Building Key-Private Public-Key Encryption Schemes

Kenneth G. Paterson and Sriramkrishnan Srinivasan

Information Security Group,
Royal Holloway, University of London,
Egham, Surrey TW20 0EX, U.K.
{kenny.paterson,s.srinivasan}@rhul.ac.uk

Abstract. In the setting of identity-based encryption with multiple trusted authorities, TA anonymity formally models the inability of an adversary to distinguish two ciphertexts corresponding to the same message and identity, but generated using different TA master public-keys. This security property has applications in the prevention of traffic analysis in coalition networking environments. In this paper, we examine the implications of TA anonymity for key-privacy for normal public-key encryption (PKE) schemes. Key-privacy for PKE captures the requirement that ciphertexts should not leak any information about the public-keys used to perform encryptions. Thus key-privacy guarantees recipient anonymity for a PKE scheme. Canetti, Halevi and Katz (CHK) gave a generic transform which constructs an IND-CCA secure PKE scheme using an identity-based encryption (IBE) scheme that is selective-id IND-CPA secure and a strongly secure one-time signature scheme. Their transform works in the standard model (i.e. does not require the use of random oracles). Here, we prove that if the underlying IBE scheme in the CHK transform is TA anonymous, then the resulting PKE scheme enjoys key-privacy. Whilst IND-CCA secure, key-private PKE schemes are already known in the standard-model, our result gives the first generic method of constructing a key-private PKE scheme in the standard model. We then go on to investigate the TA anonymity of multi-TA versions of well-known standard model secure IBE schemes. In particular, we prove the TA anonymity and selective-id IND-CPA security of a multi-TA version of Gentry's IBE scheme. Applying the CHK transform, we obtain a new, efficient key-private, IND-CCA secure PKE scheme in the standard model.

Keywords: public-key encryption, key-privacy, identity-based encryption, multiple trusted authorities, TA anonymity, standard model.

1 Introduction

Building public-key encryption (PKE) schemes that are secure in a very strong sense, satisfying indistinguishability against chosen ciphertext attacks or IND-CCA secure, remains a very active area of research. Only a handful of approaches [20,13,11] are known for constructing IND-CCA secure PKE schemes

without resorting to the Random Oracle Model [3]. In the usual public-key setting, the security notion termed key-privacy has also gained increasing importance in recent years, in the context of anonymous communications [2]. While specific schemes such as ElGamal, Cramer-Shoup and RSA-based schemes are known to be key-private [2], no generic method is known for constructing a key-private PKE scheme.

Following the results of Cocks [12], and the pairing-based solutions of Sakai, Ohgishi and Kasahara [22] and Boneh and Franklin [7], identity-based cryptography (IBC) [23] has become one of the most active areas of cryptographic research. Canetti *et al.* [11] give a generic construction, now called the CHK transform, to obtain an IND-CCA secure PKE scheme from an IBE scheme that is selective-id IND-CPA secure, and a strong one time signature scheme. No mention is made in [11] of the key-privacy of the PKE schemes arising from the CHK transform.

In the world of identity-based encryption (IBE), the setting of multiple trusted authorities has recently been treated rigorously [21]. In this setting, the relatively new security property termed trusted authority (TA) anonymity captures the inability of an adversary to distinguish two ciphertexts corresponding to the same message and identity, but generated using different TA master public-keys. At a high level, in this paper, we prove that a key-private PKE scheme is obtained from the CHK transform if the underlying IBE scheme has a weak form of TA anonymity. Our result gives the first generic method for constructing a PKE scheme in the standard model that is both key-private and IND-CCA secure.

Based on our current results, we argue that the relatively new notion of TA anonymity is not only of interest in the area of anonymous communications, for example in thwarting traffic analysis [21], but also has rather subtle cryptographic implications for schemes that use IBE as a building block. It therefore merits further study. We investigate the TA anonymity of multi-TA versions of well-known IBE schemes in the standard model and we are easily able to show that they do not satisfy the notion of TA anonymity. By contrast, we prove that a multi-TA version of Gentry's IBE scheme [15] is TA anonymous. Instantiating the CHK transform with this multi-TA Gentry scheme gives us a concrete and reasonably efficient key-private, IND-CCA secure PKE scheme in the standard model.

2 Background

Anonymous encryption was historically motivated in the context of mobile communication. In the standard public-key setting, entities $\mathcal{A}$ and $\mathcal{B}$ exchange encrypted messages using each others' public-keys, over a broadcast medium, where eavesdroppers can see all ciphertexts on the network. It is reasonable to assume that $\mathcal{A}$ and $\mathcal{B}$ will want to keep their identities hidden from such eavesdroppers and this is possible only when ciphertexts do not leak information about the public-keys used to create them, a notion formalized as key-privacy in [2].

In almost all the existing literature on IBE, with a small number of exceptions, there is a single TA that issues keys to all the users in the system, and all ciphertexts are created using the public parameters of that single TA. This TA is also known as the private key generator (PKG) in the literature. In this traditional single-TA

identity-based setting, the notions of security roughly equivalent to the IND-CPA and IND-CCA security notions for PKE were first formalized in [7]. In the IND-CPA and IND-CCA games for IBE, the adversary is also given access to a private key extraction oracle with suitable restrictions on its use.

In IBE, the security notion equivalent to key-privacy in PKE is termed recipient anonymity. The systematic study of recipient anonymity was initiated in [1], motivated both by its intrinsic interest in IBE and for its application in constructing public-key encryption with keyword search (PEKS) schemes. Recipient anonymity models the requirement that ciphertexts should not leak the identity of their intended recipients.

It is possible to envisage scenarios with multiple, independent TAs perhaps sharing some common system parameters. The systematic study of security of IBE in this multi-TA setting was initiated in [21]. In such a setting, in addition to the usual IBE security notions of indistinguishability and recipient anonymity, the notion of TA anonymity arises naturally. TA anonymity captures the requirement that an adversary should find it difficult to distinguish ciphertexts produced using different TA master public-keys, even if the ciphertext is for the same message and identity string. TA anonymity has practical significance, again in the context of anonymous communications. For example, if a coalition of TAs operate in a wireless setting where all ciphertexts can be captured from the network by an adversary, and if the ciphertext were to somehow leak the identity of the TA, this would open up avenues for traffic analysis [21]. However, the cryptographic implications of TA anonymity for schemes that use IBE as a building block have as yet not been studied.

3 Our Contributions

We consider the CHK transform in the setting of multiple public keys that is needed when studying key-privacy. This quite naturally gives rise to a multi-TA IBE setting of the type considered in [21]. We show how to modify the CHK construction to reflect this setting. We then prove that the key-privacy of the PKE scheme resulting from our modified CHK transform follows from a weak form of TA anonymity for the underlying multi-TA IBE scheme. Our result gives us the first generic method of constructing a PKE scheme in the standard model that is key-private, as well as being IND-CCA secure.

We note that the transform of Boneh and Katz [8] builds on the ideas of [11] to give a more efficient construction of PKE from IBE. We can prove similar results for the Boneh-Katz transform. (The results from both [11,8] appear in [6].) Due to constraints of space and bearing in mind that the proof of security in [8] is more involved and that the our aim in this paper has been to highlight the significance of TA anonymity, especially with relation to building key-private PKE schemes, we have limited our discussions to the original CHK transform.

