

Flexible and Secure Communications in an Identity-Based Coalition Environment

Kent D. Boklan*, Zev Klagsbrun†, Kenneth G. Paterson‡ and Sriramkrishnan Srinivasan‡

*Queens College, City University of New York, USA. Email: boklan@cs.qc.cuny.edu

†University of California - Irvine, USA. Email: zev.klagsbrun@gmail.com

‡Royal Holloway, University of London, Egham, Surrey, UK. Email: {kenny.paterson,s.srinivasan}@rhul.ac.uk

Abstract—In this paper we consider the problem of how entities operating under distinct roots of trust in a coalition environment can flexibly and securely communicate with one another. We consider the identity-based setting, with each entity being pre-configured with a private key from a particular Trusted Authority (TA), but where multiple, independent TAs are involved in the coalition. Our solution to the problem adapts the Boneh-Franklin identity-based encryption (IBE) scheme. It allows any entity to securely communicate with any other entity, even without knowing the TA with which the intended recipient is associated. To enable this, we assume that the TAs co-operate to distribute certain additional public information to all entities which allows entities to decrypt a ciphertext that was composed using the public parameters of one TA, using a private key issued by another. We include a security analysis of our new approach.

Keywords: Identity-based cryptography, MANET, multi-TA IBE

I. INTRODUCTION

Complex environments involving cooperation between entities formed under distinct roots of trust are very relevant in mobile ad hoc networks (MANETs). We consider the fundamental and important question of how to enable secure communications between disparate entities within such heterogeneous, potentially resource-constrained environments.

Traditional public key cryptography is not well-suited for such networks since it utilises large amounts of energy and bandwidth and requires a constant connection to a public key infrastructure (PKI) to look up public keys, certificates and revocation data. In addition, transmitting, storing and verifying certificates puts extra strain on already limited resources. Identity-based encryption (IBE) is an attractive choice for such resource-constrained environments since it eliminates the need for public key lookups, does not need certificates, and allows revocation to be simplified by using time-based identifiers.

Although the concept of IBE was introduced by Shamir [16] in 1984, it was not until the breakthrough papers of Sakai *et al.* [15] and Boneh and Franklin [5] that practical and secure constructions were demonstrated, using the mathematics of pairings. The paper by Boneh and Franklin also proposed the first security models for IBE and gave schemes secure in the random oracle model [4]. This led to an explosion of interest in the field of identity-based cryptography (IBC). In IBC, identifiers such as email addresses, IP addresses or any unique bit string can be used to form public keys for users.

A Trusted Authority (TA), who possesses a master secret, generates private keys for users as a function of the master secret and the user's identifier. If a party A wishes to encrypt a message for B , he only needs to know B 's identifier and the system-wide public parameters; he need not have any relation with the TA to send the message, and need not authenticate B 's public key. A number of authors have studied the applicability of IBE to MANETs [11], [6], [12], [9], [2].

Almost all prior work in this area has been limited to the case of a single TA. In more complex settings, such as those encountered in dynamic coalition forming, we may have multiple TAs (from different administrative domains) and hence multiple roots of trust. For example, in the application scenario envisioned in [14], coalition forces controlled by one coalition member need the ability to communicate securely with individuals and entities associated with other members of the coalition. In this type of scenario, it is necessary to find methods that allow entities that are under different roots of trust to securely communicate with each other. A particular challenge then arises when coalitions are formed *dynamically*. In this situation pre-configuration of devices with all of the required static security data (such as system-wide public parameters for each of the coalition TAs) before a mission commences will not, in general, be possible: the exact set of coalition members may be unpredictable in advance, and their security data may need to be updated during the course of a mission. Issuing a fresh set of system-wide public parameters for each new coalition, along with private keys for all of the coalition members, is an unattractive solution since it involves high communication costs and needs secure channels for distribution of the new private keys. This is particularly true when coalitions are short-lived, forming and re-forming rapidly. Another approach would be for coalition members to distribute the required security data amongst themselves as and when necessary, in an *ad hoc* fashion. However, this would require each member to maintain a complete set of data, would involve significant storage overhead, and could result in problems if secure communication was urgent but the required security data had not yet been received.

