Attribute-Based Encryption Implies Identity-Based Encryption

Javier Herranz

Dept. Matemàtiques
Universitat Politècnica de Catalunya
c. Jordi Girona 1-3, 08034, Barcelona, Spain

Abstract. In this short paper we formally prove that designing attribute-based encryption schemes cannot be easier than designing identity-based encryption schemes. In more detail, we show how an attribute-based encryption scheme which admits, at least, AND policies can be combined with a collision-resistant hash function to obtain an identity-based encryption scheme.
Even if this result may seem natural, not surprising at all, it has not been explicitly written anywhere, as far as we know. Furthermore, it may be an unknown result for some people: Odelu and Das [7] have proposed an attribute-based encryption scheme in the Discrete Logarithm setting, without bilinear pairings, admitting AND policies. If this scheme was secure, then by using the implication that we prove in this paper, we would obtain a secure identity-based encryption scheme in the Discrete Logarithm setting, without bilinear pairings, which would be a breakthrough in the area. Unfortunately, we present here a complete attack of the scheme proposed by Odelu and Das in [7].

1 Introduction

In a classical encryption scheme, for both the symmetric and asymmetric settings, a message is encrypted so that a single user, in possession of a secret key, can decrypt and recover the original plaintext. In the last years, other cryptographic paradigms have been proposed so that the sender of the message encrypts it in such a way that, later, many different users will be able to decrypt, as long as their identities, attributes or credentials are enough. In particular, maybe a user who is not registered in the system at the time where a message is encrypted can later decrypt it. Identity-based and attribute-based encryption are perhaps the two instantiations of these alternative paradigms that have attracted more attention from the cryptographic community. These paradigms are suitable for situations where many different kinds of users and data are in place: social networks, the Internet of Things, Cloud storage and Cloud computation, analysis of big data, etc.

A ciphertext computed by an identity-based encryption (IBE, for short) scheme for a specific identity id can be decrypted only by the user(s) holding this exact identity id. Other users, having secret keys for other identities $\mathsf{id}' \neq \mathsf{id}$, must obtain nothing useful on the plaintext. In a ciphertext-policy attribute-based encryption (ABE, for short) scheme, decryption can be performed only by users who hold a subset of attributes $A \subset \mathcal{P}$ that satisfy some policy $\Gamma \subset 2^{\mathcal{P}}$ chosen by the sender of the message. An adversary who obtains secret keys for other subsets of attributes $B_1, \ldots, B_q$ cannot obtain any information on the plaintext, if $B_i \notin \Gamma$ holds for all $i = 1, \ldots, q$. This collusion-resistance property must hold even if the union of (some of) the subsets B_i gives a subset in Γ, maybe the whole set $\mathcal{P}$ of attributes.

Both identity and attribute-based encryption are particular instantiations of more general notions that have been introduced later, like predicate encryption or functional encryption. The notion of identity-based encryption [9] was generalized to the notion of fuzzy identity-based encryption [8], which was then generalized to the notion of attribute-based encryption [3, 1], with two flavours: key-policy and ciphertext-policy. Therefore, it seems natural to believe that identity-based encryption is a particular case, an instantiation, of attribute-based encryption. However, the two aforementioned generalizations add some modifications to the initial notion of identity-based encryption, which could potentially affect this natural chain of implications. Namely, in identity-based encryption, the set of possible identities may have exponential size on the security parameter, whereas in attribute-based encryption the set of attributes has polynomial size. Therefore, the

natural instantiation of seeing each identity as an attribute, does not work. Even the solution of assigning to each identity $\mathsf{id} = (id_1, \ldots, id_\ell) \in \{0,1\}^\ell$ the set of attributes $B_{\mathsf{id}} = \{\mathsf{at}_i \mid id_i = 1\}$ does not work, because two different identities $\mathsf{id} \neq \mathsf{id}'$ can potentially lead to two subsets of attributes $B_{\mathsf{id}}, B_{\mathsf{id}'}$ such that $B_{\mathsf{id}} \subset B_{\mathsf{id}'}$, which prevents us from reducing the security of one scheme to the security of the other one, for instance if $(\mathsf{id}, B_{\mathsf{id}})$ correspond to the challenge identity and $(\mathsf{id}', B_{\mathsf{id}'})$ correspond to an extraction query.

1.1 Our Contributions

Although these first two attempts to prove that ABE implies IBE do not work, we provide a way of proving such implication. We start with an ABE scheme which admits, at least, AND policies on a set of 2ℓ attributes; by using a collision-resistance hash function $H : \{0,1\}^* \to \{0,1\}^\ell$, we derive an IBE scheme for arbitrary identities. The security of the resulting IBE scheme is the same as that of the initial ABE scheme.

We detail this construction for the case of ciphertext-policy ABE, but from the construction itself it is clear that the same can be done by starting from a key-policy ABE scheme. Therefore, the first contribution of the paper is to formally prove that any ABE scheme which admits at least AND policies can be transformed into an IBE scheme. As a direct consequence of this result, by using [2], we obtain that designing (meaningful) ABE schemes from public-key encryption or trapdoor permutations, in a black-box way, is impossible.