To obtain concrete PKE schemes that are key-private and IND-CCA secure, we study the TA anonymity properties of multi-TA versions of the known standard-model IND-CPA secure IBE schemes. We are able to prove that a multi-TA version of the scheme of Gentry [15] is TA anonymous. We are also

able to show that multi-TA versions of the two popular standard model schemes of Boneh and Boyen in [5], termed BB1 and BB2 in the literature, and multi-TA versions of the schemes related to the BB1 scheme, such as those of Waters [24] and Naccache [19], trivially do not meet the notion of TA anonymity.

4 Definitions

In this section, we provide basic definitions needed for the remainder of the paper. Here, we omit standard definitions for PKE, IBE and strongly secure one-time signatures which can be found in the full version of this paper or for example in [6].

Definition 1. A pairing-friendly group generator *PairingGen* is a polynomial time algorithm with input 1^k and output a tuple (G, G_T, e, p, g) . Here G, G_T are groups of prime order p , g generates G , and $e : G \times G \rightarrow G_T$ is a bilinear, non-degenerate and efficiently computable map.

For ease of presentation, we work exclusively in the setting where e is symmetric; our definitions and results can be generalised to the asymmetric setting where $e : G_1 \times G_2 \rightarrow G_T$, with G_1 and G_2 being different groups. Further details concerning the basic choices that are available when using pairings in cryptography can be found in [14].

Definition 2. We define the advantage of an algorithm $\mathcal{A}$ in solving the Truncated Decisional ℓ -Augmented Bilinear Diffie-Hellman Exponent (ℓ -TDABDHE) problem in (G, G_T) to be:

$$\begin{aligned} \text{Adv}_{\mathcal{A}}^{\ell\text{-TDABDHE}}(k) = & |\Pr(\mathcal{A}(g', g'_{(l+2)}, g_1, g_2, \dots, g_l, e(g_{(l+1)}, g')) = 1) \\ & - \Pr(\mathcal{A}(g', g'_{(l+2)}, g_1, g_2, \dots, g_l, Z) = 1)| \end{aligned}$$

where $\alpha \xleftarrow{\$} \mathbb{Z}_p^*$, $Z \xleftarrow{\$} \mathbb{G}_T$, $g' \xleftarrow{\$} \mathbb{G}$, $g_i = g^{(\alpha^i)}$ and $g'_i = g'^{(\alpha^i)}$. Here, we implicitly assume that parameters (G, G_T, e, p, g) are given to $\mathcal{A}$ as additional inputs. The distribution on the left is referred to as P_{ABDHE} and that on the right is referred to as R_{ABDHE} .

We note that this is the same assumption that is used to prove the security of the IBE scheme presented in [15].

Definition 3. We say that the (t, ϵ, ℓ) -TDABDHE problem is hard in (G, G_T) if no t -time algorithm has advantage at least ϵ in solving the ℓ -TDABDHE problem in (G, G_T) .

Definition 4. A function $\epsilon(k)$ is said to be negligible if, for every c , there exists k_c such that $\epsilon(k) \leq k^{-c}$ for every $k \geq k_c$.

5 Multi-TA IBE

A multi-TA IBE scheme is defined in [21] in terms of five algorithms:

- **CommonSetup**: On input 1^k , outputs $params$, a set of system parameters shared by all TAs; $\mathcal{T} = \{ta_i : 1 \leq i \leq n\}$ will represent the set of (labels of) TAs, where $n = n(k) \in \mathbb{N}$.
- **TASetup**: On input $params$, outputs a master public-key mpk (which includes $params$), and a master secret key msk . This algorithm is randomized and executed independently for each TA in $\mathcal{T}$.
- **KeyDer, Enc, Dec**: These are all as per a normal IBE scheme.

Following the reasoning in [21] we note that for the concrete schemes and transforms considered in this paper, common parameters are needed in order to achieve the notion of TA anonymity; doing so without having some (non-trivial) common parameters is an interesting open problem. We note that it is not unreasonable to assume that the different TAs may share some common system parameters (e.g. the output of a pairing parameter generator) [21]. In fact the possibility of TAs sharing parameters becomes much more likely when we consider the greater complexity of setting up an IBE scheme (where consideration has to be given, among other things, to the choice of elliptic curves, groups used in Pairings, the representation of elements etc.) compared to the “relative simplicity” of setting up, say an RSA scheme. The IEEE P1363.3 working group aims to produce a set of standards specific to identity-based cryptography to address these difficulties. Indeed, sharing of common parameters is inevitable if any kind of interoperability is desired between TAs and such scenarios are becoming more and more desirable [4].

5.1 Security Models for Multi-TA IBE

In all the security games that follow, we associate to an adversary $\mathcal{A}$ and a bit $b \in \{0, 1\}$, the advantage of the adversary for a “notion-attack” combination, which is defined to be:

$$\text{Adv}_{\mathcal{A}}^{\text{notion-atk-b}}(k) = \left| \Pr[\text{Exp}_{\mathcal{A}}^{\text{notion-atk-1}}(k) = 1] - \Pr[\text{Exp}_{\mathcal{A}}^{\text{notion-atk-0}}(k) = 1] \right|.$$

A scheme is said to be “notion-atk”-secure if the advantage of all PPT adversaries is negligible as a function of the security parameter k .

We detail the m-IND-RA-TAA-CCA experiment as defined in [21] that simultaneously captures message indistinguishability, recipient anonymity and TA anonymity in the multi-TA IBE setting, against chosen ciphertext adversaries. This model also gives the adversary access to a **Corrupt** oracle that returns the master secret key for a TA of the adversary’s choice.

In the security game defined below, $TASet$ represents the set of TAs that have been corrupted, i.e. queried for their master secret keys, $IDSet_{ta}$ represents the set of identities queried for private keys for each $ta \in \mathcal{T}$, while $CSet_{ta}$ represents the set of identity/ciphertext pairs on which decryption queries have been performed for each $ta \in \mathcal{T}$. In these games, $MPK = \{mpk_{ta} : ta \in \mathcal{T}\}$ and $MSK = \{msk_{ta} : ta \in \mathcal{T}\}$ represent the set of all master public-keys and all master secret keys, respectively.