In this paper we present an alternative solution to the problem of enabling secure communications in dynamic coalitions. Our solution enables any entity A to securely communicate with any other coalition entity B , even if B is associated with a different TA and A is unable to obtain authentic public parameters for B 's TA (for example, if that TA is

simply unknown). The cost of our solution is that when a coalition forms the TAs must broadcast a small amount of additional information to each coalition member. However, this broadcast need not itself be encrypted. Now A can use its own TA's public parameters (along with B 's identifier) to perform the encryption to B . In fact, A can use a set of authentic public parameters from any coalition TA. Meanwhile, B , upon receipt of the broadcast information from the TAs, is able to perform key translations: B can convert its existing private key issued by its TA into private keys that are valid for the same identifier under each and every one of the other TAs in the coalition. This means that B is able to decrypt A 's message no matter which public parameters were selected by A during encryption. On the other hand, no other entity is able to translate its private key to enable decryption of the message intended for B .

We provide security models that are appropriate to this kind of multi-TA application scenario, and give a specific instantiation of our solution that adapts the Boneh-Franklin `BasicIdent` IBE scheme [5]. This results in a highly efficient encryption scheme in which the size of the coalition-enabling broadcast is linear in the number of coalition partners.

In comparison to the approach based on distributing a fresh set of system-wide public parameters for each new coalition, our approach eliminates the need for the secure channel to distribute new private keys, as well as the need for bespoke communication between TAs and individual nodes. Compared with the *ad hoc* approach, our method may be more reliable (since nodes in receipt of the single broadcast will immediately have all the information needed to enable secure communications). In the specific instantiation we provide in this paper, our approach also requires less communication. In general, our approach does make use of a network-wide authentic broadcast, and assumes that entities do know one another's identifiers even if the relevant TA public parameters are not available. It also requires the various TAs to share some common cryptographic parameters (but not master secrets), which reduces its flexibility.

A. Related Work

To the best of our knowledge, multi-TA IBE schemes where the TAs collaborate to enable secure communications in dynamic coalitions have not been investigated in the literature prior to this work. All of the work done in IBE has been considered in the single-TA setting. HIBE schemes [10], [8] do employ multiple TAs but our work is distinct from the notion of a HIBE in that in HIBE schemes the setup of the lower level TAs is closely bound to the upper level TAs, whereas we envisage a scenario where existing single-TA deployments can be made to interact dynamically with minimum overhead, i.e. without reissuing private keys or similar infrastructure overhaul.

A recent paper [13] makes a systematic study of IBE in the setting of multiple TAs, but without TA collaboration of the type considered here. The authors of [13] consider multiple TAs sharing some common parameters, but with each TA generating its own master secret and public parameters. They

provide security notions and security models for this multi-TA setting and their work can be seen as providing natural extensions of security notions and security models for the single-TA setting. Their work does not consider the type of solutions enabling dynamic coalition forming that we propose in this paper.

In the situation where there are no prior relationships between the different TAs except that they may share a subset of common public parameters, if the sender of a message can obtain authentic public parameters for all the TAs, then cross-TA communication can be enabled by encrypting the same message for the identifier in question using the public parameters of each of the TAs. However, such schemes typically lead to significant ciphertext expansion and additional cryptographic computation, which may render them unsuitable in resource-constrained environments. This kind of scenario has been addressed in [17], where a scheme is given that only needs a single pairing operation during encryption. However, the security model used in [17] is the usual one for the single TA setting, and no consideration is given as to how security may be affected by encrypting the same message using multiple, different public parameters. In addition, the schemes of [17] reuse randomness to enhance efficiency, and this is not formally addressed in the security analysis. Barbosa and Farshim [3] do consider the security of multi-recipient IBE with randomness re-use, but only in the single-TA setting.

II. BACKGROUND AND DEFINITIONS

In this section, we provide basic definitions needed for the remainder of the paper.