The result that designing (minimally useful) ABE schemes cannot be easier than designing IBE schemes may sound as a natural, folkloric, well-known result; however, we have not found any explicit formal proof of it. Furthermore, there seems to be some people in the cryptographic community who are not aware of this implication. Namely, Odelu and Das [7] have proposed a CP-ABE scheme, admitting AND policies, in the Discrete Logarithm setting, without bilinear pairings, along with a proof of security for it. With the result that we prove in this paper, a direct consequence would be the first IBE scheme in the Discrete Logarithm setting, without bilinear pairings. This would be a breakthrough result in cryptography, so we may suspect that some of the proofs (security of the ABE scheme by Odelu and Das, or security of the ABE $\Rightarrow$ IBE implication) is incorrect. Indeed, we give an explicit attack against the CP-ABE scheme proposed by Odelu and Das in [7]. In particular, the existence of such an attack means that the security analysis provided in [7] must be wrong at some point.

Therefore, the existence of IBE or ABE schemes in the Discrete Logarithm setting, without bilinear pairings, remains as an open problem. We note that relaxed versions of these notions, such as bounded-IBE and bounded-ABE, can be obtained in this setting [10, 4, 5].

1.2 Organization of the Paper

The rest of the paper is organized as follows. In Section 2 we recall the notions of identity-based encryption and ciphertext-policy attribute-based encryption schemes: we give the syntax definition and the required security properties for such schemes. We describe in Section 3 the transformation that constructs, from a ciphertext-policy attribute-based encryption scheme, an identity-based encryption scheme. We formally prove that this transformation preserves security: if the initial ABE scheme is secure, so it is the resulting IBE scheme. In Section 4 we present an explicit attack that totally breaks the attribute-based encryption scheme of Odelu and Das [7]. We conclude the paper in Section 5, with some final remarks and (hard) open problems.

2 IBE and ABE: Protocols and Security

2.1 IBE: Syntactic Definition

An identity-based encryption (IBE, from now on) scheme IBE consists of four probabilistic polynomial-time algorithms:

- $\mathsf{IBE.Setup}(1^\lambda)$. The setup algorithm takes as input a security parameter λ; it outputs some public parameters pms and a master secret key msk.

– $\mathsf{IBE.KeyGen}(\mathsf{id}, \mathsf{msk}, \mathsf{pms})$. The key generation algorithm takes as input the master secret key msk, the public parameters pms and an identity $\mathsf{id} \in \{0,1\}^*$. The output is a private key $\mathsf{sk}_{\mathsf{id}}$.
– $\mathsf{IBE.Encrypt}(m, \mathsf{id}, \mathsf{pms})$. The encryption algorithm takes as input the public parameters pms, a message m and an identity id. The output is a ciphertext C.
– $\mathsf{IBE.Decryption}(C, \mathsf{id}, \mathsf{sk}_{\mathsf{id}}, \mathsf{pms})$. The decryption algorithm takes as input a ciphertext C, an identity id, a secret key $\mathsf{sk}_{\mathsf{id}}$ and the public parameters pms. The output is a message $\tilde{m}$.

The property of correctness requires that, if the following four protocols are run: $(\mathsf{msk}, \mathsf{pms}) \leftarrow \mathsf{IBE.Setup}(1^\lambda)$, $\mathsf{sk}_i d \leftarrow \mathsf{IBE.KeyGen}(\mathsf{id}, \mathsf{msk}, \mathsf{pms})$, $C \leftarrow \mathsf{IBE.Encrypt}(m, \mathsf{id}, \mathsf{pms})$ and $\tilde{m} \leftarrow \mathsf{IBE.Decryption}(C, \mathsf{id}, \mathsf{sk}_{\mathsf{id}}, \mathsf{pms})$, then it holds $\tilde{m} = m$.

2.2 IBE: Security Definition

The usual security notion for encryption schemes is indistinguishability of ciphertexts under chosen plaintext attacks (IND-CPA security). In the setting of IBE, an attacker is also allowed to query for secret keys for identities different from the identity id^* that will be used to generate the challenge ciphertext.

To formally define the resulting security notion, we consider the following experiment involving a challenger and an adversary $\mathcal{A}_{\mathsf{IBE}}$.

1. The challenger chooses a random bit $b \xleftarrow{R} \{0,1\}$, runs $(\mathsf{pms}, \mathsf{msk}) \leftarrow \mathsf{IBE.Setup}(1^\lambda)$ and sends pms to $\mathcal{A}_{\mathsf{IBE}}$.
2. $\mathcal{A}_{\mathsf{IBE}}$ can make secret key queries for identities $\mathsf{id} \in \{0,1\}^*$ of his choice. To answer such a query, the challenger runs $\mathsf{sk}_{\mathsf{id}} \leftarrow \mathsf{IBE.KeyGen}(\mathsf{id}, \mathsf{msk}, \mathsf{pms})$ and sends $\mathsf{sk}_{\mathsf{id}}$ to $\mathcal{A}_{\mathsf{IBE}}$.
3. At some point, $\mathcal{A}_{\mathsf{IBE}}$ sends two plaintexts $m^{(0)} \neq m^{(1)}$ and a challenge identity id^* to the challenger, where $\mathsf{id}^* \neq \mathsf{id}$, for all the identities id for which a secret key has been queried in step 3.
4. The challenger runs $C^* \leftarrow \mathsf{IBE.Encrypt}(m^{(b)}, \mathsf{id}^*, \mathsf{pms})$ and sends the challenge ciphertext C^* to $\mathcal{A}_{\mathsf{IBE}}$.
5. $\mathcal{A}_{\mathsf{IBE}}$ can make more secret key queries for more identities id, as long as $\mathsf{id} \neq \mathsf{id}^*$.
6. Finally, $\mathcal{A}_{\mathsf{IBE}}$ outputs a bit $b' \in \{0,1\}$.

