

Towards Black-Box Accountable Authority IBE with Short Ciphertexts and Private Keys

Benoît Libert¹ and Damien Vergnaud^{2,*}

¹ Université Catholique de Louvain, Microelectronics Laboratory
Place du Levant, 3 – 1348 Louvain-la-Neuve – Belgium

² Ecole Normale Supérieure – C.N.R.S. – I.N.R.I.A.
45, Rue d’Ulm – 75230 Paris CEDEX 05 – France

Abstract. At Crypto’07, Goyal introduced the concept of *Accountable Authority Identity-Based Encryption* as a convenient tool to reduce the amount of trust in authorities in Identity-Based Encryption. In this model, if the Private Key Generator (PKG) maliciously re-distributes users’ decryption keys, it runs the risk of being caught and prosecuted. Goyal proposed two constructions: the first one is efficient but can only trace well-formed decryption keys to their source; the second one allows tracing obfuscated decryption boxes in a model (called *weak black-box* model) where cheating authorities have no decryption oracle. The latter scheme is unfortunately far less efficient in terms of decryption cost and ciphertext size. In this work, we propose a new construction that combines the efficiency of Goyal’s first proposal with a very simple weak black-box tracing mechanism. Our scheme is described in the selective-ID model but readily extends to meet all security properties in the adaptive-ID sense, which is not known to be true for prior black-box schemes.

Keywords: Identity-based encryption, traceability, efficiency.

1 Introduction

Identity-based cryptography, first proposed by Shamir [36], alleviates the need for digital certificates used in traditional public-key infrastructures. In such systems, users’ public keys are public identifiers (*e.g.* email addresses) and the matching private keys are derived by a trusted party called Private Key Generator (PKG). The first practical construction for *Identity-Based Encryption* (IBE) was put forth by Boneh and Franklin [7] – despite the bandwidth-demanding proposal by Cocks [15] – and, since then, a large body of work has been devoted to the design of schemes with additional properties or relying on different algorithmic assumptions [4,5,6,8,11,19,20,32,38].

The original version of this chapter was revised: The copyright line was incorrect. This has been corrected. The Erratum to this chapter is available at DOI: [10.1007/978-3-642-00468-1_29](https://doi.org/10.1007/978-3-642-00468-1_29)

* The first author acknowledges the Belgian National Fund for Scientific Research (F.R.S.-F.N.R.S.) for their financial support and the BCRYPT Interuniversity Attraction Pole. The second author is supported by the European Commission through the IST Program under Contract ICT-2007-216646 ECRYPT II and by the French *Agence Nationale de la Recherche* through the PACE project.

In spite of its appealing advantages, identity-based encryption has not undergone rapid adoption as a standard. The main reason is arguably the fact that it requires unconditional trust in the PKG: the latter can indeed decrypt any ciphertext or, even worse, re-distribute users' private keys. The key escrow problem can be mitigated as suggested in [7] by sharing the master secret among multiple PKGs, but this inevitably entails extra communication and infrastructure. Related paradigms [2,18] strived to remove the key escrow problem but only did so at the expense of losing the benefit of human-memorizable public keys: these models get rid of escrow authorities but both involve traditional (though not explicitly certified) public keys that are usually less convenient to work with than easy-to-remember public identifiers.

In 2007, Goyal [21] explored a new approach to deter rogue actions from authorities. With the *Accountable Authority Identity-Based Encryption* (A-IBE) primitive, if the PKG discloses a decryption key associated with some identity over the Internet, it runs the risk of being caught and sued by the user. A-IBE schemes achieve this goal by means of an interactive private key generation protocol between the user and the PKG. For each identity, there are exponentially-many families of possible decryption keys. The key generation protocol provides the user with a single decryption key while concealing to the PKG the family that this key belongs to. From this private key, the user is computationally unable to find one from a different family. Hence, for a given identity, a pair of private keys from distinct families serves as evidence of a fraudulent PKG. The latter remains able to passively eavesdrop communications but is discouraged to reveal users' private keys. Also, users cannot falsely accuse an honest PKG since they are unable to compute a new key from a different family using a given key.

PRIOR WORKS. Two constructions were given in [21]. The first one (that we call *Goyal-1* hereafter) builds on Gentry's IBE [19] and, while efficient, only allows tracing well-formed decryption keys. This white-box model seems unlikely to suffice in practice since malicious parties can rather release an imperfect and/or obfuscated program that only decrypts with small but noticeable probability. The second scheme of [21] (let us call it *Goyal-2*), constructed on the Sahai-Waters fuzzy IBE [32], can be extended so as to provide weak black-box traceability: even an imperfect pirate decryption box can be traced (based on its input/output behavior) back to its source although traceability is only guaranteed against dishonest PKGs that have no decryption oracle in the attack game. However, *Goyal-2* is somewhat inefficient as decryption requires a number of pairing calculations that is linear in the security parameter. For the usually required security level, ciphertexts contain more than 160 group elements and decryption calculates a product of about 160 pairings.

Subsequently, Au *et al.* [3] described another A-IBE scheme providing retrievability (*i.e.*, a property that prevents the PKG from revealing more than one key for a given identity without exposing its master key) but remained in the white-box model. More recently, Goyal *et al.* [22] modified the *Goyal-2* system using attribute-based encryption techniques [23,32] to achieve full black-box traceability: unlike *Goyal-2*, the scheme of [22] preserves security against dishonest PKGs

that have access to a decryption oracle in the model. While definitely desirable, this property is currently only achievable at the expense of the same significant penalty as in *Goyal-2* [21] in decryption cost and ciphertext size.

OUR CONTRIBUTIONS. We present a very efficient and conceptually simple scheme with weak black-box traceability. We prove its security (in the standard model) under the same assumption as *Goyal-2*. Decryption keys and ciphertexts consist of a constant number of group elements and their length is thus linear in the security parameter λ (instead of quadratic as in *Goyal-2*). Encryption and decryption take $O(\lambda^3)$ -time (w.r.t. $O(\lambda^4)$ in *Goyal-2*) with only two pairing computations as for the latter (against more than 160 in *Goyal-2*).

While presented in the selective-ID security model (where adversaries must choose the identity that will be their prey at the outset of the game) for simplicity, our scheme is easily adaptable to the adaptive-ID model of [7]. In contrast, one of the security properties (*i.e.*, the infeasibility for users to frame innocent PKGs) was only established in the selective-ID setting for known schemes in the black-box model (*i.e.*, *Goyal-2* and its fully black-box extension [22]). Among such schemes, ours thus appears to be the first one that can be tweaked so as to achieve adaptive-ID security against dishonest users.

Our scheme performs almost as well as *Goyal-1* (the main overhead being a long master public key *à la* Waters [38] to obtain the adaptive-ID security). In comparison with the latter, that was only analyzed in a white-box model of traceability, our system provides several other advantages:

- Its security relies on a weaker assumption. So far, the only fully practical A-IBE scheme was resting on assumptions whose strength grows with the number of adversarial queries, which can be as large as 2^{30} as commonly assumed in the literature. Such assumptions are subject to a limited attack [14] that requires a careful adjustment of group sizes (by as much as 50% additional bits) to guarantee a secure use of schemes.
- It remains secure when many users want to run the key generation protocol in a concurrent fashion. *Goyal-1* has a key generation protocol involving zero-knowledge proofs. As its security reductions require to rewind adversaries at each key generation query, security is only guaranteed when the PKG interacts with users sequentially. In inherently concurrent environments like the Internet, key generation protocols should remain secure when executed by many users willing to register at the same time. By minimizing the number of rewinds in reductions, we ensure that our scheme remains secure in a concurrent setting. In these regards, the key generation protocol of *Goyal-2* makes use of oblivious transfers (OT) in sub-protocols. It thus supports concurrency whenever the underlying OT protocol does. As already mentioned however, our scheme features a much better efficiency than *Goyal-2*.
- In a white-box model of traceability, it can be made secure against dishonest PKGs equipped with a decryption oracle [4]. In the following, we nevertheless focus on the (arguably more interesting) weak black-box traceability aspect.

¹ We believe that the *Goyal-1* system can also be modified so as to obtain this property.

ORGANIZATION. In the rest of the paper, section 2 recalls the A-IBE security model defined in [21]. We first analyze the white-box version of our scheme in section 3 and then describe a weak black-box tracing mechanism in section 4.

2 Background and Definitions

SYNTACTIC DEFINITION AND SECURITY MODEL. We recall the definition of A-IBE schemes and their security properties as defined in [21].

Definition 1. An Accountable Authority Identity-Based Encryption scheme (A-IBE) is a tuple (**Setup**, **Keygen**, **Encrypt**, **Decrypt**, **Trace**) of efficient algorithms or protocols such that:

- **Setup** takes as input a security parameter and outputs a master public key mpk and a matching master secret key msk .
- $\text{Keygen}^{(\text{PKG}, \text{U})}$ is an interactive protocol between the public parameter generator PKG and the user U:
 - the common input to PKG and U are: the master public key mpk and an identity ID for which the decryption key has to be generated;
 - the private input to PKG is the master secret key msk .
 Both parties may use a sequence of private coin tosses as additional inputs. The protocol ends with U receiving a decryption key d_{ID} as his private output.
- **Encrypt** takes as input the master public key mpk , an identity ID and a message m and outputs a ciphertext.
- **Decrypt** takes as input the master public key mpk , a decryption key d_{ID} and a ciphertext C and outputs a message.
- **Trace** given the master public key mpk , a decryption key d_{ID} , this algorithm outputs a key family number n_F or the special symbol \perp if d_{ID} is ill-formed.