Experiment $\text{Exp}_A^{\text{m-IND-RA-TAA-CCA-b}}(k)$ $params \leftarrow \text{CommonSetup}(1^k)$ $TASet \leftarrow \emptyset$ $\forall ta \in \mathcal{T}, (mpk_{ta}, msk_{ta}) \leftarrow \text{TASetup}(params),$ $IDSet_{ta} \leftarrow \emptyset$ and $CSet_{ta} \leftarrow \emptyset$ $(ta_0, ta_1, id_0, id_1, m_0, m_1, state) \leftarrow$ $\mathcal{A}^{\text{Corrupt, KeyDer, Dec}}(\text{find}, MPK)$ $c^* \leftarrow \text{Enc}(mpk_{ta_b}, id_b, m_b)$ $b' \leftarrow \mathcal{A}^{\text{Corrupt, KeyDer, Dec}}(\text{guess}, c^*, state)$ If $\{m_0, m_1\} \not\subseteq \text{MsgSp}$ or $ m_0 \neq m_1 $ then Return 0 If $(ta_0 = ta_1$ and $id_0 = id_1$ and $m_0 = m_1)$ then Return 0 If $ta_0 \notin TASet, ta_1 \notin TASet, id_0 \notin IDSet_{ta_0},$ $id_1 \notin IDSet_{ta_1}, (id_0, c^*) \notin CSet_{ta_0}$ and $(id_1, c^*) \notin$ $CSet_{ta_1}$ then Return b' else Return 0	Oracle $\text{Corrupt}(ta)$ $TASet \leftarrow TASet \cup \{ta\}$ Return msk_{ta} Oracle $\text{KeyDer}(ta, id)$ $IDSet_{ta} \leftarrow IDSet_{ta} \cup \{id\}$ $usk_{id, ta} \leftarrow \text{KeyDer}(msk_{ta}, id)$ Return $usk_{id, ta}$ Oracle $\text{Dec}(ta, id, c)$ $CSet_{ta} \leftarrow CSet_{ta} \cup \{(id, c)\}$ $usk_{id, ta} \leftarrow \text{KeyDer}(msk_{ta}, id)$ $m \leftarrow \text{Dec}(mpk_{ta}, usk_{id, ta}, c)$ Return m

By placing suitable restrictions on the m-IND-RA-TAA-CCA security notion, we can define other, weaker security notions appropriate to the multi-TA IBE setting. For example, CPA secure versions can be defined by removing the adversary's access to the decryption oracle. By removing the adversary's access to the **Corrupt** oracle we define "restricted" versions, and appropriate selective-id versions can be defined by having the adversary commit ahead of time to the identities used in the challenge query. Furthermore, setting $m_0 = m_1$, $id_0 = id_1$ or $ta_0 = ta_1$ gives security notions appropriate to specific circumstances. We will elaborate on some of these security models as we encounter them in this paper.

6 Key-Privacy of the CHK Transform

Canetti *et al.* [11] give a construction that builds an IND-CCA secure PKE scheme from a selective-id IND-CPA secure IBE scheme and a strongly secure one-time signature scheme.

We first describe a security notion for PKE which we term IND-IK-CCA and which simultaneously captures message indistinguishability and key-privacy. We then define the selective-id r-m-IND-TAA-CPA security notion for IBE, this being a weakened version of the m-IND-RA-TAA-CCA notion defined above. We then modify the CHK construction from [11] to reflect the setting of multiple users. Finally we show that the IND-IK-CCA security of the public-key encryption scheme built using the (modified) CHK transform follows from the selective-id r-m-IND-TAA-CPA security of the underlying IBE scheme.

We note that we do not require recipient anonymity of the IBE scheme to obtain our result. Rather, the security property needed from the IBE scheme is the form TA anonymity which is captured in our selective-id r-m-IND-TAA-CPA security notion. In section 7.1 we will prove that a multi-TA version of Gentry's IBE scheme meets the stronger m-IND-RA-TAA-CPA security notion. This is sufficient

for the application of our result. Instantiating the CHK transform with the multi-TA Gentry scheme and any strongly secure one-time signature scheme will give us a concrete construction of a key-private and IND-CCA secure PKE scheme.

6.1 IND-IK-CCA Security for PKE

Bellare *et al.* [2] define two notions, IK-CPA and IK-CCA security, that capture the notions of key-privacy under chosen plaintext attacks and chosen ciphertext attacks, respectively. For our purposes, we define a combined security notion which simultaneously captures both message indistinguishability and key-privacy. We term this IND-IK-CCA security.

Experiment $\text{Exp}_A^{\text{IND-IK-CCA-b}}(k)$ $I \xleftarrow{\\$} \text{CommonSetup}(1^k)$ $(PK_0, SK_0) \xleftarrow{\\$} \text{KeyGen}(I)$ $(PK_1, SK_1) \xleftarrow{\\$} \text{KeyGen}(I)$ $CSet_{SK_0} \leftarrow \emptyset, CSet_{SK_1} \leftarrow \emptyset$ $(m_0, m_1, state) \leftarrow \mathcal{A}^{\text{Dec}}(\text{find}, PK_0, PK_1)$ $c^* \leftarrow \text{Enc}(PK_b, m_b)$ $b' \leftarrow \mathcal{A}^{\text{Dec}}(\text{guess}, c^*, state)$ If $m_0 \neq m_1$ or $m_0 = m_1$ then Return 0 If $c^* \notin CSet_{SK_0}$ and $c^* \notin CSet_{SK_1}$ then Return b' else Return 0	Oracle $\text{Dec}(PK_b, c)$ $CSet_{SK_b} \leftarrow CSet_{SK_b} \cup \{c\}$ $m \leftarrow \text{Dec}(SK_b, c)$ Return m

6.2 Security for Multi-TA IBE

We define the selective-id r-m-IND-TAA-CPA security notion for multi-TA IBE. A single identity is used in the challenge phase in this model, i.e. $id_0 = id_1$. Furthermore, the adversary commits to this identity at the start of the game. The adversary is not allowed to make decryption or **Corrupt** queries.

Experiment $\text{Exp}_A^{\text{s-id r-m-IND-TAA-CPA-b}}(k)$ $id^* \leftarrow \mathcal{A}(1^k)$ $params \leftarrow \text{CommonSetup}(1^k)$ $TASet \leftarrow \emptyset$ $\forall ta \in \mathcal{T}, (mpk_{ta}, msk_{ta}) \leftarrow \text{TASetup}(params),$ $IDSet_{ta} \leftarrow \emptyset$ $(ta_0, ta_1, m_0, m_1, state) \leftarrow$ $\quad \mathcal{A}^{\text{KeyDer}}(\text{find}, MPK)$ $c^* \leftarrow \text{Enc}(mpk_{ta_b}, id^*, m_b)$ $b' \leftarrow \mathcal{A}^{\text{KeyDer}}(\text{guess}, c^*, state)$ If $\{m_0, m_1\} \not\subseteq \text{MsgSp}$ or $m_0 \neq m_1$ or $m_0 = m_1$ then Return 0 If $ta_0 = ta_1$ then Return 0 If $ta_0 \notin TASet, ta_1 \notin TASet, id^* \notin IDSet_{ta_0},$ $id^* \notin IDSet_{ta_1}$ then Return b' else Return 0	Oracle $\text{KeyDer}(ta, id)$ $IDSet_{ta} \leftarrow IDSet_{ta} \cup \{id\}$ $usk_{id, ta} \leftarrow \text{KeyDer}(msk_{ta}, id)$ Return $usk_{id, ta}$

6.3 The Modified CHK Transform

Let $\Pi' = \{\text{CommonSetup}', \text{TAS} \text{etup}, \text{KeyDer}, \text{Enc}', \text{Dec}'\}$ be a multi-TA IBE scheme for identities of length n .

Let $\text{Sig} = \{\text{Gen}, \text{Sgn}, \text{Vrfy}\}$ be a one-time signature scheme in which the verification keys output by Gen have length n .