The advantage of $\mathcal{A}_{\mathsf{IBE}}$ in breaking the IND-CPA property of the IBE scheme is defined as

$$\mathsf{Adv}^{\mathsf{ind\text{-}cpa}}_{\mathcal{A}_{\mathsf{IBE}}}(\lambda) = \left| \Pr[b' = b] - \frac{1}{2} \right|.$$

Definition 1. *An identity-based encryption scheme is indistinguishable under adaptive chosen-plaintext attacks (IND-CPA secure) if, for any adversary $\mathcal{A}_{IBE}$ that runs in polynomial time, the advantage $\mathsf{Adv}^{\mathsf{ind\text{-}cpa}}_{\mathcal{A}_{IBE}}(\lambda)$ is negligible in the security parameter λ.*

We recall that a function $f(\lambda)$ is said to be negligible if it decreases (as λ increases) faster than the inverse of any polynomial.

A weaker security notion for IBE schemes is selective IND-CPA security, which is defined by a very similar game. The difference is that the attacker must choose the identity id^* at the very beginning, before step 1 of the experiment.

2.3 CP-ABE: Syntactic Definition

A ciphertext-policy attribute-based encryption (CP-ABE, from now on) scheme ABE consists of four probabilistic polynomial-time algorithms:

– $\mathsf{ABE.Setup}(1^\lambda, \mathcal{U}, \mathcal{F})$. The setup algorithm takes as input a security parameter λ, the total universe of attributes $\mathcal{U} = \{\mathsf{at}_1, \ldots, \mathsf{at}_n\}$ and the family $\mathcal{F}$ of decryption policies that the scheme supports. It outputs some public parameters pms and a master secret key msk.

– $\mathsf{ABE.KeyGen}(A, \mathsf{msk}, \mathsf{pms})$. The key generation algorithm takes as input the master secret key msk, the public parameters pms and a set of attributes $A \subset \mathcal{U}$ satisfied by the user. The output is a private key sk_A.
– $\mathsf{ABE.Encrypt}(m, \mathcal{P}, \Gamma, \mathsf{pms})$. The encryption algorithm takes as input the public parameters pms, a message m and a decryption policy $(\mathcal{P}, \Gamma)$ where $\mathcal{P} \subset \mathcal{U}$ and $\Gamma \subset 2^{\mathcal{P}}$ satisfies $\Gamma \in \mathcal{F}$. The output is a ciphertext C.
– $\mathsf{ABE.Decryption}(C, \mathcal{P}, \Gamma, \mathsf{sk}_S, \mathsf{pms})$. The decryption algorithm takes as input a ciphertext C, a decryption policy $(\mathcal{P}, \Gamma)$, a secret key sk_A and the public parameters pms. The output is a message $\tilde{m}$.

The property of correctness requires that, if the following four protocols are run: $(\mathsf{msk}, \mathsf{pms}) \leftarrow \mathsf{ABE.Setup}(1^\lambda, \mathcal{U}, \mathcal{F})$, $\mathsf{sk}_A \leftarrow \mathsf{ABE.KeyGen}(A, \mathsf{msk}, \mathsf{pms})$, $C \leftarrow \mathsf{ABE.Encrypt}(m, \mathcal{P}, \Gamma, \mathsf{pms})$ and $\tilde{m} \leftarrow \mathsf{ABE.Decryption}(C, \mathcal{P}, \Gamma, \mathsf{sk}_S, \mathsf{pms})$, then it holds $\tilde{m} = m$, if $A \cap \mathcal{P} \in \Gamma$ and $\Gamma \in \mathcal{F}$.

Regarding the family $\mathcal{F}$ of admitted decryption policies, $\mathcal{F}$ may for instance contain all the possible policies, $\mathcal{F} = \{\Gamma \subset 2^{\mathcal{U}}\}$, or may contain all the monotone increasing policies, $\mathcal{F} = \{\Gamma \subset 2^{\mathcal{U}}, \Gamma \text{ is monotone increasing}\}$, where Γ is monotone increasing if $A \in \Gamma, A \subset B$ implies $B \in \Gamma$. Some schemes may support only threshold decryption policies, $\mathcal{F} = \{\Gamma_{(t,\mathcal{P})}, \mathcal{P} \subset \mathcal{U}, 1 \leq t \leq |\mathcal{P}|\}$, and $\Gamma_{(t,\mathcal{P})} = \{A \subset \mathcal{P} \text{ s.t. } |A| \geq t\}$. A particular, more restrictive, case of threshold policies corresponds to AND policies, of the form $\Gamma_{(|\mathcal{P}|,\mathcal{P})} = \{\mathcal{P}\}$, containing only one subset, $\mathcal{P} \subset \mathcal{U}$. Since the CP-ABE schemes studied in this paper support AND policies, we will refer to these policies as $\mathcal{F}_{\mathsf{AND}} = \{\Gamma_{(|\mathcal{P}|,\mathcal{P})} \text{ s.t. } \mathcal{P} \subset \mathcal{U}\}$.

2.4 CP-ABE: Security Definition

In the setting of CP-ABE, an attacker against the IND-CPA security of the scheme is allowed to query for secret keys for different subsets of users, as long as none of these subsets is authorized for the decryption policy $(\mathcal{P}, \Gamma)$ which will be used to generate the challenge ciphertext.