Correctness requires that, for any outputs (mpk, msk) of **Setup**, any plaintext m and any identity ID, whenever $d_{\text{ID}} \leftarrow \text{Keygen}^{(\text{PKG}(\text{msk}), \text{U})}(\text{mpk}, \text{ID})$, we have

$$\begin{aligned} & \text{Trace}(\text{mpk}, d_{\text{ID}}) \neq \perp, \\ & \text{Decrypt}(\text{mpk}, d_{\text{ID}}, \text{Encrypt}(\text{mpk}, \text{ID}, m)) = m. \end{aligned}$$

The above definition is for the white-box setting. In a black-box model, **Trace** takes as input an identity ID, the corresponding user’s well-formed private key d_{ID} and a decryption box \mathbb{D} that successfully opens a non-negligible fraction ε of ciphertexts encrypted under ID. The output of **Trace** is either “PKG” or “User” depending on which party is found guilty for having crafted \mathbb{D} .

Goyal formalized three security properties for A-IBE schemes. The first one is the standard notion of privacy [7] for IBE systems. As for the other ones, the FindKey game captures the intractability for the PKG to create a decryption key of the same family as the one obtained by the user during the key generation protocol. Finally, the ComputeNewKey game models the infeasibility for users to generate a key $d_{\text{ID}}^{(2)}$ outside the family of the legally obtained one $d_{\text{ID}}^{(1)}$.

Definition 2. An A-IBE scheme is deemed secure if all probabilistic polynomial time (PPT) adversaries have negligible advantage in the following games.

1. **The IND-ID-CCA game.** For any PPT algorithm \mathcal{A} , the model considers the following game, where $\lambda \in \mathbb{N}$ is a security parameter:

$$\begin{array}{l}
 \boxed{\text{Game}_{\mathcal{A}}^{\text{IND-ID-CCA}}(\lambda)} \\
 (\text{mpk}, \text{msk}) \leftarrow \text{Setup}(\lambda) \\
 (m_0, m_1, \text{ID}^*, s) \leftarrow \mathcal{A}^{\text{Dec,KG}}(\text{find}, \text{mpk}) \\
 \left. \begin{array}{l}
 \text{Dec} : (C, \text{ID}) \\
 \quad \rightarrow \text{Decrypt}(\text{mpk}, \text{msk}, \text{ID}, C); \\
 \text{KG} : \text{ID} \rightarrow \text{Keygen}^{(\text{PKG}(\text{msk}), \mathcal{A})}(\text{mpk}, \text{ID}) \\
 \quad \quad \quad // \text{ID} \neq \text{ID}^*
 \end{array} \right| \\
 d^* \xleftarrow{\$} \{0, 1\} \\
 C^* \leftarrow \text{Encrypt}(\text{mpk}, \text{ID}^*, m_{d^*}) \\
 d \leftarrow \mathcal{A}^{\text{Dec,KG}}(\text{guess}, s, C^*) \\
 \left. \begin{array}{l}
 \text{Dec} : (C, \text{ID}) \rightarrow \text{Decrypt}(\text{mpk}, \text{msk}, \text{ID}, C); \\
 \quad \quad \quad // (C, \text{ID}) \neq (C^*, \text{ID}^*) \\
 \text{KG} : \text{ID} \rightarrow \text{Keygen}^{(\text{PKG}(\text{msk}), \mathcal{A})}(\text{mpk}, \text{ID}) \\
 \quad \quad \quad // \text{ID} \neq \text{ID}^*
 \end{array} \right| \\
 \text{return } 1 \text{ if } d = d^* \text{ and } 0 \text{ otherwise.}
 \end{array}$$

\mathcal{A} 's advantage is measured by $\text{Adv}_{\mathcal{A}}^{\text{CCA}}(\lambda) = |\Pr[\text{Game}_{\mathcal{A}}^{\text{CCA}} = 1] - 1/2|$.

The weaker definition of chosen-plaintext security (IND-ID-CPA) is formalized in the same way in [7] but \mathcal{A} is not granted access to a decryption oracle.

2. **The FindKey game.** Let \mathcal{A} be a PPT algorithm. We consider the following game, where $\lambda \in \mathbb{N}$ is a security parameter:

$$\begin{array}{l}
 \boxed{\text{Game}_{\mathcal{A}}^{\text{FindKey}}(\lambda)} \\
 (\text{mpk}, \text{ID}, s_1) \leftarrow \mathcal{A}(\text{setup}, \lambda) \\
 (d_{\text{ID}}^{(1)}, s_2) \leftarrow \text{Keygen}^{(\mathcal{A}(s_1), \cdot)}(\text{mpk}, \text{ID}) \\
 d_{\text{ID}}^{(2)} \leftarrow \mathcal{A}(\text{findkey}, s_1, s_2) \\
 \text{return } 1 \text{ if } \text{Trace}(\text{mpk}, d_{\text{ID}}^{(1)}) = \text{Trace}(\text{mpk}, d_{\text{ID}}^{(2)}) \\
 \quad \quad \quad 0 \text{ otherwise.}
 \end{array}$$

\mathcal{A} 's advantage is now defined as $\text{Adv}_{\mathcal{A}}^{\text{FindKey}}(\lambda) = \Pr[\text{Game}_{\mathcal{A}}^{\text{FindKey}} = 1]$.

Here, the adversary \mathcal{A} acts as a cheating PKG and the challenger emulates the honest user. Both parties engage in a key generation protocol where the challenger obtains a private key for an identity ID chosen by \mathcal{A} . The latter aims at producing a private key corresponding to ID and belonging to the same family as the key obtained by the challenger in the key generation protocol. Such a successful dishonest PKG could disclose user keys without being caught.

Note that, at the beginning of the experiment, \mathcal{A} generates mpk without revealing the master key msk and the challenger runs a sanity check on mpk .

As noted in [21], it makes sense to provide \mathcal{A} with a decryption oracle that undoes ciphertexts using $d_{\text{ID}}^{(1)}$ (and could possibly leak information on the latter’s family) between steps 2 and 3 of the game. We call this enhanced notion FindKey-CCA (as opposed to the weaker one which we call FindKey-CPA).

Finally, in the black-box model, instead of outputting a new key $d_{\text{ID}}^{(2)}$, the dishonest PKG comes up with a decryption box \mathbb{D} that correctly decrypts ciphertexts intended for ID with non-negligible probability ε and wins if the tracing algorithm returns “User” when run on $d_{\text{ID}}^{(1)}$ and with oracle access to \mathbb{D} .

3. The ComputeNewKey game. For a PPT algorithm \mathcal{A} , the model finally considers the following game:

Game $_{\mathcal{A}}^{\text{ComputeNewKey}}(\lambda)$

$(\text{mpk}, \text{msk}) \leftarrow \text{Setup}(\lambda)$

$(d_{\text{ID}^*}^{(1)}, d_{\text{ID}^*}^{(2)}, \text{ID}^*) \leftarrow \mathcal{A}^{\text{KG}}(\text{mpk})$

| $\text{KG} : \text{ID} \dashrightarrow \text{Keygen}^{(\text{PKG}(\text{msk}), \mathcal{A})}(\text{mpk}, \text{ID})$

return 1 if $\text{Trace}(\text{mpk}, d_{\text{ID}^*}^{(1)}) \neq \perp$ and

$\text{Trace}(\text{mpk}, d_{\text{ID}^*}^{(2)}) \notin \{\perp, \text{Trace}(\text{mpk}, d_{\text{ID}^*}^{(1)})\}$

0 otherwise.

\mathcal{A} ’s advantage is $\text{Adv}_{\mathcal{A}}^{\text{ComputeNewKey}}(\lambda) = \Pr[\text{Game}_{\mathcal{A}}^{\text{ComputeNewKey}} = 1]$.

The ComputeNewKey game involves an adversary interacting with a PKG in executions of the key generation protocol and obtaining private keys associated with *distinct* identities of her choosing. The adversary is declared successful if, for some identity that may have been queried for key generation, she is able to find *two* private keys from *distinct* families. Such a pair would allow her to trick a judge into wrongly believing in a misbehavior of the PKG.

In the black-box scenario, the output of the dishonest user consist of a key $d_{\text{ID}^*}^{(1)}$ and a pirate decryption box \mathbb{D} that yields the correct answer with probability ε when provided with a ciphertext encrypted for ID^* . In this case, the adversary wins if the output of $\text{Trace}^{\mathbb{D}}(\text{mpk}, d_{\text{ID}^*}^{(1)})$ is “PKG”.

In [12], Canetti, Halevi and Katz suggested relaxed notions of IND-ID-CCA and IND-ID-CPA security where the adversary has to choose the target identity ID^* ahead of time (even before seeing the master public key mpk). This relaxed model, called “selective-ID” model (or IND-sID-CCA and IND-sID-CPA for short), can be naturally extended to the ComputeNewKey notion.