Definition 1: A pairing-friendly group generator `PairingGen` is a polynomial time algorithm with input 1^k and output a tuple $(\mathbb{G}, \mathbb{G}_T, e, q, P)$. Here \mathbb{G}, \mathbb{G}_T are groups of prime order q , P generates G , and $e : \mathbb{G} \times \mathbb{G} \rightarrow \mathbb{G}_T$ is a bilinear, non-degenerate and efficiently computable map. By convention, \mathbb{G} is an additive group and \mathbb{G}_T multiplicative.

For ease of presentation, we work exclusively in the setting where e is symmetric; our definitions and results can be generalised to the asymmetric setting where $e : \mathbb{G}_1 \times \mathbb{G}_2 \rightarrow \mathbb{G}_T$, with \mathbb{G}_1 and \mathbb{G}_2 being different groups. Further details concerning the basic choices that are available when using pairings in cryptography can be found in [7]

Definition 2: A function $\epsilon(k)$ is said to be *negligible* if, for every c , there exists k_c such that $\epsilon(k) \leq k^{-c}$ for every $k \geq k_c$.

Definition 3: We define the advantage of an algorithm \mathcal{A} in solving the Bilinear Diffie-Hellman (BDH) problem in $(\mathbb{G}, \mathbb{G}_T)$ to be:

$$\text{Adv}_{\mathcal{A}}^{\text{BDH}}(k) = \Pr(\mathcal{A}(aP, bP, cP) = e(P, P)^{abc})$$

where $a, b, c \stackrel{\$}{\leftarrow} \mathbb{Z}_q$. Here, we implicitly assume that parameters $(\mathbb{G}, \mathbb{G}_T, e, q, P)$ are given to \mathcal{A} as additional inputs. We say that the BDH problem is *hard* in $(\mathbb{G}, \mathbb{G}_T)$ if no polynomial-time algorithm that solves the BDH problem in $(\mathbb{G}, \mathbb{G}_T)$ has a non-negligible advantage.

Definition 4: An IBE scheme is defined in terms of four algorithms:

- **Setup**: On input 1^k , outputs a master public key mpk and a master secret key msk . We assume that mpk contains descriptions of the message and ciphertext spaces, MsgSp and CtSp .
- **KeyDer**: A key derivation algorithm that on input mpk , msk and identifier $id \in \{0,1\}^*$, returns a private key usk_{id} . This algorithm may or may not be randomized.
- **Enc**: An encryption algorithm that on input mpk , identifier $id \in \{0,1\}^*$ and message $m \in \text{MsgSp}$, returns a ciphertext $c \in \text{CtSp}$.
- **Dec**: A decryption algorithm that on input mpk , a private key usk_{id} and a ciphertext $c \in \text{CtSp}$, returns either a message $m \in \text{MsgSp}$ or a failure symbol \perp .

These algorithms must satisfy the standard consistency requirement that decryption undoes encryption, i.e. $\forall m \in \text{MsgSp}$,

$$\text{Dec}(mpk, usk_{id}, c) = m \text{ where } c = \text{Enc}(mpk, id, m).$$

A. Security notions for IBE

The standard security notions for IBE, viz. IND-ID-CCA and IND-ID-CPA security, were first established in [5]. (Henceforth we drop the ‘‘ID’’ as we are working exclusively in the identity-based setting.) These security notions model the requirement that by looking at the ciphertext an adversary should not be able to determine to which one of two chosen messages it corresponds. An additional desirable security property is recipient anonymity, which models the requirement that an adversary should not be able to determine who the intended recipient of a ciphertext is, i.e. the ciphertext should not leak the identifier of the intended recipient.

We describe the ANO-IND-CCA security notion which simultaneously captures both the property of message indistinguishability and the property of recipient anonymity, and describe the other security notions by placing certain restrictions on ANO-IND-CCA security.