Define $\Pi = \{\text{CommonSetup}, \text{KeyGen}, \text{Enc}, \text{Dec}\}$ as follows

- **CommonSetup**: Runs $\text{CommonSetup}'$ to obtain params .
- **KeyGen**: Runs $\text{TAS} \text{etup}$ to obtain mpk, msk . The public-key is $PK = \text{mpk}$ (PK includes params , as mpk by definition includes params) and the secret key is $SK = \text{msk}$.
- **Enc**: To encrypt a message m using public-key PK , the sender first runs Gen to obtain a verification key vk and the corresponding signing key sk (with $|vk| = n$). Then, the sender computes $c = \text{Enc}(PK, m) = \text{Enc}'(\text{mpk}, vk, m)$ (i.e. the sender encrypts the message m with respect to identity vk for recipient with public-key $PK = \text{mpk}$) and $\sigma = \text{Sgn}(sk, c)$. The final ciphertext is (vk, c, σ) .
- **Dec**: To decrypt (vk, c, σ) using the secret key msk , the recipient first checks whether $\text{Vrfy}(vk, c, \sigma) \stackrel{?}{=} 1$. If not, the receiver outputs $\perp$. Otherwise, the receiver computes $\text{usk}_{vk} = \text{KeyDer}(\text{msk}, vk)$ and outputs $m = \text{Dec}(SK, c) = \text{Dec}'(\text{mpk}, \text{usk}_{vk}, c)$.

Theorem 1. *If Π' is an IBE scheme which is selective-id r -m-IND-TAA-CPA secure and Sig is a strongly secure one-time signature scheme, then Π is an IND-IK-CCA secure PKE scheme.*

Proof. Our proof follows closely the proof of [11] with suitable modifications to reflect the setting of multiple users. In the following, expressions of the form $\Pr_{A,S}[\text{Event}]$ denote the probability that an **Event** occurs when an adversary $\mathcal{A}$ interacts with a scheme S in a specified security game.

Let $\mathcal{A}$ be an IND-IK-CCA adversary against Π . We say a ciphertext (vk, c, σ) is valid if $\text{Vrfy}(vk, c, \sigma) = 1$. Let (vk^*, c^*, σ^*) denote the challenge ciphertext received by $\mathcal{A}$ during a particular run of the experiment and let **Forge** denote the event that $\mathcal{A}$ submits a valid ciphertext (vk^*, c, σ) to its decryption oracle.

Claim 1: $\Pr_{A,\Pi}[\text{Forge}]$ is negligible.

Proof of Claim 1: $\mathcal{A}$ is an IND-IK-CCA adversary against the PKE scheme Π . We use $\mathcal{A}$ to construct an adversary $\mathcal{F}$ that forges a signature with respect to the one-time signature scheme Sig , with probability $\Pr_{A,\Pi}[\text{Forge}]$.

$\mathcal{F}$ is given a verification key vk . $\mathcal{F}$ first runs KeyGen to obtain (PK_0, SK_0) and (PK_1, SK_1) . It gives $\mathcal{A}$ the two public-keys PK_0 and PK_1 . Note that $\mathcal{F}$ can answer any decryption queries of $\mathcal{A}$.

If $\mathcal{A}$ happens to submit a valid ciphertext (vk^*, c, σ) to its decryption oracle before requesting the challenge ciphertext then $\mathcal{F}$ simply outputs the forgery (c, σ) and stops.

Otherwise, when $\mathcal{A}$ outputs messages m_0 and m_1 , it chooses a random bit b and computes $c^* = \text{Enc}'(mpk_b, vk^*, m_b)$ and obtains from its signing oracle a signature σ^* on the message c^* , i.e. $\sigma^* = \text{Sgn}(sk, c^*)$ where sk is the signing key corresponding to vk . $\mathcal{F}$ gives $\mathcal{A}$ the challenge ciphertext (vk^*, c^*, σ^*) .

Subsequently, if $\mathcal{A}$ submits a valid ciphertext (vk^*, c, σ) to its decryption oracle, (note that we must have $(c, \sigma) \neq (c^*, \sigma^*)$) $\mathcal{F}$ simply outputs (c, σ) as its forgery.

It is easy to see that $\mathcal{F}$'s success probability is exactly $\Pr_{\mathcal{A}, \Pi}[\text{Forge}]$.

Claim 2: $|\Pr_{\mathcal{A}, \Pi}[\text{Succ} \wedge \overline{\text{Forge}}] + \frac{1}{2} \Pr_{\mathcal{A}, \Pi}[\text{Forge}] - \frac{1}{2}|$ is negligible.

Proof of Claim 2: We now use $\mathcal{A}$ to construct a selective-id r-m-IND-TAA-CPA attacker $\mathcal{B}$ against the IBE scheme Π' .

Adversary $\mathcal{B}$ acts as a Challenger for $\mathcal{A}$ as follows.

$\mathcal{B}$ runs $\text{Gen}(1^k)$ to obtain (sk^*, vk^*) and outputs a target identity $id^* = vk^*$ to its Challenger $\mathcal{C}$.

$\mathcal{C}$ gives $\mathcal{B}$ MPK , the set of all master public-keys in the multi-TA IBE scheme. Adversary $\mathcal{B}$ gives $\mathcal{A}$ the two public-keys $PK_0 = mpk_{ta_0}$ and $PK_1 = mpk_{ta_1}$.

$\mathcal{A}$ is a IND-IK-CCA attacker against the public-key scheme. When $\mathcal{A}$ makes decryption queries on ciphertexts of the form (vk, c, σ) , it specifies whether it wants the decryption corresponding to PK_0 or PK_1 . $\mathcal{B}$ answers decryption queries as follows.

- If $vk = vk^*$ then $\mathcal{B}$ checks whether $\text{Vrfy}(vk^*, c, \sigma) = 1$. In this case, $\mathcal{B}$ does not know the corresponding IBE secret key corresponding to the identity vk^* and it is not allowed to make this query to its Challenger $\mathcal{C}$. Consequently, $\mathcal{B}$ aborts and outputs a random bit. If $\text{Vrfy}(vk^*, c, \sigma) \neq 1$ then $\mathcal{B}$ responds with $\perp$.
- If $vk \neq vk^*$ and $\text{Vrfy}(vk, c, \sigma) \neq 1$ then $\mathcal{B}$ responds with $\perp$.
- If $vk \neq vk^*$ and $\text{Vrfy}(vk, c, \sigma) = 1$ then $\mathcal{B}$,
  - Makes the oracle query $\text{KeyDer}(ta_i, vk)$ where PK_i is specified in the challenge query and obtains usk_{vk, ta_i} .
  - Computes $m = \text{Dec}'(mpk_{ta_i}, usk_{vk, ta_i}, c)$ and responds with m .

At some point during the simulation $\mathcal{A}$ outputs two equal length messages m_0 and m_1 . $\mathcal{B}$ forwards (ta_0, m_0) and (ta_1, m_1) to its Challenger. $\mathcal{B}$ is given the challenge ciphertext $c^* = \text{Enc}'(mpk_{ta_b}, id^*, m_b)$. $\mathcal{B}$ computes $\sigma^* = \text{Sgn}(sk^*, c^*)$ and gives $\mathcal{A}$ (vk^*, c^*, σ^*) .

$\mathcal{A}$ continues to make decryption oracle queries which are answered by $\mathcal{B}$ as before.

Finally $\mathcal{A}$ outputs a guess b' and this same guess is output by $\mathcal{B}$ and $\mathcal{B}$ wins if $b' = b$. We note that $\mathcal{B}$ provides a perfect simulation for $\mathcal{A}$ as well as a legal strategy for attacking the IBE scheme. In particular it never requests the secret key corresponding to the target identity vk^* for either of the target TAs.