BILINEAR MAPS AND COMPLEXITY ASSUMPTIONS. We use prime order groups $(\mathbb{G}, \mathbb{G}_T)$ endowed with an efficiently computable map $e : \mathbb{G} \times \mathbb{G} \rightarrow \mathbb{G}_T$ 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. $e(g, h) \neq 1_{\mathbb{G}_T}$ whenever $g, h \neq 1_{\mathbb{G}}$.

In such *bilinear groups*, we assume the hardness of the (now classical) Decision Bilinear Diffie-Hellman problem that has been widely used in the recent years.

Definition 3. Let $(\mathbb{G}, \mathbb{G}_T)$ be bilinear groups of prime order p and $g \in \mathbb{G}$. The **Decision Bilinear Diffie-Hellman Problem (DBDH)** is to distinguish the distributions of tuples $(g^a, g^b, g^c, e(g, g)^{abc})$ and $(g^a, g^b, g^c, e(g, g)^z)$ for random values $a, b, c, z \xleftarrow{\$} \mathbb{Z}_p^*$. The advantage of a distinguisher \mathcal{B} is measured by

$$\text{Adv}_{\mathbb{G}, \mathbb{G}_T}^{\text{DBDH}}(\lambda) = |\Pr[a, b, c \xleftarrow{\$} \mathbb{Z}_p^* : \mathcal{B}(g^a, g^b, g^c, e(g, g)^{abc}) = 1] \\ - \Pr[a, b, c, z \xleftarrow{\$} \mathbb{Z}_p^* : \mathcal{B}(g^a, g^b, g^c, e(g, g)^z) = 1]|.$$

For convenience, we use an equivalent formulation – called *modified DBDH* – of the problem which is to distinguish $e(g, g)^{ab/c}$ from random given (g^a, g^b, g^c) .

3 The Basic Scheme

The scheme mixes ideas from the “commutative-blinding” [4] and “exponent-inversion” [33] frameworks. Private keys have the same shape as in commutative-blinding-based schemes [4, 5, 11, 38]. At the same time, their first element is a product of two terms, the first one of which is inspired from Gentry’s IBE [19].

According to a technique applied in [21], private keys contain a family number t that cannot be tampered with while remaining hidden from the PKG. This family number t is determined by combining two random values t_0 and t_1 respectively chosen by the user and the PKG in the key generation protocol. The latter begins with the user sending a commitment R to t_0 . Upon receiving R , the PKG turns it into a commitment to $t_0 + t_1$ and uses the modified commitment to generate a “blinded” private key d'_{ID} . The user obtains his final key d_{ID} by “unblinding” d'_{ID} thanks to the randomness that was used to compute R .

A difference with *Goyal-1* is that the key family number is perfectly hidden to the PKG and the FindKey-CPA security is unconditional. In the key generation protocol, the user’s first message is a perfectly hiding commitment that comes along with a witness-indistinguishable (WI) proof of knowledge of its opening. In *Goyal-1*, users rather send a deterministic (and thus non-statistically hiding) commitment and knowledge of the underlying value must be proven in zero-knowledge because a proof of knowledge of a discrete logarithm must be simulated (by rewinding the adversary) in the proof of FindKey-CPA security. In the present scheme, the latter does not rely on a specific assumption and we do not need to simulate knowing the solution of a particular problem instance. Therefore, we can dispense with perfectly ZK proofs and settle for a more efficient 3-move WI proof (such as Okamoto’s variant [30] of Schnorr [35]) whereas 4 rounds are needed² using zero-knowledge proofs of knowledge.

² A similar modification can be brought to the key generation protocol of *Goyal-1* to statistically hide the key family number to the PKG and avoid the need for 4-round ZK proofs.

3.1 Description

Setup: given $\lambda \in \mathbb{N}$, the PKG selects bilinear groups $(\mathbb{G}, \mathbb{G}_T)$ of prime order $p > 2^\lambda$ with a random generator $g \xleftarrow{\$} \mathbb{G}$. It chooses $h, Y, Z \xleftarrow{\$} \mathbb{G}$ and $x \xleftarrow{\$} \mathbb{Z}_p^*$ at random. It defines its master key as $\text{msk} := x$ and the master public key is chosen as $\text{mpk} := (X = g^x, Y, Z, h)$.

Keygen^(PKG,U) : to obtain a private key for his identity ID, a user U interacts with the PKG in the following key generation protocol.

1. The user U draws $t_0, \theta \xleftarrow{\$} \mathbb{Z}_p^*$, provides the PKG with a commitment $R = h^{t_0} \cdot X^\theta$ and also gives an interactive witness indistinguishable proof of knowledge of the pair (t_0, θ) , which he retains for later use.
2. The PKG outputs \perp if the proof of knowledge fails to verify. Otherwise, it picks $r', t_1 \xleftarrow{\$} \mathbb{Z}_p^*$ and returns

$$d'_{\text{ID}} = (d'_1, d'_2, d'_3) = \left((Y \cdot R \cdot h^{t_1})^{1/x} \cdot (g^{\text{ID}} \cdot Z)^{r'}, X^{r'}, t_1 \right). \quad (1)$$

3. U picks $r'' \xleftarrow{\$} \mathbb{Z}_p^*$ and computes $d_{\text{ID}} = (d'_1/g^\theta \cdot (g^{\text{ID}} \cdot Z)^{r''}, d'_2 \cdot X^{r''}, d'_3 + t_0)$ which should equal

$$d_{\text{ID}} = (d_1, d_2, d_3) = \left((Y \cdot h^{t_0+t_1})^{1/x} \cdot (g^{\text{ID}} \cdot Z)^r, X^r, t_0 + t_1 \right) \quad (2)$$

where $r = r' + r''$. Then, U checks whether d_{ID} satisfies the relation

$$e(d_1, X) = e(Y, g) \cdot e(h, g)^{d_3} \cdot e(g^{\text{ID}} \cdot Z, d_2). \quad (3)$$

If so, he sets his private key as d_{ID} and the latter belongs to the family of decryption keys identified by $n_F = d_3 = t_0 + t_1$. He outputs \perp otherwise.

Encrypt: to encrypt $m \in \mathbb{G}_T$ given mpk and ID, choose $s \xleftarrow{\$} \mathbb{Z}_p^*$ and compute

$$C = (C_1, C_2, C_3, C_4) = \left(X^s, (g^{\text{ID}} \cdot Z)^s, e(g, h)^s, m \cdot e(g, Y)^s \right).$$

Decrypt: given $C = (C_1, C_2, C_3, C_4)$ and $d_{\text{ID}} = (d_1, d_2, d_3)$, compute

$$m = C_4 \cdot \left(\frac{e(C_1, d_1)}{e(C_2, d_2) \cdot C_3^{d_3}} \right)^{-1} \quad (4)$$

Trace: given a purported private key $d_{\text{ID}} = (d_1, d_2, d_3)$ and an identity ID, check the validity of d_{ID} w.r.t. ID using relation (3). If valid, d_{ID} is declared as a member of the family identified by $n_F = d_3$.

The correctness of the scheme follows from the fact that well-formed private keys always satisfy relation (3). By raising both members of (3) to the power $s \in \mathbb{Z}_p^*$, we see that the quotient of pairings in (4) actually equals $e(g, Y)^s$.

The scheme features about the same efficiency as classical IBE schemes derived from the commutative-blinding framework [4]. Encryption demands no pairing calculation since $e(g, h)$ and $e(g, Y)$ can both be cached as part of the system

parameters. Decryption requires to compute a quotient of two pairings which is significantly faster than two independent pairing evaluations when optimized in the same way as modular multi-exponentiations.

In comparison with the most efficient standard model scheme based on the same assumption (which is currently the first scheme of [4]), the only overhead is a slightly longer ciphertext and an extra exponentiation in \mathbb{G}_T at both ends.

3.2 Security

SELECTIVE-ID SECURITY. We first prove the IND-sID-CPA security under the modified DBDH assumption (mDBDH).

Theorem 1. *The scheme is IND-sID-CPA under the mDBDH assumption.*

Proof. We show how a simulator \mathcal{B} can interact with a selective-ID adversary \mathcal{A} to solve a mDBDH instance ($T_a = g^a, T_b = g^b, T_c = g^c, T \stackrel{?}{=} e(g, g)^{ab/c}$). At the outset of the game, \mathcal{A} announces the target identity ID^* . To prepare mpk , \mathcal{B} chooses $\alpha, \gamma, t^* \stackrel{\$}{\leftarrow} \mathbb{Z}_p^*$ and sets $X = T_c = g^c, h = T_b = g^b, Y = X^\gamma \cdot h^{-t^*}$, and $Z = g^{-\text{ID}^*} \cdot X^\alpha$. The adversary’s view is simulated as follows.

