ASIACRYPT 2015. LNCS, vol. 9452, pp. 470–491. Springer, Heidelberg (2015). doi:10.1007/978-3-662-48797-6_20
7. Boneh, D., Boyen, X.: Short signatures without random oracles. In: Cachin, C., Camenisch, J.L. (eds.) EUROCRYPT 2004. LNCS, vol. 3027, pp. 56–73. Springer, Heidelberg (2004)
8. Boneh, D., Boyen, X.: Efficient selective-ID secure identity-based encryption without random oracles. In: Cachin, C., Camenisch, J.L. (eds.) EUROCRYPT 2004. LNCS, vol. 3027, pp. 223–238. Springer, Heidelberg (2004)
9. Boneh, D., Boyen, X., Shacham, H.: Short group signatures. In: Franklin, M. (ed.) CRYPTO 2004. LNCS, vol. 3152, pp. 41–55. Springer, Heidelberg (2004)
10. Boneh, D., Boyen, X., Goh, E.-J.: Hierarchical identity based encryption with constant size ciphertext. In: Cramer, R. (ed.) EUROCRYPT 2005. LNCS, vol. 3494, pp. 440–456. Springer, Heidelberg (2005)
11. Boneh, D., Goh, E.-J., Nissim, K.: Evaluating 2-DNF formulas on ciphertexts. In: Kilian, J. (ed.) TCC 2005. LNCS, vol. 3378, pp. 325–341. Springer, Heidelberg (2005)
12. Boneh, D., Gentry, C., Waters, B.: Collusion resistant broadcast encryption with short ciphertexts and private keys. In: Shoup, V. (ed.) CRYPTO 2005. LNCS, vol. 3621, pp. 258–275. Springer, Heidelberg (2005)
13. Camenisch, J., Kohlweiss, M., Soriente, C.: An accumulator based on Bilinear maps and efficient revocation for anonymous credentials. In: Jarecki, S., Tsudik, G. (eds.) PKC 2009. LNCS, vol. 5443, pp. 481–500. Springer, Heidelberg (2009)
14. Catalano, D., Fiore, D.: Concise vector commitments and their applications to zero-knowledge elementary databases. In: Cryptology ePrint Archive: Report 2011/495 (2011)
15. Chen, J., Gay, R., Wee, H.: Improved dual system ABE in prime-order groups via predicate encodings. In: Oswald, E., Fischlin, M. (eds.) EUROCRYPT 2015. LNCS, vol. 9057, pp. 595–624. Springer, Heidelberg (2015)
16. Chase, M., Meiklejohn, S.: Déjà Q: using dual systems to revisit q-type assumptions. In: Nguyen, P.Q., Oswald, E. (eds.) EUROCRYPT 2014. LNCS, vol. 8441, pp. 622–639. Springer, Heidelberg (2014)
17. Cheon, J.H.: Security analysis of the strong Diffie-Hellman problem. In: Vaudenay, S. (ed.) EUROCRYPT 2006. LNCS, vol. 4004, pp. 1–11. Springer, Heidelberg (2006)
18. Delerablée, C.: Identity-based broadcast encryption with constant size ciphertexts and private keys. In: Kurosawa, K. (ed.) ASIACRYPT 2007. LNCS, vol. 4833, pp. 200–215. Springer, Heidelberg (2007)
19. Dodis, Y., Yampolskiy, A.: A verifiable random function with short proofs and keys. In: Vaudenay, S. (ed.) PKC 2005. LNCS, vol. 3386, pp. 416–431. Springer, Heidelberg (2005)
20. Goyal, V., Pandey, O., Sahai, A., Waters, B.: Attribute-based encryption for fine-grained access control of encrypted data. In: ACM CCS 2006, pp. 89–98 (2006)
21. Izabachène, M., Libert, B., Vergnaud, D.: Block-wise P-signatures and non-interactive anonymous credentials with efficient attributes. In: Chen, L. (ed.) IMACC 2011. LNCS, vol. 7089, pp. 431–450. Springer, Heidelberg (2011)

22. Katz, J., Sahai, A., Waters, B.: Predicate encryption supporting disjunctions, polynomial equations, and inner products. In: Smart, N.P. (ed.) EUROCRYPT 2008. LNCS, vol. 4965, pp. 146–162. Springer, Heidelberg (2008)
23. Lewko, A., Sahai, A., Waters, B.: Revocation systems with very small private keys. In: IEEE Symposium on Security and Privacy 2010, pp. 273–285. IEEE Computer Society (2010)