To formally define the resulting security notion, we consider the following experiment involving a challenger and an adversary $\mathcal{A}_{\mathsf{ABE}}$.

1. The adversary $\mathcal{A}_{\mathsf{ABE}}$ chooses the universe of attributes $\mathcal{U}$ and the family of policies $\mathcal{F}$.
2. The challenger chooses a random bit $b \xleftarrow{R} \{0,1\}$, runs $(\mathsf{pms}, \mathsf{msk}) \leftarrow \mathsf{ABE.Setup}(1^\lambda, \mathcal{U}, \mathcal{F})$ and sends pms to $\mathcal{A}_{\mathsf{ABE}}$.
3. $\mathcal{A}_{\mathsf{ABE}}$ can make secret key queries for subsets of attributes $A_i \subset \mathcal{U}$ of his choice. To answer such a query, the challenger runs $\mathsf{sk}_{A_i} \leftarrow \mathsf{ABE.KeyGen}(A_i, \mathsf{msk}, \mathsf{pms})$ and sends sk_{A_i} to $\mathcal{A}_{\mathsf{ABE}}$.
4. At some point, $\mathcal{A}_{\mathsf{ABE}}$ sends two plaintexts $m^{(0)} \neq m^{(1)}$ and a decryption policy $(\mathcal{P}^*, \Gamma^*)$ to the challenger, where $\mathcal{P}^* \subset \mathcal{U}$, $\Gamma^* \in \mathcal{F}$ and $A_i \cap \mathcal{P}^* \notin \Gamma^*$, for all the subsets A_i for which a secret key has been queried in step 3.
5. The challenger runs $C^* \leftarrow \mathsf{ABE.Encrypt}(m^{(b)}, \mathcal{P}^*, \Gamma^*, \mathsf{pms})$ and sends the challenge ciphertext C^* to $\mathcal{A}_{\mathsf{ABE}}$.
6. $\mathcal{A}_{\mathsf{ABE}}$ can make more secret key queries for subsets of attributes A_i, as long as $A_i \cap \mathcal{P}^* \notin \Gamma^*$.
7. Finally, $\mathcal{A}_{\mathsf{ABE}}$ outputs a bit $b' \in \{0,1\}$.

The advantage of $\mathcal{A}_{ABE}$ in breaking the IND-CPA property of the CP-ABE scheme is defined as

$$\mathsf{Adv}^{\mathsf{ind\text{-}cpa}}_{\mathcal{A}_{\mathsf{ABE}}}(\lambda) = \left| \Pr[b' = b] - \frac{1}{2} \right|.$$

Definition 2. *A ciphertext-policy attribute-based encryption scheme is indistinguishable under adaptive chosen-plaintext attacks (IND-CPA secure) if, for any adversary $\mathcal{A}_{ABE}$ that runs in polynomial time, the advantage $\mathsf{Adv}^{\mathsf{ind\text{-}cpa}}_{\mathcal{A}_{ABE}}(\lambda)$ is negligible in the security parameter λ.*

A weaker security notion for CP-ABE schemes is selective IND-CPA security, which is defined by a very similar game. The difference is that the attacker must choose the policy $(\mathcal{P}, \Gamma)$ at the very beginning, in step 1 of the experiment.

3 CP-ABE Implies IBE

The main result of this paper is that ABE implies IBE. This implication holds for the two existing flavours of ABE, ciphertext-policy and key-policy. We will detail the result (construction and security proof) for the case of ciphertext-policy ABE, but since the policies involved in the construction are AND policies, which can be thought as a "having all the attributes from a specific subset", it is clear that the same result holds if we start from a key-policy ABE, by swapping the roles of secret keys and ciphertexts.

3.1 The Transformation

Let $\mathsf{ABE} = (\mathsf{ABE.Setup}, \mathsf{ABE.KeyGen}, \mathsf{ABE.Encrypt}, \mathsf{ABE.Decrypt})$ be a CP-ABE scheme admitting, at least, AND policies. Let $H : \{0,1\}^* \to \{0,1\}^\ell$ be a hash function.

An IBE scheme $\mathsf{IBE} = (\mathsf{IBE.Setup}, \mathsf{IBE.KeyGen}, \mathsf{IBE.Encrypt}, \mathsf{IBE.Decrypt})$ is constructed as follows.

- $\mathsf{IBE.Setup}$: run $(\mathsf{pms}_{\mathsf{ABE}}, \mathsf{msk}_{\mathsf{ABE}}) \leftarrow \mathsf{ABE.Setup}(1^\lambda, \mathcal{U}, \mathcal{F})$, for $\mathcal{U} = \{\mathsf{at}_{1,0}, \mathsf{at}_{1,1}, \ldots, \mathsf{at}_{\ell,0}, \mathsf{at}_{\ell,1}\}$ ($n = 2\ell$ attributes), and $\mathcal{F} \supset \mathcal{F}_{\mathsf{AND}} = \{\Gamma_{(|\mathcal{P}|,\mathcal{P})} \text{ s.t. } \mathcal{P} \subset \mathcal{U}\}$. The master secret key of IBE is $\mathsf{msk}_{\mathsf{IBE}} = \mathsf{msk}_{\mathsf{ABE}}$, the public parameters of IBE are $\mathsf{pms}_{\mathsf{IBE}} = (\mathsf{pms}_{\mathsf{ABE}}, H)$.
- $\mathsf{IBE.KeyGen}$: given an identity $\mathsf{id} \in \{0,1\}^*$, let $H(\mathsf{id})_i$ denote the i-th bit of $H(\mathsf{id})$, for $i = 1, \ldots, \ell$. Define the subset $A_{\mathsf{id}} \subset \mathcal{U}$, which contains ℓ attributes, as