Queries: at any time, \mathcal{A} may trigger an execution of the key generation protocol for an identity $\text{ID} \neq \text{ID}^*$ of her choosing. She then supplies an element $R = h^{t_0} \cdot X^\theta$ along with a WI proof of knowledge of (t_0, θ) . The simulator \mathcal{B} verifies the proof but does not need to rewind the adversary as it can answer the query without knowing (t_0, θ) . To do so, it picks $t_1 \stackrel{\$}{\leftarrow} \mathbb{Z}_p^*$ at random and defines $W = Y \cdot R \cdot h^{t_1}, d'_3 = t_1$. Elements d'_1 and d'_2 are generated as

$$(d'_1, d'_2) = \left((g^{\text{ID}} \cdot Z)^{r'} \cdot W^{-\frac{\alpha}{\text{ID} - \text{ID}^*}}, X^{r'} \cdot W^{-\frac{1}{\text{ID} - \text{ID}^*}} \right) \tag{5}$$

using a random $r' \stackrel{\$}{\leftarrow} \mathbb{Z}_p^*$. If we set $\tilde{r}' = r' - \frac{w}{c(\text{ID} - \text{ID}^*)}$, where $w = \log_g(W)$, we observe that (d'_1, d'_2) has the correct distribution since

$$\begin{aligned} W^{1/c} \cdot (g^{\text{ID}} \cdot Z)^{\tilde{r}'} &= W^{1/c} \cdot (g^{\text{ID} - \text{ID}^*} \cdot X^\alpha)^{\tilde{r}'} \\ &= W^{1/c} \cdot (g^{\text{ID} - \text{ID}^*} \cdot X^\alpha)^{r'} \cdot (g^{\text{ID} - \text{ID}^*})^{-\frac{w}{c(\text{ID} - \text{ID}^*)}} \cdot X^{-\frac{w\alpha}{c(\text{ID} - \text{ID}^*)}} \\ &= (g^{\text{ID}} \cdot Z)^{r'} \cdot W^{-\frac{\alpha}{\text{ID} - \text{ID}^*}} \end{aligned}$$

and $X^{\tilde{r}'} = X^{r'} \cdot (g^c)^{-\frac{w}{c(\text{ID} - \text{ID}^*)}} = X^{r'} \cdot W^{-\frac{1}{\text{ID} - \text{ID}^*}}$. Finally, the “partial private key” (d'_1, d'_2, d'_3) is returned to \mathcal{A} . Note that the above calculation can be carried out without knowing $w = \log_g(W)$ or the representation (t_0, θ) of R w.r.t. to (h, X) and \mathcal{B} does *not* need to rewind \mathcal{A} .

Challenge: when the first stage is over, \mathcal{A} outputs $m_0, m_1 \in \mathbb{G}_T$. At this point, \mathcal{B} picks $r^* \stackrel{\$}{\leftarrow} \mathbb{Z}_p^*$ and defines a private key $(d_1, d_2, d_3) = (g^\gamma \cdot X^{\alpha r^*}, X^{r^*}, t^*)$ for the identity ID^* . It flips a fair coin $d^* \stackrel{\$}{\leftarrow} \{0, 1\}$ and encrypts m_{d^*} as

$$C_1^* = T_a = g^a \quad C_2^* = T_a^\alpha \quad C_3^* = T \quad C_4^* = m_{d^*} \cdot \frac{e(C_1^*, d_1)}{e(C_2^*, d_2) \cdot C_3^{d_3^*}}.$$

We see that (d_1, d_2, d_3) is a valid key for ID^* . Since $g^{\text{ID}^*} \cdot Z = X^\alpha = T_c^\alpha$ and $h = g^b$, $C^* = (C_1^*, C_2^*, C_3^*, C_4^*)$ is a valid encryption of m_{d^*} (with the exponent $s = a/c$) if $T = e(g, g)^{ab/c}$. If T is random, we have $T = e(g, h)^{s'}$ for some random $s' \in \mathbb{Z}_p^*$ and thus $C_4^* = m_{d^*} \cdot e(Y, g)^s \cdot e(g, h)^{(s-s')t^*}$, which means that m_{d^*} is perfectly hidden since t^* is independent of \mathcal{A} 's view.

As usual, \mathcal{B} outputs 1 (meaning that $T = e(g, g)^{ab/c}$) if \mathcal{A} successfully guesses $d' = d^*$ and 0 otherwise. □

In the above proof, the simulator does not rewind the adversary at any time. The scheme thus remains IND-sID-CPA in concurrent environments, where a batch of users may want to simultaneously run the key generation protocol.

Also, the simulator knows a valid private key for each identity. This allows using Cramer-Shoup-like techniques [16,17] as in [19,26] to secure the scheme against chosen-ciphertext attacks. The advantage of this approach, as we show in appendix A, is to provide FindKey-CCA security in a white-box setting.

Unlike the *Goyal-1* scheme, the basic system provides unconditional FindKey-CPA security: after an execution of the key generation protocol, even an all powerful PKG does not have any information on the component d_3 that is eventually part of the private key obtained by the new user.

Theorem 2. *In the information theoretic sense, no adversary has an advantage in the FindKey-CPA game.*

Proof. The proof directly follows from the perfect hiding property of Pedersen's commitment [31] and the perfect witness indistinguishability of the protocol [30] for proving knowledge of a discrete logarithm representation. Since the commitment $R = h^{t_0} \cdot X^\theta$ and the proof of knowledge of (t_0, θ) perfectly hide t_0 to the PKG, all elements of \mathbb{Z}_p^* are equally likely values of $d_3 = t_0 + t_1$ as for the last part of the user's eventual private key. □

Appendix A describes a hybrid variant of the scheme that provides white-box FindKey-CCA security using authenticated symmetric encryption in the fashion of [27,37,24] so as to reject all invalid ciphertexts with high probability.

Theorem 3. *In the selective-ID ComputeNewKey game, any PPT adversary has negligible advantage assuming that the Diffie-Hellman assumption holds.*

Proof. For simplicity, we prove the result using an equivalent formulation of the Diffie-Hellman problem which is to find $h^{1/x}$ given $(g, h, X = g^x)$.

At the outset of the game, \mathcal{A} declares the identity ID^* for which she aims at finding two private keys $d_{\text{ID}^*}^{(1)}, d_{\text{ID}^*}^{(2)}$ comprising distinct values of $d_3 = t$. Then, the simulator \mathcal{B} prepares the PKG's public key as follows. Elements h and X are taken from the modified Diffie-Hellman instance (g, h, X) . As in the proof of theorem 1, \mathcal{B} defines $Z = g^{-\text{ID}^*} \cdot X^\alpha$ for a randomly chosen $\alpha \xleftarrow{\$} \mathbb{Z}_p^*$. To define Y , it chooses random values $\gamma, t'_1 \xleftarrow{\$} \mathbb{Z}_p^*$ and sets $Y = X^\gamma \cdot h^{-t'_1}$.

Queries: in this game, \mathcal{A} is allowed to query executions of the key generation protocol w.r.t. any identity, including ID^* . The only requirement is that queried identities be distinct.

- For an identity $ID \neq ID^*$, \mathcal{B} can proceed *exactly* as suggested by relation (5) in the proof of theorem 1 and does not need to rewind \mathcal{A} .
- When $ID = ID^*$, \mathcal{B} conducts the following steps. When \mathcal{A} supplies a group element $R = h^{t_0} \cdot X^\theta$ along with a WI proof of knowledge of (t_0, θ) , \mathcal{B} uses the knowledge extractor of the proof of knowledge that allows extracting a representation (t_0, θ) of R by rewinding \mathcal{A} . Next, \mathcal{B} computes $t_1 = t'_1 - t_0$ picks $r \xleftarrow{\$} \mathbb{Z}_p^*$ and returns

$$(d'_1, d'_2, d'_3) = (g^{\gamma+\theta} \cdot (g^{ID^*} \cdot Z)^r, X^r, t_1). \quad (6)$$

To see that the above tuple has the appropriate shape, we note that

$$(Y \cdot R \cdot h^{t_1})^{1/x} = (Y \cdot h^{t_0+t_1} \cdot X^\theta)^{1/x} = (Y \cdot h^{t'_1} \cdot X^\theta)^{1/x} = g^{\gamma+\theta}.$$

Output: upon its termination, \mathcal{A} is expected to come up with distinct valid private keys $d_{ID^*}^{(1)} = (d_1^{(1)}, d_2^{(1)}, d_3^{(1)})$ and $d_{ID^*}^{(2)} = (d_1^{(2)}, d_2^{(2)}, d_3^{(2)})$, such that $t = d_3^{(1)} \neq d_3^{(2)} = t'$, for the identity ID^* . Given that we must have

$$\begin{aligned} d_1^{(1)} &= (Y \cdot h^t)^{1/x} \cdot X^{\alpha r} & d_2^{(1)} &= X^r \\ d_1^{(2)} &= (Y \cdot h^{t'})^{1/x} \cdot X^{\alpha r'} & d_2^{(2)} &= X^{r'} \end{aligned}$$

for some values $r, r' \in \mathbb{Z}_p$, \mathcal{B} can extract $h^{1/x} = \left(\frac{d_1^{(1)}/d_2^{(1)\alpha}}{d_1^{(2)}/d_2^{(2)\alpha}} \right)^{\frac{1}{t-t'}}$. \square

We note that, in the above proof, the simulator does not have to rewind all executions of the key generation protocol but only one, when the adversary asks for a private key corresponding to the target identity ID^* (recall that all queries involve distinct identities). Given that the number of rewinds is constant, the proof still goes through when the simulator is presented with many concurrent key generation queries. If other executions of the protocol (that necessarily involve identities $ID \neq ID^*$) are nested within the one being rewinded when dealing with ID^* , the simulator simply runs them as an honest verifier would in the proof of knowledge and calculates the PKG's output as per relation (5) in the proof of theorem 1. Thus, the initial rewind does not trigger any other one and the simulation still takes polynomial time in a concurrent setting.