Therefore we have

$$|\Pr_{\mathcal{B}, \Pi'}[\text{Succ}] - \frac{1}{2}| = |\Pr_{\mathcal{A}, \Pi}[\text{Succ} \wedge \overline{\text{Forge}}] + \frac{1}{2} \cdot \Pr_{\mathcal{A}, \Pi}[\text{Forge}] - \frac{1}{2}|.$$

Claim 2 then follows as we know the left hand side of the above equation is negligible by the assumed security of the IBE scheme.

Finally, we have

$$\begin{aligned} & |\Pr_{\mathcal{A}, \Pi}[\text{Succ}] - \tfrac{1}{2}| \\ & \leq |\Pr_{\mathcal{A}, \Pi}[\text{Succ} \wedge \text{Forge}] - \tfrac{1}{2} \cdot \Pr_{\mathcal{A}, \Pi}[\text{Forge}]| \\ & \quad + |\Pr_{\mathcal{A}, \Pi}[\text{Succ} \wedge \overline{\text{Forge}}] + \tfrac{1}{2} \Pr_{\mathcal{A}, \Pi}[\text{Forge}] - \tfrac{1}{2}| \\ & \leq \tfrac{1}{2} \cdot \Pr_{\mathcal{A}, \Pi}[\text{Forge}] + |\Pr_{\mathcal{A}, \Pi}[\text{Succ} \wedge \overline{\text{Forge}}] + \tfrac{1}{2} \Pr_{\mathcal{A}, \Pi}[\text{Forge}] - \tfrac{1}{2}|. \end{aligned}$$

The proof of the result follows from the proofs of claims 1 and 2.

7 Anonymity of Standard Model IBE Schemes

To obtain a standard-model-secure key-private PKE scheme by applying the CHK transform, we need a multi-TA IBE scheme that is suitably TA anonymous under chosen plaintext attacks, in the standard model. While the TA anonymity of multi-TA versions of some popular IBE schemes in the Random Oracle Model has been previously studied in [21], the TA anonymity of multi-TA versions of standard model IBE schemes has not as yet been investigated.

A multi-TA version of the BB1 IBE scheme from [5] can be sketched similar to the multi-TA version of Gentry's IBE scheme in section 7.1

We can easily show that such a multi-TA BB1 scheme is not TA anonymous. Let us consider an adversary that requests the encryption of a message m to identity id in either ta_0 or ta_1 , in its challenge. The challenge ciphertext it receives is of the form:

$$c^* = (\hat{e}(g_1, g_2)^s \cdot m, g^s, F(id)^s) = (A, B, C).$$

Since

$$params_{ta_0} = (params, g_1, g_2, \hat{e}(g_1, g_2), h, F)$$

and

$$params_{ta_1} = (params, g'_1, g'_2, \hat{e}(g'_1, g'_2), h', F')$$

the adversary simply checks if

$$\hat{e}(B, g_1^{id} \cdot h) = \hat{e}(C, g) \quad \text{or} \quad \hat{e}(B, g_1'^{id} \cdot h') = \hat{e}(C, g)$$

to find, with overwhelming probability, which TA's parameters were used.

Multi-TA analogues of schemes related to the BB1 scheme, such as those of Waters [24] and Naccache [19], as well as the multi-TA analogue of the BB2 scheme [5], are also not TA anonymous for similar reasons.

The original scheme by Gentry [15] is recipient anonymous and we now show that a multi-TA version of this scheme is TA anonymous.

7.1 Multi-TA Gentry

We first sketch a multi-TA version of Gentry's IBE scheme. We assume identities are elements in $\mathbb{Z}_p^*$ and messages are elements in $\mathbb{G}_T$. Later, we will need identities that are bit-strings of a fixed length; such identities can easily and securely be converted into elements of $\mathbb{Z}_p^*$ by applying a suitable collision-resistant hash function.

CommonSetup(1^k): – $(G, G_T, e, p, g) \leftarrow \text{PairingGen}(1^k)$. – Output $params = (G, G_T, e, p, g)$. TASetup($params$): – Pick $\alpha \xleftarrow{\\$} \mathbb{Z}_p^*$. Set $g_1 = g^\alpha$. – Pick $h \xleftarrow{\\$} \mathbb{G}$. – Define function $F : \mathbb{Z}_p^* \rightarrow \mathbb{G}$ st $F(x) = g_1 \cdot g^{-x}$. – Set $mpk = (params, g_1, h, e(g, g), e(g, h), F)$. – Set $msk = \alpha$. – Output (mpk, msk).	KeyDer(ta, id): – Pick $r_{id} \xleftarrow{\\$} \mathbb{Z}_p^*$. – Output $usk_{ta, id} = (r_{id}, h_{id})$ where $h_{id} = (h \cdot g^{-r_{id}})^{\frac{1}{\alpha - id}}$. Enc($ta, id, m$): – Pick $s \xleftarrow{\\$} \mathbb{Z}_p^*$. – Output $c = (F(id)^s, e(g, g)^s, e(g, h)^{-s} \cdot m)$. Dec($ta, id, c$): – Parse c as (u, v, w). – Parse $usk_{ta, id}$ as (d_0, d_1). – Output $m = w \cdot e(u, h_{id})v^{r_{id}}$.

The Multi-TA Gentry scheme.

Anonymity of Multi-TA Gentry: We will first show that the multi-TA version of Gentry's IBE scheme meets the r-m-IND-RA-TAA-CPA security notion under the q -TDABDHE assumption. This gives us a reduction that has tightness similar to the original single-TA scheme. Our proof follows closely the proof of [15] with suitable modifications to reflect the multi-TA setting.

Theorem 2. *Let $q = q_{id} + 1$ where q_{id} is the maximum number of private key extraction queries allowed by the adversary per TA. Assume the (t, ϵ, q) -TDABDHE assumption holds in $(\mathbb{G}, \mathbb{G}_T)$. Then, the above multi-TA IBE scheme is (t', ϵ', q_{id}) r-m-IND-RA-TAA-CPA secure for $t' = t - O(t_{exp} \cdot n \cdot q^2)$ and $\epsilon' = \epsilon + (4/p)$ where t_{exp} is the time required to exponentiate in $\mathbb{G}$.*

Proof. The proof is given in the appendix.

The above proof can be modified slightly to enable $\mathcal{B}$ to respond to **Corrupt** queries as well, thereby giving us a proof of security for the m-IND-RA-TAA-CPA security for the multi-TA version of Gentry's IBE scheme, under the same assumptions:

Theorem 3. *Let $q = q_{id} + 1$ where q_{id} is the maximum number of private key extraction queries allowed by the adversary per TA. Assume the (t, ϵ, q) -TDABDHE assumption holds in $(\mathbb{G}, \mathbb{G}_T)$. Then, the above multi-TA IBE scheme is (t', ϵ', q_{id}) m -IND-RA-TAA-CPA secure for $t' = t - O(t_{exp} \cdot n \cdot q^2)$ and $\epsilon' = (\epsilon + (4/p)) \cdot \binom{n}{2}$ where t_{exp} is the time required to exponentiate in $\mathbb{G}$.*

Proof. $\mathcal{B}$ simply generates two related q -TDABDHE challenges from the original input challenge and uses these to respond to private key extraction queries for two specific TAs indexed by $ta_x, ta_y \in \mathcal{T}$. It cannot respond to **Corrupt** queries on these two TAs and the success of the proof relies on $\mathcal{A}$ choosing these two TAs in its Challenge query (thereby reducing the tightness of the reduction).