$$A_{\mathsf{id}} = \{\mathsf{at}_{i,H(\mathsf{id})_i}, 1 \leq i \leq \ell\}.$$

  The secret key for identity id is defined as $\mathsf{sk}_{\mathsf{id}} \leftarrow \mathsf{ABE.KeyGen}(\mathsf{pms}_{\mathsf{ABE}}, A_{\mathsf{id}}, \mathsf{msk}_{\mathsf{ABE}})$, the ABE secret key for the subset of attributes A_{id}.
- $\mathsf{IBE.Encrypt}$: given a plaintext m and an identity id, consider the AND policy $\Gamma_{(|\mathcal{P}|,\mathcal{P})}$ for the subset $\mathcal{P} = A_{\mathsf{id}}$. The ciphertext for this pair (m, id) is defined as $C \leftarrow \mathsf{ABE.Encrypt}(m, \mathcal{P}, \Gamma_{(|\mathcal{P}|,\mathcal{P})}, \mathsf{pms}_{\mathsf{ABE}})$.
- $\mathsf{IBE.Decrypt}$: given a ciphertext C for an identity id and the secret key $\mathsf{sk}_{\mathsf{id}}$ for that identity, the decryption protocol outputs the plaintext obtained by running $m' \leftarrow \mathsf{ABE.Decrypt}(C, \mathcal{P}, \Gamma_{(|\mathcal{P}|,\mathcal{P})}, \mathsf{sk}_{\mathsf{id}}, \mathsf{pms}_{\mathsf{ABE}})$.

3.2 Security of the Transformation

We are going to prove that, if the scheme ABE is secure, then the scheme IBE is secure, too. We are going to consider adaptive (full) security for both schemes. The analogous result with selective security for both the CP-ABE and IBE schemes may be proved in a very similar way.

Theorem 1. *If ABE is IND-CPA secure and H is a collision-resistant hash function, then IBE is IND-CPA secure.*

Proof. To prove this result, we assume the existence of a successful adversary $\mathcal{A}_{\mathsf{IBE}}$ against the IND-CPA security of the scheme IBE and we design an adversary $\mathcal{A}_{\mathsf{ABE}}$ against the IND-CPA security of the CPA-ABE scheme ABE.

In the IND-CPA experiment for $\mathcal{A}_{ABE}$, a bit $b \xleftarrow{R} \{0,1\}$ is first chosen at random. Then the adversary $\mathcal{A}_{ABE}$ that we are designing asks for the execution of $\mathsf{ABE.Setup}$, for a security parameter λ, a universe $\mathcal{U} = \{\mathsf{at}_{1,0}, \mathsf{at}_{1,1}, \ldots, \mathsf{at}_{\ell,0}, \mathsf{at}_{\ell,1}\}$ of $n = 2\ell$ attributes and for a family $\mathcal{F}$ of decryption policies which contains AND policies, $\mathcal{F} \supset \mathcal{F}_{\mathsf{AND}} = \{\Gamma_{(|\mathcal{P}|,\mathcal{P})} \text{ s.t. } \mathcal{P} \subset \mathcal{U}\}$. As a result of this execution of $\mathsf{ABE.Setup}$, the adversary $\mathcal{A}_{\mathsf{ABE}}$ receives some public parameters $\mathsf{pms}_{\mathsf{ABE}}$. At this point, $\mathcal{A}_{ABE}$ initializes an execution of the adversary $\mathcal{A}_{\mathsf{IBE}}$, by providing him with his initial input, the public parameeters of the identity-based scheme IBE, which are set to be $\mathsf{pms}_{\mathsf{IBE}} = (\mathsf{pms}_{\mathsf{ABE}}, H)$.

The adversary $\mathcal{A}_{\mathsf{IBE}}$ can, from this point on, make secret key queries for identities id of his choice. To answer such queries, our adversary $\mathcal{A}_{\mathsf{ABE}}$ proceeds as follows.

- Let $H(\mathsf{id})_i$ denote the i-th bit of $H(\mathsf{id})$. Define the subset $A_{\mathsf{id}} \subset \mathcal{U}$ as $A_{\mathsf{id}} = \{\mathsf{at}_{i,H(\mathsf{id})_i}, 1 \leq i \leq \ell\}$.

– Making use of the secret key queries that $\mathcal{A}_{\mathsf{ABE}}$ can make, he asks for a secret key for the subset of attributes A_{id}. As the answer, $\mathcal{A}_{\mathsf{ABE}}$ gets $\mathsf{sk}_{A_{\mathsf{id}}} \leftarrow \mathsf{ABE.KeyGen}(\mathsf{pms}_{\mathsf{ABE}}, A_{\mathsf{id}}, \mathsf{msk}_{\mathsf{ABE}})$.
– Send to $\mathcal{A}_{\mathsf{IBE}}$ the secret key $\mathsf{sk}_{\mathsf{id}} = \mathsf{sk}_{A_{\mathsf{id}}}$.

At some point $\mathcal{A}_{\mathsf{IBE}}$ outputs two different plaintexts, $m^{(0)}, m^{(1)}$ and a challenge identity id^* such that $\mathsf{id}^* \neq \mathsf{id}$, for all the identities id for which $\mathcal{A}_{\mathsf{IBE}}$ made a secret key query.