ADAPTIVE-ID SECURITY. The scheme can obviously be made IND-ID-CPA if Waters' "hash function" $F(ID) = u' \prod_{j=1}^n u_i^{i_j}$ – where $ID = i_1 \dots i_n \in \{0, 1\}^n$ and $(u', u_1, \dots, u_n) \in \mathbb{G}^{n+1}$ is part of mpk – supersedes the Boneh-Boyen identity hashing $F(ID) = g^{ID} \cdot Z$. The function F is chosen so as to equal $F(ID) = g^{J_1(ID)} \cdot X^{J_2(ID)}$ for integer-valued functions J_1, J_2 that are computable by the simulator. The security proof relies on the fact that J_1 is small in absolute value and cancels with non-negligible probability proportional to $1/q(n+1)$, where q is the number of key generation queries.

When extending the proof of theorem 3 to the adaptive setting, an adversary with advantage ε allows solving CDH with probability $O(\varepsilon/q^2(n+1))$. The reason is that the simulator has to guess beforehand which key generation query will

involve the target identity ID^* . If ID^* is expected to appear in the j^{th} query, when the latter is made, \mathcal{B} rewinds \mathcal{A} to extract (t_0, θ) and uses the special value t'_1 to answer the query as per (6). With probability $1/q$, \mathcal{B} is fortunate when choosing $j \xleftarrow{\$} \{1, \dots, q\}$ at the beginning and, again, $J_1(ID^*)$ happens to cancel with probability $O(1/q(n+1))$ for the target identity.

4 Weak Black-Box Traceability

Theorem 3 showed the infeasibility for users to compute another key from a different family given their private key. In these regards, a decryption key implements a “1-copyrighted function” – in the terminology of [25,29] – for the matching identity. Using this property and the perfect white-box FindKey-CPA security, we describe a black-box tracing mechanism that protects the user from a dishonest PKG as long as the latter is withheld access to a decryption oracle.

The tracing strategy is the one used by Kiayias and Yung [25] in 2-user traitor tracing schemes, where the tracer determines which one out of two subscribers produced a pirate decoder. In our setting, one rather has to decide whether an ε -useful decryption device stems from the PKG or the user himself.

Trace ^{\mathbb{D}} (mpk, d_{ID} , ε): given a well-formed private key $d_{ID} = (d_1, d_2, d_3)$ belonging to a user of identity ID and oracle access to a decoder \mathbb{D} that decrypts ciphertexts encrypted for ID with probability ε , conduct the following steps.

- a. Initialize a counter $ctr \leftarrow 0$ and repeat the next steps $L = 16\lambda/\varepsilon$ times:
 1. Choose distinct exponents $s, s' \xleftarrow{\$} \mathbb{Z}_p^*$ at random, compute $C_1 = X^s$, $C_2 = (g^{ID} \cdot Z)^s$ and $C_3 = e(g, h)^{s'}$.
 2. Calculate $C_4 = m \cdot e(C_1, d_1) / (e(C_2, d_2) \cdot C_3^{d_3})$ for a randomly chosen message $m \in \mathbb{G}_T$.
 3. Feed the decryption device \mathbb{D} with (C_1, C_2, C_3, C_4) . If \mathbb{D} outputs $m' \in \mathbb{G}_T$ such that $m' = m$, increment ctr .
- b. If $ctr < 4\lambda$, incriminate the PKG. Otherwise, incriminate the user.

The soundness of this algorithm is proved using a similar technique to [1]. To ensure the independence of iterations, we assume (as in [1]) that pirate devices are stateless, or resettable, and do not retain information from prior queries: each query is answered as if it were the first one.

Theorem 4. *Under the mDBDH assumption, dishonest users have negligible chance to produce a decryption device \mathbb{D} that makes the tracing algorithm incriminate the PKG in the selective-ID ComputeNewKey game.*

Proof. The tracing algorithm points to the PKG if it ends up with a too small value of ctr . The latter can be seen as the sum of $L = 16\lambda/\varepsilon$ independent random variables $X_i \in \{0, 1\}$ having the same expected value p_1 . We have $\mu = \mathbf{E}[ctr] = Lp_1$. The Chernoff bound tells us that, for any real number ω such that $0 \leq \omega \leq 1$, $\Pr[ctr < (1 - \omega)\mu] < \exp(-\mu\omega^2/2)$. Under the mDBDH assumption,

we certainly have $\mathbf{Adv}^{\text{mDBDH}}(\lambda) \leq \varepsilon/2$ (since $\varepsilon/2$ is presumably non-negligible). Lemma 1 shows that $p_1 \geq \varepsilon - \mathbf{Adv}^{\text{mDBDH}}(\lambda)$, which implies

$$\mu = Lp_1 \geq L(\varepsilon - \mathbf{Adv}^{\text{mDBDH}}(\lambda)) \geq \frac{L\varepsilon}{2} = 8\lambda. \tag{7}$$

With $\omega = 1/2$, the Chernoff bound guarantees that

$$\Pr[\text{ctr} < 4\lambda] = \Pr[\text{ctr} < \mu/2] < \exp(-\mu/8) = \exp(-\lambda). \quad \square$$

Lemma 1. *In the selective-ID ComputeNewKey game, if \mathbb{D} correctly opens well-formed ciphertexts with probability ε , the probability that an iteration of the tracing algorithm increases ctr is at least $p_1 \geq \varepsilon - \mathbf{Adv}^{\text{mDBDH}}(\lambda)$.*

Proof. We consider two games called Game₀ and Game₁ where the adversary \mathcal{A} is faced with a ComputeNewKey challenger \mathcal{B} and produces a decryption device \mathbb{D} which is provided with ciphertexts during a tracing stage. In Game₀, \mathbb{D} is given a properly formed encryption of some plaintext m whereas it is given a ciphertext C where C_3 has been changed in Game₁. In either case, we call p_i (with $i \in \{0, 1\}$) the probability that \mathbb{D} returns the plaintext m chosen by \mathcal{B} .

In the beginning of Game₀, \mathcal{A} chooses a target identity ID^* and the challenger \mathcal{B} defines the system parameters as $X = g^c$, $h = g^b$, $Y = X^\gamma \cdot h^{-t^*}$ and $Z = g^{-\text{ID}^*} \cdot X^\alpha$ for random $\alpha, \gamma, t^* \stackrel{\$}{\leftarrow} \mathbb{Z}_p^*$. Then, \mathcal{A} starts making key generation queries that are treated using the same technique as in the proof of theorem 3. Again, \mathcal{B} only has to rewind the WI proof when the query pertains to ID^* .

At the end of the game, \mathcal{A} outputs a decryption box \mathbb{D} that correctly decrypts a fraction ε of ciphertexts. Then, \mathcal{B} constructs a ciphertext C as

$$C_1 = g^a, \quad C_2 = (g^a)^\alpha, \quad C_3 = T, \quad C_4 = m \cdot \frac{e(C_1, d_1)}{e(C_2, d_2) \cdot C_3^{t^*}}$$

where $T \in \mathbb{G}_T$. In Game₀, \mathcal{B} sets $T = e(g, g)^{ab/c}$ so that we have $C_3 = e(g, h)^{a/c}$ and C is a valid ciphertext (for the encryption exponent $s = a/c$) that \mathbb{D} correctly decrypts with probability ε . In this case, \mathbb{D} thus outputs $m' = m \in \mathbb{G}_T$ with probability $p_0 = \varepsilon$. In Game₁, T is chosen as a random element of \mathbb{G}_T and $C = (C_1, C_2, C_3, C_4)$ has the distribution of a ciphertext produced by the tracing stage and \mathbb{D} must output a plaintext $m' = m$ with probability p_1 . It is clear that $|p_0 - p_1| \leq \mathbf{Adv}^{\text{mDBDH}}(\lambda)$ and we thus have $p_1 \geq \varepsilon - \mathbf{Adv}^{\text{mDBDH}}(\lambda)$. \square

The proofs of theorem 4 and lemma 1 readily extend to the adaptive-ID setting using the same arguments as in the last paragraph of section 3. The system thus turns out to be the first scheme that is amenable for weak black-box traceability against dishonest users in the adaptive-ID sense. Due to their reliance on attribute-based encryption techniques (for which only selective-ID adversaries were dealt with so far), earlier black-box or weakly black-box A-IBE proposals [21, 22] are only known to provide selective-ID security against dishonest users.

As for the security against dishonest PKGs, we observed that, in the FindKey-CPA game, the last part $d_3^{(1)} = t$ of the user’s private key is perfectly hidden

to the malicious PKG after the key generation protocol. Then, a pirate decoder \mathbb{D} made by the PKG has negligible chance of decrypting ciphertexts where C_3 is random in the same way as the user would. When the user comes across \mathbb{D} and takes it to the court, the latter runs the tracing algorithm using \mathbb{D} and the user's well-formed key $d_{\mathbb{D}}^{(1)} = (d_1^{(1)}, d_2^{(1)}, d_3^{(1)})$ for which $d_3^{(1)}$ is independent of \mathbb{D} .