The ANO-IND-CCA game is defined in terms of the following game between an adversary \mathcal{A} and a challenger \mathcal{C} . \mathcal{C} takes as input the security parameter 1^k , runs algorithm **Setup** of the IBE scheme, gives \mathcal{A} mpk , and keeps msk to itself. \mathcal{A} then runs in two phases:

- **Phase 1**: \mathcal{A} issues a series of adaptively selected key derivation and decryption queries on identities id and identifier/ciphertext combinations (id, c) of its choice. These are replied to by \mathcal{C} by using algorithms **KeyDer** and **Dec** and knowledge of msk .
- **Challenge**: After \mathcal{A} decides to end Phase 1, it outputs two tuples $(id_0, m_0), (id_1, m_1)$ where m_0 and m_1 are equal length messages. \mathcal{C} selects $b \xleftarrow{\$} \{0,1\}$, sets $c^* = \text{Enc}(mpk, id_b, m_b)$ and gives c^* to \mathcal{A} . We require that id_0 and id_1 not be the subject of any key derivation query in Phase 1.
- **Phase 2**: This phase proceeds as in Phase 1, with the constraints that id_0 and id_1 not be the subject of any key derivation query and that (id_0, c^*) and (id_1, c^*) not be the subject of any decryption query.
- **Guess**: \mathcal{A} outputs a bit b' and wins the game if $b' = b$.

The advantage of \mathcal{A} against the IBE scheme in the above ANO-IND-CCA security game is defined to be:

$$\text{Adv}_{\mathcal{A}}^{\text{ANO-IND-CCA}}(k) = \left| \Pr[b = b'] - \frac{1}{2} \right|$$

where the probability is measured over the random choices of coins of \mathcal{A} and \mathcal{C} . An IBE scheme is said to be ANO-IND-CCA secure if the function $\text{Adv}_{\mathcal{A}}^{\text{ANO-IND-CCA}}(k)$ is negligible for all polynomial time adversaries \mathcal{A} .

Removing the adversary’s access to the decryption oracle gives the weaker ANO-IND-CPA security notion.

Setting $id_0 = id_1$ gives the IND-CCA security notion. Setting $id_0 = id_1$ and removing the adversary’s access to the decryption oracle gives the IND-CPA security notion.

Setting $m_0 = m_1$ gives the ANO-CCA security notion. Setting $m_0 = m_1$ and removing the adversary’s access to the decryption oracle gives the ANO-CPA security notion.

Lemma 1: Let IBE be an IBE scheme that is ANO-CPA-secure and IND-CPA-secure. Then IBE is also ANO-IND-CPA-secure.

Proof: The proof (informally) follows by noting that if IBE is IND-CPA-secure, then the challenger may replace the tuple (m_0, id_0) with (m_1, id_0) in its response to the challenge query without the adversary being able to detect the change. Likewise, using ANO-CPA security, the challenger may then replace (m_1, id_0) with (m_1, id_1) . This informal argument can be made rigorous using a sequence of games. ■

B. Boneh-Franklin IBE

We describe the **BasicIdent** IBE scheme from [5].

- **Setup**: On input 1^k the algorithm works as follows
 - Runs algorithm **PairingGen** to obtain $(\mathbb{G}, \mathbb{G}_T, e, q, P)$.
 - Picks $s \xleftarrow{\$} \mathbb{Z}_q^*$.
 - Sets master secret key $msk = s$ and $P_{pub} = sP$.
 - Chooses cryptographic hash functions $H_1 : \{0,1\}^* \rightarrow \mathbb{G}^*$ and $H_2 : \mathbb{G}_T \rightarrow \{0,1\}^n$ for some $n = n(k)$.
 - Sets

$$mpk = (\mathbb{G}, \mathbb{G}_T, e, q, P, P_{pub}, H_1, H_2, n).$$

The message space is $\text{MsgSp} = \{0,1\}^n$.

The ciphertext space is $\text{CtSp} = \mathbb{G}^* \times \{0,1\}^n$.