Since H is assumed to be a collision resistance hash function, we have $H(\mathsf{id}^*) \neq H(\mathsf{id})$, for all queried identities m. Therefore, the subset of attributes $\mathcal{P}^* = A_{\mathsf{id}^*} = \{\mathsf{at}_{i,H(\mathsf{id}^*)_i}, 1 \leq i \leq \ell\}$ satisfies that $\mathcal{P}^* \not\subset A_{\mathsf{id}}$, for all queried identities id. Indeed, for each queried identity id, let $j \in \{1, \ldots, \ell\}$ be such that $H(\mathsf{id}^*)_j \neq H(\mathsf{id})_j$. We have $\mathsf{at}_{j,H(\mathsf{id}^*)_j} \in \mathcal{P}^*$ and $\mathsf{at}_{j,H(\mathsf{id}^*)_j} \notin A_{\mathsf{id}}$, so $\mathcal{P}^* \not\subset A_{\mathsf{id}}$. Therefore, there cannot be any inclusion relation between A_{id} and $\mathcal{P}^*$.

This means that, for all queried identities id, we have that the subset A_{id} does not satisfy the AND policy $\Gamma_{(|\mathcal{P}^*|,\mathcal{P}^*)}$ and thus, $\mathcal{A}_{\mathsf{ABE}}$ can choose $\Gamma_{(|\mathcal{P}|,\mathcal{P})}$ as the challenge policy. $\mathcal{A}_{\mathsf{ABE}}$ chooses the same two messages $m^{(0)}, m^{(1)}$. $\mathcal{A}_{\mathsf{ABE}}$ sends these two messages, along with the subset of attributes $\mathcal{P}^*$ and the policy $\Gamma^* = \Gamma_{(|\mathcal{P}|,\mathcal{P})}$, and gets as answer a challenge ciphertext $C^* \leftarrow \mathsf{ABE.Encrypt}(m^{(b)}, \mathcal{P}^*, \Gamma_{(|\mathcal{P}|,\mathcal{P})}, \mathsf{pms}_{\mathsf{ABE}})$.

$\mathcal{A}_{\mathsf{ABE}}$ gives to $\mathcal{A}_{\mathsf{IBE}}$ the challenge ciphertext C^*. If $\mathcal{A}_{\mathsf{IBE}}$ makes more secret key queries, for $\mathsf{id} \neq \mathsf{id}^*$, they are answered as the previous ones, and the same argument that the subsets of attributes A_{id} do not satisfy the AND policy Γ^* is valid.

Finally, when $\mathcal{A}_{\mathsf{IBE}}$ outputs a bit b', our adversary $\mathcal{A}_{\mathsf{ABE}}$ outputs the same bit b'.

Obviously, we have that $\mathsf{Adv}^{\mathsf{ind\text{-}cpa}}_{\mathcal{A}_{\mathsf{ABE}}}(\lambda) = \mathsf{Adv}^{\mathsf{ind\text{-}cpa}}_{\mathcal{A}_{\mathsf{IBE}}}(\lambda)$, which concludes the proof. □

4 An Attack against the CP-ABE Scheme in [7]

In [7], Odelu and Das propose a CP-ABE scheme which supports AND decryption policies and which works in the classical, pairing-free, Discrete Logarithm setting. They prove that the scheme enjoys selective IND-CPA security. In this section we show that their security analysis must be incorrect at some point, because we provide an explicit attack against the scheme.

Since the attack we are going to present is a key-recovery attack, stronger than breaking IND-CPA security, we will describe only (the necessary parts of) the Setup and Key Generation protocols of their CP-ABE scheme, which will be enough to, later, understand the proposed attack. The details of the Encrypt and Decrypt protocols are thus not necessary; we simply assume that they satisfy the standard correctness property.

4.1 Description of the CP-ABE Scheme in [7]

The typical Discrete Logarithm framework consists of a cyclic group $\mathbb{G}$ of prime order p. Examples of such groups are some groups of points in elliptic curves or subgroups of $\mathbb{Z}_q$, when $p|q-1$. We will use additive notation in $\mathbb{G}$, to follow the same notation as in [7]. That is, $\mathbb{G} = \{aP,\ a \in \{0, 1, \ldots, p-1\}\}$, where P is a generator of $\mathbb{G}$.

Their scheme supports AND decryption policies $\Gamma_{(|\mathcal{P}|,\mathcal{P})}$, defined on the total universe $\mathcal{U} = \{\mathsf{at}_1, \ldots, \mathsf{at}_n\}$ of attributes. They use the following notation: any subset $A \subset \mathcal{U}$ will be represented by an n-bit string $a_1 a_2 \ldots a_n$, where $a_i = 1$ if $\mathsf{at}_i \in A$, and $a_i = 0$ if $\mathsf{at}_i \notin A$. For example, if $n = 4$ and $A = \{\mathsf{at}_1, \mathsf{at}_4\}$, then the bit string corresponding to A is 1001.

With this notation, an AND decryption policy $\Gamma_{(|\mathcal{P}|,\mathcal{P})}$ may be represented by the bit string $b_1 b_2 \ldots b_n$ corresponding to subset $\mathcal{P}$. If A is a subset of attributes (held by a user), with bit string $a_1 a_2 \ldots a_n$, the condition that must be satisfied in order for that user to decrypt, $A \cap \mathcal{P} \in \Gamma_{(|\mathcal{P}|,\mathcal{P})}$, which is equivalent to $\mathcal{P} \subset A$, becomes $a_i \geq b_i$, $\forall i = 1, \ldots, n$.