For all other TAs $\{ta_i \in \mathcal{T} : i \neq x, i \neq y\}$, $\mathcal{B}$ simply generates the master public-keys and master secret keys itself and can therefore respond to private key extraction and **Corrupt** queries on these TAs. Further details are similar to the proof of Theorem 2.

Final Observations: The multi-TA version of Gentry’s IBE scheme that we have given meets stronger notions of security than those required for the application of Theorem 1. We can therefore instantiate the modified CHK transform with the multi-TA version of Gentry’s IBE scheme (and any strongly secure one-time signature scheme) to obtain an IND-IK-CCA PKE scheme. This scheme is quite efficient. Ciphertexts consist of 3 group elements plus a signature and a verification key from the one-time signature scheme. Encryption and decryption cost roughly the same as encryption and decryption in Gentry’s scheme, with the additional requirement of generating or verifying one-time signatures.

8 Conclusion and Future Work

We have shown that the key-privacy of the PKE scheme resulting from the application of the CHK transform follows from the TA anonymity of the underlying IBE scheme, giving us the first generic method to construct a key-private PKE scheme. We have investigated various IBE schemes in the standard model and shown that a multi-TA version of Gentry’s IBE scheme meets the notion of TA anonymity. We have also constructed a key-private PKE scheme by instantiating the CHK transform with the multi-TA version of Gentry’s IBE scheme.

We believe that the relatively new notion of TA anonymity in the setting of multiple TAs has rather subtle cryptographic implications on schemes that use IBE as a building block, but these have not been studied rigorously. For example, Holt [16] also considered security of IBE in the multi-TA setting, motivated by earlier work on anonymous credential systems [17,9]. However, the TA anonymity requirements for these applications are yet to be formally investigated.

Acknowledgements

This research was sponsored in part by the US Army Research Laboratory and the UK Ministry of Defence and was accomplished under Agreement Number

W911NF-06-3-0001. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the US Army Research Laboratory, the US Government, the UK Ministry of Defense, or the UK Government. The US and UK Governments are authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation hereon.

The second author is supported by a Dorothy Hodgkin Postgraduate Award, funded by EPSRC and Vodafone and administered by Royal Holloway, University of London.

We are grateful to the anonymous referees for valuable comments and suggestions and Gaven Watson for proofreading parts of this work.

References

1. Abdalla, M., Bellare, M., Catalano, D., Kiltz, E., Kohno, T., Lange, T., Malone-Lee, J., Neven, G., Paillier, P., Shi, H.: Searchable encryption revisited: Consistency properties, relation to anonymous IBE, and extensions. In: Shoup, V. (ed.) CRYPTO 2005. LNCS, vol. 3621, pp. 205–222. Springer, Heidelberg (2005)
2. Bellare, M., Boldyreva, A., Desai, A., Pointcheval, D.: Key-privacy in public-key encryption. In: Boyd, C. (ed.) ASIACRYPT 2001. LNCS, vol. 2248, pp. 566–582. Springer, Heidelberg (2001)
3. Bellare, M., Rogaway, P.: Random oracles are practical: A paradigm for designing efficient protocols. In: ACM Conference on Computer and Communications Security, pp. 62–73 (1993)
4. Boklan, K.D., Klagsbrun, Z., Paterson, K.G., Srinivasan, S.: Flexible and Secure Communications in an Identity-Based Coalition Environment. In: IEEE Military Communications Conference, 2008. MILCOM 2008, pp. 1–6 (2008)
5. Boneh, D., Boyen, X.: Efficient selective-ID secure identity-based encryption without random oracles. In: Cachin and Camenisch [10], pp. 223–238
6. Boneh, D., Canetti, R., Halevi, S., Katz, J.: Chosen-ciphertext security from identity-based encryption. SIAM J. Comput. 36(5), 1301–1328 (2007)
7. Boneh, D., Franklin, M.K.: Identity-based encryption from the Weil pairing. In: Kilian, J. (ed.) CRYPTO 2001. LNCS, vol. 2139, pp. 213–229. Springer, Heidelberg (2001)
8. Boneh, D., Katz, J.: Improved efficiency for CCA-secure cryptosystems built using identity-based encryption. In: Menezes, A. (ed.) CT-RSA 2005. LNCS, vol. 3376, pp. 87–103. Springer, Heidelberg (2005)
9. Bradshaw, R.W., Holt, J.E., Seamons, K.E.: Concealing complex policies with hidden credentials. In: Atluri, V., Pfizmann, B., McDaniel, P.D. (eds.) ACM Conference on Computer and Communications Security, pp. 146–157. ACM, New York (2004)
10. Cachin, C., Camenisch, J.L. (eds.): EUROCRYPT 2004. LNCS, vol. 3027. Springer, Heidelberg (2004)
11. Canetti, R., Halevi, S., Katz, J.: Chosen-ciphertext security from identity-based encryption. In: Cachin and Camenisch [10], pp. 207–222
12. Cocks, C.: An identity based encryption scheme based on quadratic residues. In: Honary, B. (ed.) Cryptography and Coding 2001. LNCS, vol. 2260, pp. 360–363. Springer, Heidelberg (2001)

13. Cramer, R., Shoup, V.: A practical public key cryptosystem provably secure against adaptive chosen ciphertext attack. In: Krawczyk [18], pp. 13–25
14. Galbraith, S.D., Paterson, K.G., Smart, N.P.: Pairings for cryptographers. Cryptology ePrint Archive, Report 2006/165 (2006), <http://eprint.iacr.org/>
15. Gentry, C.: Practical identity-based encryption without random oracles. In: Vaudenay, S. (ed.) EUROCRYPT 2006. LNCS, vol. 4004, pp. 445–464. Springer, Heidelberg (2006)
16. Holt, J.E.: Key privacy for identity based encryption. Cryptology ePrint Archive, Report 2006/120 (2006), <http://eprint.iacr.org/>
17. Holt, J.E., Bradshaw, R.W., Seamons, K.E., Orman, H.K.: Hidden credentials. In: Jajodia, S., Samarati, P., Syverson, P.F. (eds.) WPES, pp. 1–8. ACM Press, New York (2003)
18. Krawczyk, H. (ed.): CRYPTO 1998. LNCS, vol. 1462. Springer, Heidelberg (1998)
19. Naccache, D.: Secure and practical identity-based encryption. Information Security, IET 1(2), 59–64 (2007)
20. Naor, M., Yung, M.: Public-key cryptosystems provably secure against chosen ciphertext attacks. In: STOC, pp. 427–437. ACM Press, New York (1990)
21. Paterson, K.G., Srinivasan, S.: Security and anonymity of identity-based encryption with multiple trusted authorities. In: Galbraith, S.D., Paterson, K.G. (eds.) Pairing 2008. LNCS, vol. 5209, pp. 354–375. Springer, Heidelberg (2008)
22. Sakai, R., Ohgishi, K., Kasahara, M.: Cryptosystems based on pairing. In: The 2000 Symposium on Cryptography and Information Security, Okinawa, Japan, January 2000, pp. 26–28 (2000)
23. Shamir, A.: Identity-based cryptosystems and signature schemes. In: Blakely, G.R., Chaum, D. (eds.) CRYPTO 1984. LNCS, vol. 196, pp. 47–53. Springer, Heidelberg (1985)
24. Waters, B.: Efficient identity-based encryption without random oracles. In: Cramer, R. (ed.) EUROCRYPT 2005. LNCS, vol. 3494, pp. 114–127. Springer, Heidelberg (2005)

Appendix

Proof of Theorem 2

Proof. Let $\mathcal{A}$ be an adversary that (t', ϵ', q_{id}) breaks the r-m-IND-RA-TAA-CPA security of the multi-TA Gentry IBE scheme described. Here q_{id} is the maximum number of private key extraction queries allowed by the adversary per TA. We construct an algorithm $\mathcal{B}$ that solves the q -TDABDHE problem, as follows.