24. Lewko, A., Okamoto, T., Sahai, A., Takashima, K., Waters, B.: Fully secure functional encryption: attribute-based encryption and (Hierarchical) inner product encryption. In: Gilbert, H. (ed.) EUROCRYPT 2010. LNCS, vol. 6110, pp. 62–91. Springer, Heidelberg (2010)
25. Lewko, A., Waters, B.: New techniques for dual system encryption and fully secure HIBE with short ciphertexts. In: Micciancio, D. (ed.) TCC 2010. LNCS, vol. 5978, pp. 455–479. Springer, Heidelberg (2010)
26. Libert, B., Ramanna, S.C., Yung, M.: Functional commitment schemes: from polynomial commitments to pairing-based accumulators from simple assumptions. In: ICALP 2016 (2016, to appear)
27. Libert, B., Yung, M.: Concise mercurial vector commitments and independent zero-knowledge sets with short proofs. In: Micciancio, D. (ed.) TCC 2010. LNCS, vol. 5978, pp. 499–517. Springer, Heidelberg (2010)
28. Naor, M.: On cryptographic assumptions and challenges. In: Boneh, D. (ed.) CRYPTO 2003. LNCS, vol. 2729, pp. 96–109. Springer, Heidelberg (2003)
29. Naor, M., Pinkas, B.: Efficient trace and revoke schemes. In: Frankel, Y. (ed.) FC 2000. LNCS, vol. 1962, pp. 1–20. Springer, Heidelberg (2001)
30. Okamoto, T., Takashima, K.: Hierarchical predicate encryption for inner-products. In: Matsui, M. (ed.) ASIACRYPT 2009. LNCS, vol. 5912, pp. 214–231. Springer, Heidelberg (2009)
31. Okamoto, T., Takashima, K.: Fully secure functional encryption with general relations from the decisional linear assumption. In: Rabin, T. (ed.) CRYPTO 2010. LNCS, vol. 6223, pp. 191–208. Springer, Heidelberg (2010)
32. Okamoto, T., Takashima, K.: Adaptively attribute-hiding (Hierarchical) inner product encryption. In: Pointcheval, D., Johansson, T. (eds.) EUROCRYPT 2012. LNCS, vol. 7237, pp. 591–608. Springer, Heidelberg (2012)
33. Okamoto, T., Takashima, K.: Fully secure unbounded inner-product and attribute-based encryption. In: Wang, X., Sako, K. (eds.) ASIACRYPT 2012. LNCS, vol. 7658, pp. 349–366. Springer, Heidelberg (2012)
34. Okamoto, T., Takashima, K.: Achieving short ciphertexts or short secret-keys for adaptively secure general inner-product encryption. Des. Codes Crypt. **77**(2–3), 725–771 (2015)
35. Sahai, A., Waters, B.: Fuzzy identity-based encryption. In: Cramer, R. (ed.) EUROCRYPT 2005. LNCS, vol. 3494, pp. 457–473. Springer, Heidelberg (2005)
36. Shen, E., Shi, E., Waters, B.: Predicate privacy in encryption systems. In: Reingold, O. (ed.) TCC 2009. LNCS, vol. 5444, pp. 457–473. Springer, Heidelberg (2009)
37. Waters, B.: Dual system encryption: realizing fully secure IBE and HIBE under simple assumptions. In: Halevi, S. (ed.) CRYPTO 2009. LNCS, vol. 5677, pp. 619–636. Springer, Heidelberg (2009)
38. Wee, H.: Threshold and revocation cryptosystems via extractable hash proofs. In: Paterson, K.G. (ed.) EUROCRYPT 2011. LNCS, vol. 6632, pp. 589–609. Springer, Heidelberg (2011)
39. Wee, H.: Dual system encryption via predicate encodings. In: Lindell, Y. (ed.) TCC 2014. LNCS, vol. 8349, pp. 616–637. Springer, Heidelberg (2014)

40. Wee, H.: Déjà Q: encore! Un Petit IBE. In: Kushilevitz, E., et al. (eds.) TCC 2016-A. LNCS, vol. 9563, pp. 237–258. Springer, Heidelberg (2016). doi:10.1007/978-3-662-49099-0_9
41. Yamada, S., Attrapadung, N., Hanaoka, G., Kunihiro, N.: A framework and compact constructions for non-monotonic attribute-based encryption. In: Krawczyk, H. (ed.) PKC 2014. LNCS, vol. 8383, pp. 275–292. Springer, Heidelberg (2014)