We are now ready to describe the Setup and Key Generation protocols of the scheme in [7].

$\mathsf{Setup}(1^\lambda, \mathcal{U}, \mathcal{F}_{\mathsf{AND}})$. The setup algorithm starts by choosing a cyclic group $\mathbb{G}$ of prime order p, such that p is λ bits long, and a generator P of $\mathbb{G}$. A hash function $H_4 : \{0,1\}^* \to \mathbb{Z}_p^*$ is also chosen.

Then, three random elements $\alpha, k_1, k_2 \in \mathbb{Z}_p$ are chosen. For each $i \in \{0, 1, \ldots, n\}$, the values $P_i = \alpha^i P$, $U_i = k_1 \alpha^i P$ and $V_i = k_2 \alpha^i P$ are computed.

The master secret key is $\mathsf{msk} = (\alpha, k_1, k_2)$.

The public parameters of the system are $\mathsf{pms} = (p, \mathbb{G}, P, H_4, \{\mathcal{P}_i, U_i, V_i\}_{0 \le i \le n})$.

$\mathsf{KeyGen}(A, \mathsf{msk}, \mathsf{pms})$. The key generation algorithm takes as input a subset of attributes $A \subset \mathcal{U}$, the master secret key msk and the public parameters pms.

Let $a_1 a_2 \ldots a_n$ be the bit string corresponding to subset A. Let us define the polynomial $f(x, A) = \prod_{i=1}^{n} (x + H_4(i))^{1-a_i}$, which has degree $n - |A|$.

Two random numbers $r_u, t_u \in \mathbb{Z}_p$ are chosen. A value s_u is computed, such that the relation $\frac{1}{f(\alpha, A)} = k_1 s_u + k_2 r_u \bmod p$ is satisfied. That is:

$$s_u = \frac{1}{k_1} \cdot \left(\frac{1}{f(\alpha, A)} - k_2 r_u \right) \bmod p$$

Finally, the values $u_1 = r_u + k_1 t_u \bmod p$ and $u_2 = s_u - k_2 t_u \bmod p$ are computed and the secret key is set to be $\mathsf{sk}_A = (u_1, u_2)$.

Remark. The vulnerability of the scheme comes from the fact that a secret key $\mathsf{sk}_A = (u_1, u_2)$ does not have enough entropy. Although two random and independent values, r_u, t_u, are generated, the final two elements u_1 and u_2 are not independent. If we write in matrix notation the relation between the pairs (r_u, t_u) and (u_1, u_2), we have

$$\begin{pmatrix} u_1 \\ u_2 \end{pmatrix} = \begin{pmatrix} 1 & k_1 \\ -\frac{k_2}{k_1} & -k_2 \end{pmatrix} \cdot \begin{pmatrix} r_u \\ t_u \end{pmatrix} + \begin{pmatrix} 0 \\ \frac{1}{k_1 f(\alpha, A)} \end{pmatrix} \bmod p$$

The matrix is not invertible: the second row is equal to the first one multiplied with $-\frac{k_2}{k_1}$. Therefore, we have that

$$u_2 = -\frac{k_2}{k_1} u_1 + \frac{1}{k_1 f(\alpha, A)} \bmod p.$$

4.2 The Attack

The attack is based on three simple observations.

(1) From two secret queries for the subset of attributes A, it is easy to recover the values $X := -\frac{k_2}{k_1} \bmod p$ and $Y_A := \frac{1}{k_1 f(\alpha, A)} \bmod p$.

Indeed, according to the remark at the end of previous section, from the first secret key query we will get a pair (u_1, u_2) such that

$$u_2 = -\frac{k_2}{k_1} u_1 + \frac{1}{k_1 f(\alpha, A)} \bmod p.$$

From the second secret key query for subset A, we will get a pair (u_1', u_2') such that

$$u_2' = -\frac{k_2}{k_1} u_1' + \frac{1}{k_1 f(\alpha, A)} \bmod p.$$

We can consider the following system of equations, in matrix notation (and using the notation $X := -\frac{k_2}{k_1} \bmod p$ and $Y_A := \frac{1}{k_1 f(\alpha, A)} \bmod p$ for the unknowns):

$$\begin{pmatrix} u_1 & 1 \\ u_1' & 1 \end{pmatrix} \cdot \begin{pmatrix} X \\ Y_A \end{pmatrix} = \begin{pmatrix} u_2 \\ u_2' \end{pmatrix} \bmod p.$$

Since $u_1 \neq u_2' \bmod p$ with overwhelming probability, the matrix is invertible and we can recover (X, Y_A) from the two secret key queries.

(2) For each subset of attributes B, knowledge of the pair (X, Y_B) is enough to produce a valid secret key sk_B for subset B.

Again, from the remark at the end of previous section, the only thing we have to do is to choose $u_1 \in \mathbb{Z}_p$ at random and compute

$$u_2 = -\frac{k_2}{k_1} u_1 + \frac{1}{k_1 f(\alpha, B)} = X u_1 + Y_B \bmod p$$

The resulting key $\mathsf{sk}_B = (u_1, u_2)$ has the same probability distribution as in a real execution of $\mathsf{KeyGen}(B, \mathsf{msk}, \mathsf{pms})$.

(3) There are some basic algebraic relations between the values $f(\alpha, A)$ (and thus, between the values Y_A) for different subsets A of attributes.