Lemma 2. *In the FindKey-CPA game, one iteration of the tracing algorithm increases ctr with probability at most $1/p$.*

Proof. In an iteration of the tracing stage, \mathbb{D} is given $C = (C_1, C_2, C_3, C_4)$ such that $C_1 = X^s$, $C_2 = (g^{\mathbb{D}} \cdot Z)^s$, $C_3 = e(g, h)^{s'}$ and $C_4 = m \cdot e(g, Y)^s \cdot e(g, h)^{(s-s')t}$ for distinct $s, s' \xleftarrow{s} \mathbb{Z}_p^*$. Since \mathbb{D} has no information on $d_3^{(1)} = t$, for any plaintext $m \in \mathbb{G}_T$, there is a value $d_3^{(1)}$ that explains C_4 and it comes that \mathbb{D} returns the one chosen by the tracer with probability $1/p$. \square

Theorem 5. *In the black-box FindKey-CPA game, a dishonest PKG has negligible advantage.*

Proof. The dishonest PKG is not detected if it outputs a decryption box for which the tracing ends with a sufficiently large ctr. From lemma 2, it easily comes that $\Pr[\text{ctr} \geq 4\lambda] \leq \Pr[\text{ctr} \geq 1] \leq L/p = 16\lambda/(\varepsilon p) \leq 16\lambda/(2^\lambda \varepsilon)$. \square

To secure the scheme against chosen-ciphertext attacks and preserve the weak black-box property, we can use the Canetti-Halevi-Katz [13] technique or its optimizations [9,10] that do not affect the tracing algorithm.

5 Conclusion

We described the first A-IBE system allowing for weak black-box traceability while retaining short ciphertexts and private keys. We also suggested a white-box variant that dwells secure against dishonest PKGs equipped with a decryption oracle. In the black-box setting, it remains an open problem to achieve the latter property without significantly degrading the efficiency.

Acknowledgements

We thank Duong Hieu Phan and the anonymous referees for their comments.

References

1. Abdalla, M., Dent, A., Malone-Lee, J., Neven, G., Phan, D.-H., Smart, N.: Identity-Based Traitor Tracing. In: Okamoto, T., Wang, X. (eds.) PKC 2007. LNCS, vol. 4450, pp. 361–376. Springer, Heidelberg (2007)
2. Al-Riyami, S., Paterson, K.: Certificateless Public Key Cryptography. In: Lai, C.-S. (ed.) ASIACRYPT 2003. LNCS, vol. 2894, pp. 452–473. Springer, Heidelberg (2003)

3. Au, M.-H., Huang, Q., Liu, J.-K., Susilo, W., Wong, D.-S., Yang, G.: Traceable and Retrievable Identity-Based Encryption. In: Bellare, S.M., Gennaro, R., Keromytis, A.D., Yung, M. (eds.) ACNS 2008. LNCS, vol. 5037, pp. 94–110. Springer, Heidelberg (2008)
4. 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)
5. Boneh, D., Boyen, X.: Secure Identity-Based Encryption Without Random Oracles. In: Franklin, M. (ed.) CRYPTO 2004. LNCS, vol. 3152, pp. 443–459. Springer, Heidelberg (2004)
6. 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)
7. Boneh, D., Franklin, M.: Identity-Based Encryption from the Weil Pairing. SIAM 32(3), 586–615 (2003); earlier version in Kilian, J. (ed.) CRYPTO 2001. LNCS, vol. 2139, pp. 213–229. Springer, Heidelberg (2001)
8. Boneh, D., Gentry, C., Hamburg, M.: Space-Efficient Identity-Based Encryption Without Pairings. In: FOCS 2007, pp. 647–657 (2007)
9. 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)
10. Boyen, X., Mei, Q., Waters, B.: Direct Chosen Ciphertext Security from Identity-Based Techniques. In: ACM CCS 2005, pp. 320–329 (2005)
11. Boyen, X., Waters, B.: Anonymous Hierarchical Identity-Based Encryption (Without Random Oracles). In: Dwork, C. (ed.) CRYPTO 2006. LNCS, vol. 4117, pp. 290–307. Springer, Heidelberg (2006)
12. Canetti, R., Halevi, S., Katz, J.: A Forward-Secure Public-Key Encryption Scheme. In: Biham, E. (ed.) EUROCRYPT 2003. LNCS, vol. 2656, pp. 254–271. Springer, Heidelberg (2003)
13. Canetti, R., Halevi, S., Katz, J.: Chosen-Ciphertext Security from Identity-Based Encryption. In: Cachin, C., Camenisch, J.L. (eds.) EUROCRYPT 2004. LNCS, vol. 3027, pp. 207–222. Springer, Heidelberg (2004)
14. 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)
15. 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)
16. Cramer, R., Shoup, V.: A Practical Public-Key Cryptosystem Provably Secure Against Adaptive Chosen Ciphertext Attack. In: Krawczyk, H. (ed.) CRYPTO 1998. LNCS, vol. 1462, pp. 13–25. Springer, Heidelberg (1998)
17. Cramer, R., Shoup, V.: Universal Hash Proofs and a Paradigm for Adaptive Chosen Ciphertext Secure Public-Key Encryption. In: Knudsen, L.R. (ed.) EUROCRYPT 2002. LNCS, vol. 2332, pp. 45–64. Springer, Heidelberg (2002)
18. Gentry, C.: Certificate-Based Encryption and the Certificate Revocation Problem. In: Biham, E. (ed.) EUROCRYPT 2003. LNCS, vol. 2656, pp. 272–293. Springer, Heidelberg (2003)
19. Gentry, C.: Practical Identity-Based Encryption Without Random Oracles. In: Vaudenay, S. (ed.) EUROCRYPT 2006. LNCS, vol. 4004, pp. 445–464. Springer, Heidelberg (2006)

20. Gentry, C., Silverberg, A.: Hierarchical ID-Based Cryptography. In: Zheng, Y. (ed.) ASIACRYPT 2002. LNCS, vol. 2501, pp. 548–566. Springer, Heidelberg (2002)
21. Goyal, V.: Reducing Trust in the PKG in Identity-Based Cryptosystems. In: Menezes, A. (ed.) CRYPTO 2007. LNCS, vol. 4622, pp. 430–447. Springer, Heidelberg (2007)
22. Goyal, V., Lu, S., Sahai, A., Waters, B.: Black-Box Accountable Authority Identity Based Encryption. In: ACM-CCS 2008 (2008)
23. 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)
24. Hofheinz, D., Kiltz, E.: Secure Hybrid Encryption from Weakened Key Encapsulation. In: Menezes, A. (ed.) CRYPTO 2007. LNCS, vol. 4622, pp. 553–571. Springer, Heidelberg (2007)
25. Kiayias, A., Yung, M.: Traitor Tracing with Constant Transmission Rate. In: Knudsen, L.R. (ed.) EUROCRYPT 2002. LNCS, vol. 2332, pp. 450–465. Springer, Heidelberg (2002)
26. Kiltz, E., Vahlis, Y.: CCA2 Secure IBE: Standard Model Efficiency through Authenticated Symmetric Encryption. In: Malkin, T.G. (ed.) CT-RSA 2008. LNCS, vol. 4964, pp. 221–238. Springer, Heidelberg (2008)
27. Kurosawa, K., Desmedt, Y.: A New Paradigm of Hybrid Encryption Scheme. In: Franklin, M. (ed.) CRYPTO 2004. LNCS, vol. 3152, pp. 445–456. Springer, Heidelberg (2004)
28. Libert, B., Vergnaud, D.: Towards Black-Box Accountable Authority IBE with Short Ciphertexts and Private Keys. Computing Research Repository, <http://arxiv.org/abs/0807.1775>
29. Naccache, D., Shamir, A., Stern, J.-P.: How to Copyright a Function. In: Imai, H., Zheng, Y. (eds.) PKC 1999. LNCS, vol. 1560, pp. 188–196. Springer, Heidelberg (1999)
30. Okamoto, T.: Provably Secure and Practical Identification Schemes and Corresponding Signature Schemes. In: Brickell, E.F. (ed.) CRYPTO 1992. LNCS, vol. 740, pp. 31–53. Springer, Heidelberg (1993)
31. Pedersen, T.: Non-Interactive and Information-Theoretic Secure Verifiable Secret Sharing. In: Feigenbaum, J. (ed.) CRYPTO 1991. LNCS, vol. 576, pp. 129–140. Springer, Heidelberg (1992)
32. Sahai, A., Waters, B.: Fuzzy Identity-Based Encryption. In: Cramer, R. (ed.) EUROCRYPT 2005. LNCS, vol. 3494, pp. 457–473. Springer, Heidelberg (2005)
33. Sakai, R., Kasahara, M.: ID-based Cryptosystems with Pairing on Elliptic Curve. In: SCIS 2003 (2003), <http://eprint.iacr.org/2003/054>
34. Sarkar, P., Chatterjee, S.: Construction of a Hybrid HIBE Protocol Secure Against Adaptive Attacks. In: Susilo, W., Liu, J.K., Mu, Y. (eds.) ProvSec 2007. LNCS, vol. 4784, pp. 51–67. Springer, Heidelberg (2007)
35. Schnorr, C.P.: Efficient Identification and Signatures for Smart Cards. In: Brassard, G. (ed.) CRYPTO 1989. LNCS, vol. 435, pp. 239–252. Springer, Heidelberg (1990)
36. 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)
37. Shoup, V., Gennaro, R.: A Note on An Encryption Scheme of Kurosawa and Desmedt. Cryptology ePrint Archive: Report 2004/194 (2004)
38. Waters, B.: Efficient Identity-Based Encryption Without Random Oracles. In: Cramer, R. (ed.) EUROCRYPT 2005. LNCS, vol. 3494, pp. 114–127. Springer, Heidelberg (2005)

A A Variant with White-Box FindKey-CCA Security

To achieve IND-sID-CCA2 security, we can hybridize the scheme using an authenticated symmetric encryption scheme (as defined in appendix B) as previously considered in [34,26] in the context of identity-based encryption. The obtained variant is reminiscent of a version of Gentry's IBE described in [26].