$\mathcal{B}$ takes as input a random q -TDABDHE challenge $(g', g'_{q+2}, g_1, g_2, \dots, g_q, Z)$ where Z is either $e(g_{q+1}, g')$ or a random element of $\mathbb{G}_T$ and the expected additional inputs $(\mathbb{G}, \mathbb{G}_T, e, p, g)$. (Recall that $g_i = g^{(\alpha^i)}$.)

Algorithm $\mathcal{B}$ proceeds as follows.

For an n TA system, $\mathcal{T} = \{ta_i : 1 \leq i \leq n\}$ represents the set of (labels of) TAs, where $n = n(k) \in \mathbb{N}$. $\mathcal{B}$ uses the input challenge to generate n related q -TDABDHE challenges, one for each TA in $\mathcal{T}$. Let CHAL_i denote the q -TDABDHE corresponding to $ta_i \in \mathcal{T}$.

For $ta_i \in \mathcal{T}$, $\mathcal{B}$ first draws $\beta_i \xleftarrow{\$} \mathbb{Z}_p^*$ and sets CHAL_i equal to:

$$(g', g'_{q+2}^{(\beta_i^{(q+2)})}, g_1^{(\beta_i)}, g_2^{(\beta_i^2)}, \dots, g_q^{(\beta_i^q)}, Z^{(\beta_i^{(q+1)})})$$

or

$$(g', g'^{((\alpha\beta_i)^{q+2})}, g'^{((\alpha\beta_i))}, g'^{((\alpha\beta_i)^2)}, \dots, g'^{((\alpha\beta_i)^q)}, Z^{(\beta_i^{(q+1)})}).$$

We make a few important observations. Firstly, note that if $Z = e(g_{q+1}, g')$ then $Z^{(\beta_i^{(q+1)})}$ is the correct response for the corresponding input challenge CHAL_i . That is, if the original input q -TDABDHE challenge is drawn from P_{ABDHE} , then so are all the CHAL_i s. Similarly, if Z is random in $\mathbb{G}_T$ then so is $Z^{(\beta_i^{(q+1)})}$, i.e. if the original challenge is drawn from R_{ABDHE} then so are all the CHAL_i s.

Secondly, we note that the g' value is the same in all the n “related” challenges. This does not present a problem as g' is used only once to construct the challenge ciphertext.

- **CommonSetup:** $\mathcal{B}$ sets $params$ equal to $(\mathbb{G}, \mathbb{G}_T, e, p, g)$.
- **Setup:** For each $ta_i \in \mathcal{T}$, $\mathcal{B}$ generates a random polynomial $f_i(x) \in \mathbb{Z}_p[x]$ of degree q . It sets $h_i = g^{f_i(\alpha\beta_i)}$, computing h_i from $g, g_1^{(\beta_i)}, g_2^{(\beta_i^2)}, \dots, g_q^{(\beta_i^q)}$. It sets the public-key for ta_i to $mpk_i = (params, g_1^{(\beta_i)}, h_i, e(g, g), e(g, h_i), F_i)$ where $F_i : \mathbb{Z}_p^* \rightarrow \mathbb{G}$ is such that $F_i(x) = g_1^{\beta_i} \cdot g^{-x}$ and sends all the n master public-keys to $\mathcal{A}$. Since g, α are uniformly random, the β_i values and the polynomials $f_i(x)$ are chosen uniformly at random, the $g_1^{(\beta_i)}$ and h_i values are also uniformly random. Therefore the master public-keys have a distribution identical to that in an actual construction. (At this stage, we have essentially succeeded in using the single q -TDABDHE challenge to set up n independent TAs.)
- **Phase 1:** $\mathcal{A}$ makes key generation queries on (ta, id) . $\mathcal{B}$ responds to a query on $ta = ta_i \in \mathcal{T}$ and $id \in \mathbb{Z}_p^*$ as follows.
 $\mathcal{B}$ checks if $g^{id} = g^{\alpha\beta_j}$ in each CHAL_j . If the equality holds, this implies that $id = \alpha\beta_j$ and $\mathcal{B}$ uses $\alpha\beta_j$ to solve the q -TDABDHE challenge immediately by computing the target response to the challenge itself.
 Else, let $F_{id,i}(x)$ denote the $q-1$ degree polynomial

$$F_{id,i}(x) = (f_i(x) - f_i(id))/(x - id).$$

$\mathcal{B}$ sets the private key for id in ta_i to

$$(r_{id,i}, h_{id,i}) = (f_i(id), g^{F_{id,i}(\alpha\beta_i)}).$$

This is a valid private key since $g^{F_{id,i}(\alpha\beta_i)} = (h_i \cdot g^{-f_i(id)})^{1/(\alpha\beta_i - id)}$.

- **Challenge:** $\mathcal{A}$ outputs TAs ta_0, ta_1 (which correspond to ta_x and $ta_y \in \mathcal{T}$ respectively), identities id_0, id_1 and messages m_0, m_1 .
 Again, as in phase 1, $\mathcal{B}$ checks if either of g^{id_0} or g^{id_1} is equal to $g^{\alpha\beta_j}$ in each CHAL_j . If the equality holds, this implies that one of id_0 or id_1 is equal to $\alpha\beta_j$ and $\mathcal{B}$ uses $\alpha\beta_j$ to solve the q -TDABDHE challenge immediately by computing the target response to the challenge itself.
 Else, $\mathcal{B}$ generates a bit $b \in \{0, 1\}$ and computes a private key $d_{id_b, x} = (r_{id_b, x}, h_{id_b, x})$ for id_b in ta_x if $b = 0$ or $d_{id_b, y} = (r_{id_b, y}, h_{id_b, y})$ for id_b in ta_y if $b = 1$ as in Phase 1.

Now, let $g_2(x) = x^{(q+2)}$ and $F_{2,id_b}(x) = (g_2(x) - g_2(id_b))/(x - id_b)$ which is a $(q+1)$ degree polynomial. Then $F_{2,id_b}(x)$ can be written as

$$F_{2,id_b}(x) = \sum_{i=0}^{q+1} F_{2,id_b,i} \cdot x^i = x^{q+1} + \sum_{i=0}^q F_{2,id_b,i} \cdot x^i$$

where $F_{2,id_b,i}$ is the coefficient of x^i in $F_{2,id_b}(x)$.