For instance, let us take $n = 3$ and the subsets of attributes defined by bit strings $A_1 = 001$, $A_2 = 110$, $A_3 = 010$, $B = 101$. It is easy to check that the following equality holds

$$f(\alpha, B) = \frac{f(\alpha, A_1) \cdot f(\alpha, A_2)}{f(\alpha, A_3)} \bmod p.$$

Now, for these subsets of attributes, we have

$$Y_B = \frac{1}{k_1 f(\alpha, B)} = \frac{1}{k_1 \frac{f(\alpha, A_1) \cdot f(\alpha, A_2)}{f(\alpha, A_3)}} = \frac{\frac{1}{k_1 f(\alpha, A_1)} \cdot \frac{1}{k_1 f(\alpha, A_2)}}{\frac{1}{k_1 f(\alpha, A_3)}} = \frac{Y_{A_1} \cdot Y_{A_2}}{Y_{A_3}} \bmod p.$$

We are now ready to explain the attack, for this particular set of four subsets $A_1 = 001$, $A_2 = 110$, $A_3 = 010$, $B = 101$ in a universe with $n = 3$ attributes. We are designing a selective attack, so we can choose the policy for the challenge ciphertext in advance, as $\Gamma_{(|\mathcal{P}|, \mathcal{P})}$ for $\mathcal{P} = B = 101$.

1. Use **(1)**: make two secret queries for each of the subsets of attributes A_1, A_2, A_3. Since none of these attributes satisfy policy $\Gamma_{(|\mathcal{P}|, \mathcal{P})}$, these are all valid secret key queries. As a result, obtain the values $X, Y_{A_1}, Y_{A_2}, Y_{A_3}$.
2. Use **(3)**: compute $Y_B = \frac{Y_{A_1} \cdot Y_{A_2}}{Y_{A_3}} \bmod p$.
3. Use **(2)**: knowing (X, Y_B), compute a valid secret key sk_B for subset B.
4. Knowing sk_B, it is trivial to decrypt the challenge ciphertext and win the IND-CPA experiment.

Actually, this is a key-recovery attack, even stronger than an attack against the IND-CPA property. In any case, the conclusion is that the CP-ABE scheme in [7] is not secure.

5 Final Remarks and Conclusions

It is well-known that securely designing identity-based and attribute-based encryption schemes is a hard task. There are some black-box impossibility results, for instant proving that identity-based schemes cannot be constructed from public-key encryption or from trapdoor permutations [2]. These results were extended in [6] to some other classes of predicate encryption. The implication ABE $\Rightarrow$ IBE that we formally prove in this paper means that the same impossibility results are valid for any attribute-based encryption scheme wich admits at least AND policies.

Designing a secure IBE or ABE scheme in the classical Discrete Logarithm setting, without bilinear pairings, would not contradict these impossibility results; however, it looks like a really hard problem (the scheme by Odelu and Das in [7] is not secure, as we have proved in the previous section) and maybe some similar impossibility black-box results could be obtained in this sense. In the meanwhile, the best ABE or IBE schemes that can be designed in that setting are relaxations of the original notions, for instance bounded-collision and/or symmetric IBE and ABE [10, 4, 5].

References

1. J. Bethencourt, A. Sahai and B. Waters. Ciphertext-policy attribute-based encryption. *Proc. of IEEE Symposium on Security and Privacy*, IEEE Society Press, pp. 321–334 (2007)
2. D. Boneh, P.A. Papakonstantinou, C. Rackoff, Y. Vahlis and B. Waters. On the impossibility of basing identity based encryption on trapdoor permutations. *Proc. of FOCS'08*, ACM Press, pp. 283–292 (2008)
3. V. Goyal, O. Pandey, A. Sahai and B. Waters. Attribute-based encryption for fine-grained access control of encrypted data. *Proc. of Computer and Communications Security, CCS'06*, ACM Press, pp. 89–98 (2006)
4. J. Herranz. Attribute-based versions of Schnorr and ElGamal. *Applicable Algebra in Engineering, Communication and Computing*, Vol. **27 (1)**, pp. 17–57 (2016)
5. G. Itkis, E. Shen, M. Varia, D.A. Wilson and A. Yerukhimovich. Bounded-collusion attribute-based encryption from minimal assumptions. *Proc. of PKC'17*, accepted for publication (2017). Available at: `http://eprint.iacr.org/2017/029`
6. J. Katz and A. Yerukhimovich. On black-box constructions of predicate encryption from trapdoor permutations. *Proc. of Asiacrypt'09*, LNCS **5912**, Springer-Verlag, pp. 197–213 (2009)
7. V. Odelu and A.K. Das. Design of a new CP-ABE with constant-size secret keys for lightweight devices using elliptic curve cryptography. *Security and Communication Networks*, accepted for publication (2016). See also: `http://eprint.iacr.org/2015/841`
8. A. Sahai and B. Waters. Fuzzy identity-based encryption. *Proc. of Eurocrypt'05*, LNCS **3494**, Springer-Verlag, pp. 457–473 (2005)
9. A. Shamir. Identity-based cryptosystems and signature schemes. *Proc. of Crypto'84*, LNCS **196**, Springer-Verlag, pp. 47–53 (1984)
10. S. Tessaro and D.A. Wilson. Bounded-collusion identity-based encryption from semantically-secure public-key encryption: Generic constructions with short ciphertexts. *Proc. of PKC'14*, LNCS **8383**, Springer-Verlag, pp. 257–274 (2014)