Setup: is the same as in section 3 except that the PKG now chooses two elements $Y_A, Y_B \xleftarrow{\$} \mathbb{G}$ instead of a single one Y . An authenticated symmetric encryption scheme (E, D) of keylength $\ell \in \mathbb{N}$, a secure key derivation function $KDF : \mathbb{G}_T \rightarrow \{0, 1\}^\ell$ and a target collision-resistant hash function $H : \{0, 1\}^* \rightarrow \mathbb{Z}_p^*$ are also needed. The master key is set as $\text{msk} := x$ and the global public key is $\text{mpk} := (X = g^x, h, Y_A, Y_B, Z, H, KDF, (E, D))$.

Keygen^(PKG,U): to obtain a private key for his identity ID, a user U interacts with the PKG as follows.

1. U sends $R = h^{t_0} \cdot X^\theta$ to the PKG and proves his knowledge of the underlying pair $(t_0, \theta) \xleftarrow{\$} (\mathbb{Z}_p^*)^2$ in a witness indistinguishable fashion.
2. The PKG outputs \perp if the proof is incorrect. Otherwise, it picks random values $r'_A, t_{A,1}, r'_B, t_B \xleftarrow{\$} \mathbb{Z}_p^*$ and returns

$$d'_{\text{ID},A} = (d'_{A,1}, d'_{A,2}, d'_{A,3}) = \left((Y \cdot R \cdot h^{t_{A,1}})^{1/x} \cdot (g^{\text{ID}} \cdot Z)^{r'_A}, X^{r'_A}, t_{A,1} \right)$$

$$d'_{\text{ID},B} = (d'_{B,1}, d'_{B,2}, d'_{B,3}) = \left((Y_B \cdot h^{t_B})^{1/x} \cdot (g^{\text{ID}} \cdot Z)^{r'_B}, X^{r'_B}, t_B \right)$$

3. U computes $d_{\text{ID},A} = (d'_{A,1}/g^\theta \cdot (g^{\text{ID}} \cdot Z)^{r''_A}, d'_{A,2} \cdot X^{r''_A}, d'_{A,3} + t_0)$ as well as $d_{\text{ID},B} = (d'_{B,1} \cdot (g^{\text{ID}} \cdot Z)^{r''_B}, d'_{B,2} \cdot X^{r''_B}, d_{B,3})$, for randomly chosen $r''_A, r''_B \xleftarrow{\$} \mathbb{Z}_p^*$ so that

$$d_{\text{ID},A} = (d_{A,1}, d_{A,2}, d_{A,3}) = \left((Y_A \cdot h^{t_A})^{1/x} \cdot (g^{\text{ID}} \cdot Z)^{r_A}, X^{r_A}, t_A \right) \quad (8)$$

$$d_{\text{ID},B} = (d_{B,1}, d_{B,2}, d_{B,3}) = \left((Y_B \cdot h^{t_B})^{1/x} \cdot (g^{\text{ID}} \cdot Z)^{r_B}, X^{r_B}, t_B \right) \quad (9)$$

where $t_A = t_0 + t_{A,1}$, $r_A = r'_A + r''_A$ and $r_B = r'_B + r''_B$. He checks whether $d_{\text{ID},A}$ and $d_{\text{ID},B}$ respectively satisfy

$$e(d_{A,1}, X) = e(Y_A, g) \cdot e(h, g)^{d_{A,3}} \cdot e(g^{\text{ID}} \cdot Z, d_{A,2}) \quad (10)$$

$$e(d_{B,1}, X) = e(Y_B, g) \cdot e(h, g)^{d_{B,3}} \cdot e(g^{\text{ID}} \cdot Z, d_{B,2}). \quad (11)$$

If so, he sets his private key as $(d_{\text{ID},A}, d_{\text{ID},B})$ and the latter belongs to the family of decryption key identified by $n_F = d_{A,3} = t_A$.

Encrypt: to encrypt m given mpk and ID, choose $s \xleftarrow{\$} \mathbb{Z}_p^*$ and compute

$$C = (C_1, C_2, C_3, C_4) = \left(X^s, (g^{\text{ID}} \cdot Z)^s, e(g, h)^s, E_K(m) \right)$$

where $K = KDF(e(g, Y_A)^s \cdot e(g, Y_B)^{\kappa s})$ and $\kappa = H(C_1, C_2, C_3)$.

Decrypt: given $C = (C_1, C_2, C_3, C_4)$ and $d_{\text{ID}} = (d_{\text{ID},A}, d_{\text{ID},B})$, compute the plaintext $m = D_K(C_4)$ (which may just be \perp if C_4 is not a valid authenticated encryption) using the key

$$K = \text{KDF} \left(\frac{e(C_1, d_{A,1} \cdot d_{B,1}^\kappa)}{e(C_2, d_{A,2} \cdot d_{B,2}^\kappa) \cdot C_3^{d_{A,3} + \kappa d_{B,3}}} \right) \quad (12)$$

with $\kappa = H(C_1, C_2, C_3)$.

Trace: given an alleged private key $(d_{\text{ID},A}, d_{\text{ID},B})$, with $d_{\text{ID},A} = (d_{A,1}, d_{A,2}, d_{A,3})$, for an identity ID , check the validity of d_{ID} w.r.t. ID using relations (I0)-(I1). If valid, the key is declared as a member of the family $n_F = d_{3,A} = t_A$.

To prove the IND-sID-CCA security, we can apply the technique of [26], which in turn borrows ideas from [27,37,24].

In the chosen-ciphertext scenario, the white-box FindKey security is no longer unconditional but relies on the ciphertext integrity of the symmetric encryption scheme.

Theorem 6. *The scheme is IND-sID-CCA secure in the standard model if the modified DBDH assumption holds, if the symmetric scheme is a secure authenticated encryption scheme, if H is target collision-resistant and if the key derivation function is secure. More precisely, we have*

$$\begin{aligned} \text{Adv}_{\mathcal{A}}^{\text{CCA}}(\lambda, \ell) \leq & \frac{q_d + 2q_d^2}{p} + \text{Adv}^{\text{TCR}}(\lambda) + \text{Adv}^{\text{mDBDH}}(\lambda) + 3q_d \cdot \text{Adv}^{\text{CT-INT}}(\ell) \\ & + (2q_d + 1) \cdot \text{Adv}^{\text{KDF}}(\lambda, \ell) + \text{Adv}^{\text{IND-SYM}}(\ell) \end{aligned}$$

where q_d denotes the number of decryption queries allowed to the adversary \mathcal{A} and the advantage functions against (E, D) are defined in appendix B.

Proof. Given in the full version of the paper [28]. □

Theorem 7. *The scheme is FindKey-CCA secure assuming the security of the key derivation function and the (weak) ciphertext integrity of the symmetric encryption scheme. The advantage of an adversary \mathcal{A} making at most q_d decryption queries is bounded by*

$$\begin{aligned} \text{Adv}_{\mathcal{A}}^{\text{FindKey-CCA}}(\lambda, \ell) \leq & 2 \cdot q_d \cdot \text{Adv}^{\text{CT-INT}}(\ell) \\ & + 2 \cdot q_d \cdot \text{Adv}^{\text{KDF}}(\lambda, \ell) + \frac{2q_d^2 + q_d + 1}{p}. \end{aligned}$$

Proof. Given in appendix C. □

B Authenticated Symmetric Encryption

A symmetric encryption scheme is specified by a pair (E, D) , where E is the encryption algorithm and D is the decryption procedure, and a key space $\mathcal{K}(\ell)$

where $\ell \in \mathbb{N}$ is a security parameter. The security of authenticated symmetric encryption is defined by means of two games that capture the ciphertext indistinguishability and ciphertext (one-time) integrity properties.

Definition 4. *An symmetric encryption scheme is secure in the sense of authenticated encryption if any PPT adversary has negligible advantage in the following games.*

1. **The IND-SYM game.** For any PPT algorithm \mathcal{A} , the model considers the following game, where $\ell \in \mathbb{N}$ is a security parameter:

Game _{\mathcal{A}} ^{IND-SYM}(ℓ)

$$\begin{aligned}
 &K \xleftarrow{\$} \mathcal{K}(\ell) \\
 &(m_0, m_1, s) \leftarrow \mathcal{A}(\text{find}, \ell) \\
 &d^* \xleftarrow{\$} \{0, 1\} \\
 &c^* \leftarrow E_K(m_{d^*}) \\
 &d \leftarrow \mathcal{A}(\text{guess}, s, c^*) \\
 &\text{return } 1 \text{ if } d = d^* \text{ and } 0 \text{ otherwise.}
 \end{aligned}$$

\mathcal{A} 's advantage is $\text{Adv}_{\mathcal{A}}^{\text{IND-SYM}}(\ell) = |\Pr[\text{Game}_{\mathcal{A}}^{\text{IND-SYM}} = 1] - 1/2|$.