$\mathcal{B}$ sets the ciphertext $c = (u, v, w)$ as follows. In the following $\beta = \beta_x$ corresponding to ta_x if $b = 0$ or $\beta = \beta_y$ corresponding to ta_y if $b = 1$.

$\mathcal{B}$ sets $u = g^{(g_2(\alpha\beta) - g_2(id_b))}$, $v = Z^{(\beta^{(q+1)})} \cdot e(g', \prod_{i=0}^q g^{F_{2,id_b,i} \cdot (\alpha\beta)^i})$ and $w = m_0/(\hat{e}(u, h_{id_0,x}) \cdot v^{r_{id_0,x}})$ if $b = 0$ and $w = m_1/(\hat{e}(u, h_{id_1,y}) \cdot v^{r_{id_1,y}})$ if $b = 1$.

To see that $c = (u, v, w)$ is a valid and appropriately distributed ciphertext when $Z = e(g_{q+1}, g')$, first let $s = \log_g(g') \cdot F_{2,id_b}(\alpha\beta)$.

Note that s is uniformly random as $\log_g(g')$ is uniformly random and the β_i values are chosen uniformly at random. We will show that c is constructed using “implicit” randomness s . Now

$$g' = g^{s/(F_{2,id_b}(\alpha\beta))} = g^{(s(\alpha\beta - id_b))/(g_2(\alpha\beta) - g_2(id_b))}.$$

Therefore, $u = g^{s(\alpha\beta - id_b)}$. If Z is a random element in $\mathbb{G}_T$ then v is random in $\mathbb{G}_T$. On the other hand, if $Z = e(g_{q+1}, g')$, then $v = e(g, g)^s$ since it can be shown that

$$e(g', \prod_{i=0}^q g^{F_{2,id_b,i} \cdot (\alpha\beta)^i}) = e(g', g^{F_{2,id_b}(\alpha\beta) - (\alpha\beta)^{q+1}})$$

and therefore, it can be shown that

$$v = Z^{(\beta^{(q+1)})} \cdot e(g', \prod_{i=0}^q g^{F_{2,id_b,i} \cdot (\alpha\beta)^i}) = e(g, g)^s.$$

Finally, note that for any private key $d_{id_b,i} = (r_{id_b,i}, h_{id_b,i})$ corresponding to $ta_i \in \mathcal{T}$, it can be shown that

$$e(u, h_{id_b,i}) \cdot v^{r_{id_b,i}} = e(g, h_i)^s.$$

Therefore, $w = m_b \cdot e(g, h_x)^{-s}$ if $b = 0$ and $w = m_b \cdot e(g, h_y)^{-s}$ if $b = 1$.

- **Phase 2:** $\mathcal{A}$ continues to make key extraction queries and $\mathcal{B}$ responds as in Phase 1.
- **Guess:** Finally, the adversary outputs a guess $b' \in \{0, 1\}$. If $b = b'$ then $\mathcal{B}$ outputs 0 indicating that $Z = e(g_{q+1}, g')$; otherwise, it outputs 1.

We have already shown that the public-keys and ciphertexts are appropriately distributed. We now show that the private keys issued by $\mathcal{B}$ are appropriately distributed as well. If I_i denotes the set consisting of $\alpha\beta_i$, id_b and all the identities queried by $\mathcal{A}$ for ta_i then $|I_i| \leq (q+1)$. Then, from $\mathcal{A}$'s view the values $\{f_i(a) : a \in I\}$ are uniformly random and independent and this follows from the fact that $f_i(x)$ is a uniformly random polynomial of degree q .

Probability Analysis: As we have already seen, if $Z = e(g_{(q+1)}, g')$, then the simulation is perfect and $\mathcal{A}$ will guess the bit b with probability $(1/2) + \epsilon'$.

On the other hand, if Z is a random element in $\mathbb{G}_T$ then, u, v are uniformly random and independent elements in $\mathbb{G}, \mathbb{G}_T$ respectively. It remains to reason about w .

Now, $w = m_b / (\hat{e}(u, h_{id_b, i}) \cdot v^{r_{id_b, i}})$ where i is x or y corresponding to ta_x or ta_y . The value in the denominator can be expressed as follows:

$$e(u, h_{id_b, i}) \cdot v^{r_{id_b, i}} = e(u, h_i)^{1/(\alpha\beta_i - id_b)} \cdot (v/e(u, g)^{1/(\alpha\beta_i - id_b)})^{f_i(id_b)}.$$

Now $f_i(id_b)$ is independent of $\mathcal{A}$'s view. Therefore as long as the inequalities $v \neq e(u, g)^{1/(\alpha\beta_x - id_0)}$, $v \neq e(u, g)^{1/(\alpha\beta_x - id_1)}$ and $v \neq e(u, g)^{1/(\alpha\beta_y - id_0)}$, $v \neq e(u, g)^{1/(\alpha\beta_y - id_1)}$ hold (and they hold with probability $(1 - 4/p)$), the value $e(u, h_{id_b, i}) \cdot v^{r_{id_b, i}}$ is random and independent of $\mathcal{A}$'s view. Consequently the value w is random and independent of $\mathcal{A}$'s view. This implies that if Z is a random element then (u, v, w) can impart no information regarding the bit b .

Assuming that no queried identity equals $\alpha\beta_j$ such that $g^{\alpha\beta_j}$ is in one of the challenges CHAL_j , (which would only increase $\mathcal{B}$'s success probability), we can see that:

$$\left| \Pr(\mathcal{B}(g', g'_{(q+2)}, g_1, g_2, \dots, g_q, e(g_{(q+1)}, g'), Z) = 1) - 1/2 \right| \leq (4/p)$$

when $(g', g'_{(q+2)}, g_1, g_2, \dots, g_q, e(g_{(q+1)}, g'), Z)$ is sampled from R_{ABDHE} .

However,

$$\left| \Pr(\mathcal{B}(g', g'_{(q+2)}, g_1, g_2, \dots, g_q, e(g_{(q+1)}, g'), Z) = 1) - 1/2 \right| \geq \epsilon'$$

when $(g', g'_{(q+2)}, g_1, g_2, \dots, g_q, e(g_{(q+1)}, g'), Z)$ is sampled from P_{ABDHE} .

Thus, for uniformly random g, g', α, Z we have:

$$\begin{aligned} & \left| \Pr(\mathcal{B}(g', g'_{(l+2)}, g_1, g_2, \dots, g_l, e(g_{(l+1)}, g')) \right. \\ & \left. - \Pr(\mathcal{A}(g', g'_{(l+2)}, g_1, g_2, \dots, g_l, Z)) \right| \geq \epsilon' - (4/p). \end{aligned}$$

Time-Complexity: In the simulation, $\mathcal{B}$'s overhead is dominated by computing $g^{F_{id, i}(\alpha\beta_i)}$ in response to $\mathcal{A}$'s key generation query on identity id for $ta_i \in \mathcal{T}$, where $F_{id, i}(x)$ is a polynomial of degree $(q-1)$. Each such computation requires $O(q)$ exponentiations in $\mathbb{G}$. Since $\mathcal{A}$ makes at most $(q-1)$ such queries for each TA, $t = t' + O(t_{exp} \cdot n \cdot q^2)$.