2. **The CT-INT game.** Let \mathcal{A} be a PPT algorithm. We consider the following game, where $\ell \in \mathbb{N}$ is a security parameter:

Game _{\mathcal{A}} ^{CT-INT}(ℓ)

$$\begin{aligned}
 &K \xleftarrow{\$} \mathcal{K}(\ell) \\
 &(m, s) \leftarrow \mathcal{A}(\text{find}, \ell) \\
 &c \leftarrow E_K(m) \\
 &c' \leftarrow \mathcal{A}(\text{create}, \ell, c) \\
 &\text{return } 1 \text{ if } c' \neq c \text{ and } D_K(c') \neq \perp \\
 &\quad 0 \text{ otherwise.}
 \end{aligned}$$

\mathcal{A} 's advantage is now defined as $\text{Adv}_{\mathcal{A}}^{\text{CT-INT}}(\ell) = \Pr[\text{Game}_{\mathcal{A}}^{\text{CT-INT}} = 1]$.

The notion of weak ciphertext integrity is defined in the same way but the adversary is not allowed to see an encryption c under the challenge key K .

C Proof of Theorem 7

The proof proceeds again with a sequence of games, in all of which S_i denotes the event that the adversary \mathcal{A} wins.

Game₀: is the FindKey-CCA experiment. The dishonest PKG \mathcal{A} generates the master public key, chooses an identity ID that she wishes to be challenged upon. She interacts with the challenger in a key generation protocol, upon completion of which the challenger \mathcal{B} obtains a decryption key consisting of

two triples $d_{\text{ID},A}^{(1)} = (d_{A,1}^{(1)}, d_{A,2}^{(1)}, d_{A,3}^{(1)})$, $d_{\text{ID},B}^{(1)} = (d_{B,1}^{(1)}, d_{B,2}^{(1)}, d_{B,3}^{(1)})$ that should pass the key sanity check (otherwise, \mathcal{B} aborts). At this stage, \mathcal{A} knows $t_B^{(1)} = d_{B,3}^{(1)}$ but has no information on $d_{A,3}^{(1)} = t_A^{(1)}$ or on the values $r_A = \log_X(d_{A,2}^{(1)})$ and $r_B = \log_X(d_{B,2}^{(1)})$ (by the construction of the key generation protocol). In the next phase, \mathcal{A} starts making a number of decryption queries that the challenger handles using $(d_{\text{ID},A}^{(1)}, d_{\text{ID},B}^{(1)})$. Namely, when queried on a ciphertext $C = (C_1, C_2, C_3, C_4)$, \mathcal{B} calculates

$$\psi = \frac{e(C_1, d_{A,1}^{(1)} \cdot d_{B,1}^{(1)\kappa})}{e(C_2, d_{A,2}^{(1)} \cdot d_{B,2}^{(1)\kappa}) \cdot C_3^{d_{A,3}^{(1)} + \kappa d_{B,3}^{(1)}}},$$

where $\kappa = H(C_1, C_2, C_3)$, $K = \text{KDF}(\psi)$ and $m = \text{DK}(C_4)$ which is returned to \mathcal{A} (and may be \perp if C is declared invalid).

At the end of the game, \mathcal{A} outputs a key $(d_{\text{ID},A}^{(2)}, d_{\text{ID},B}^{(2)})$ and wins if $d_{\text{ID},A}^{(2)}$ parses into $(d_{A,1}^{(2)}, d_{A,2}^{(2)}, d_{A,3}^{(2)})$ such that $d_{A,3}^{(1)} = t_A^{(1)} = t_A^{(2)} = d_{A,3}^{(2)}$.

We note that decryption queries on well-formed ciphertexts do not reveal any information to \mathcal{A} (since all well-formed keys yield the same result). We will show that, provided all ill-formed ciphertexts are rejected by \mathcal{B} , \mathcal{A} still has negligible information on $t_A^{(1)}$ in the end of the game. For convenience, we distinguish two types of invalid ciphertexts: type I ciphertexts (C_1, C_2, C_3, C_4) are such that $\log_X(C_1) \neq \log_{F(\text{ID})}(C_2)$ (and can be told apart from valid ones by checking if $e(C_1, F(\text{ID})) \neq e(X, C_2)$), where $F(\text{ID}) = g^{\text{ID}} \cdot Z$, whereas type II ciphertexts are those for which $\log_X(C_1) = \log_{F(\text{ID})}(C_2) \neq \log_{e(g,h)}(C_3)$.

Game₁: is as Game₀ but \mathcal{B} rejects all type I invalid ciphertexts (that are publicly recognizable). Such a malformed ciphertext comprises elements $C_1 = X^{s_1}$, $C_2 = F(\text{ID})^{s_1 - s'_1}$ and $C_3 = e(g, h)^{s_1 - s''_1}$ where $s'_1 > 0$ and $s''_1 \geq 0$. Hence, the symmetric key K that \mathcal{B} calculates is derived from

$$\psi = e(g, Y_A^{s_1} \cdot Y_B^{\kappa s_1}) \cdot e(F(\text{ID}), X)^{s'_1(r_A + \kappa r_B)} \cdot e(g, h)^{s''_1(t_A^{(1)} + \kappa t_B^{(1)})} \quad (13)$$

where $\kappa = H(C_1, C_2, C_3)$. Upon termination of the key generation protocol, \mathcal{A} has no information on r_A, r_B (as \mathcal{B} re-randomizes its key). Even if κ was the same in all decryption queries (which may happen if these queries all involve identical (C_1, C_2, C_3)), the second term of the product (13) remains almost uniformly random to \mathcal{A} at each new query. Indeed, for each failed one, \mathcal{A} learns at most one value that is not $r_A + \kappa r_B$. After i attempts, $p - i$ candidates are left and the distance between the uniform distribution on \mathbb{G}_T and that of $e(F(\text{ID}), X)^{s'_1(r_A + \kappa r_B)}$ becomes at most $i/p \leq q_d/p$. Then, the only way for \mathcal{A} to cause the new rejection rule to apply is to forge a symmetric authenticated encryption for an essentially random key K . A standard argument shows that, throughout all queries, the probability of \mathcal{B} not rejecting a type I ciphertext is smaller than $q_d \cdot (\mathbf{Adv}^{\text{CT-INT}}(\ell) + \mathbf{Adv}^{\text{KDF}}(\lambda, \ell) + q_d/p)$. It easily comes that $|\Pr[S_1] - \Pr[S_0]| \leq q_d \cdot (\mathbf{Adv}^{\text{CT-INT}}(\lambda) + \mathbf{Adv}^{\text{KDF}}(\lambda, \ell) + q_d/p)$.

We now consider type II invalid queries. While \mathcal{A} knows $t_B^{(1)}$, she has initially no information on $t_A^{(1)}$ and the last term of the product (13) is unpredictable

to her at the first type II query. Each such rejected query allows \mathcal{A} to rule out at most one candidate as for the value $t_A^{(1)}$. After $i \leq q_d$ unsuccessful type II queries, she is left with at least $p - i$ candidates at the next type II query, where the distance between the uniform distribution on \mathbb{G}_T and that of ψ (calculated as per (13)) becomes smaller than $i/p \leq q_d/p$. Again, one can show that, throughout all queries, the probability of \mathcal{B} not rejecting a type II ciphertext is at most $q_d \cdot (\mathbf{Adv}^{\text{CT-INT}}(\ell) + \mathbf{Adv}^{\text{KDF}}(\lambda, \ell) + q_d/p)$. Let us call **type-2** the latter event. If all invalid ciphertexts are rejected, \mathcal{A} 's probability of success is given by $\Pr[S_1 | \neg \text{type-2}] \leq 1/(p - q_d) \leq (q_d + 1)/p$. Since

$$\begin{aligned} \Pr[S_1] &= \Pr[S_1 \wedge \text{type-2}] + \Pr[S_1 \wedge \neg \text{type-2}] \\ &\leq \Pr[\text{type-2}] + \Pr[S_1 | \neg \text{type-2}] \Pr[\neg \text{type-2}] \\ &\leq \Pr[\text{type-2}] + \Pr[S_1 | \neg \text{type-2}] \\ &\leq q_d \cdot (\mathbf{Adv}^{\text{CT-INT}}(\ell) + \mathbf{Adv}^{\text{KDF}}(\lambda, \ell) + \frac{q_d}{p}) + \frac{q_d + 1}{p} \end{aligned}$$

and $|\Pr[S_0] - \Pr[S_1]| \leq q_d \cdot (\mathbf{Adv}^{\text{CT-INT}}(\lambda) + \mathbf{Adv}^{\text{KDF}}(\lambda, \ell) + q_d/p)$, the claimed upper bound follows. \square