- **KeyDer**: On input mpk , msk and $id \in \{0,1\}^*$, sets the private key corresponding to id to be $usk_{id} = sH_1(id)$.
- **Enc**: On input mpk , to encrypt a message $m \in \text{MsgSp}$ under the identifier $id \in \{0,1\}^*$,
 - Chooses $r \xleftarrow{\$} \mathbb{Z}_q^*$.
 - Sets the ciphertext to be

$$c = (U, V) = (rP, m \oplus H_2(g_{id}^r))$$

$$\text{where } g_{id} = e(H_1(id), P_{pub}).$$

- **Dec**: On input mpk , $c = (U, V)$, and the private key usk_{id} corresponding to identifier $id \in \{0,1\}^*$, computes

$$m = V \oplus H_2(e(usk_{id}, U)).$$

Lemma 2: BasicIdent is ANO-IND-CPA secure assuming the hardness of the BDH problem in groups generated by PairingGen.

Proof: The BasicIdent scheme was shown in [5] to be IND-CPA secure in the random oracle model, assuming the hardness of the BDH problem in groups generated by PairingGen. The BasicIdent scheme was shown to be ANO-CPA secure in Theorem 4.4 of [1]. Now apply Lemma 1. ■

III. IBE FOR COALITION ENVIRONMENTS

In our approach, each user obtains a single private key corresponding to its identifier from their respective TAs, and later use additional information broadcast by the TAs to translate its private key. The translation allows a user to convert an existing private key into one that would have been issued by any of the other TAs. This in turn allows a sender to encrypt messages to a recipient using the recipient's identifier and the public parameters of any one of the TAs. The recipient can then use translation to obtain a private key that allows it to decrypt the ciphertext. An encryption scheme with these properties can be formally specified by the following 5 algorithms:

- **Setup:** On input a security parameter 1^k and an index $n = n(k)$, $n \geq 2$ denoting the number of TAs
 - Let $\mathcal{TA} = \{ta_i : 1 \leq i \leq n\}$ denote the set of labels of all the TAs.
 - Returns a set of common system parameters $params$ shared by all TAs. We assume that $params$ includes a description of the message space $MsgSp$ and the ciphertext space $CtSp$.
 - For each $ta \in \mathcal{TA}$, generates a unique master public key component mpk_{ta} and a master secret key msk_{ta} .
 - For every pair of TAs, $ta_i, ta_j \in \mathcal{TA}$ takes as input $params, msk_{ta_i}, msk_{ta_j}$, and returns a relation rel_{ta_i, ta_j} from a space \mathcal{R} of relations. Let \mathcal{R}_{ta} denote the set of all relations pertaining to $ta \in \mathcal{TA}$.
 - The master public key for $ta \in \mathcal{TA}$ is

$$mpk_{ta} = (params, mpk_{ta}).$$

- **KeyDer:** On input mpk_{ta}, msk_{ta} for $ta \in \mathcal{TA}$ and an identifier $id \in \{0, 1\}^*$, returns the private key $usk_{id, ta}$ corresponding to id for TA ta .
- **TranslateKey:** On input usk_{id, ta_i} and $rel_{ta_i, ta_j} \in \mathcal{R}_{ta_i}$ returns usk_{id, ta_j} .
- **Enc:** On input mpk_{ta} for $ta \in \mathcal{TA}$, $id \in \{0, 1\}^*$, $m \in MsgSp$, returns a ciphertext $c \in CtSp$.
- **Dec:** On input mpk_{ta} for $ta \in \mathcal{TA}$, $c \in CtSp$ and private key $usk_{id, ta}$, returns $m \in MsgSp$ or a failure symbol \perp .

These algorithms must satisfy the standard consistency requirement that $\forall m \in MsgSp, \forall id \in \{0, 1\}^*$ and $\forall ta \in \mathcal{TA}$,

$$Dec(c, usk_{id, ta}) = m \text{ where } c = Enc(mpk_{ta}, id, m).$$

We assume here that the relations \mathcal{R}_{ta} pertaining to each $ta \in \mathcal{TA}$ are broadcast to all users in the system. After this is done, a user with identifier id and private key usk_{id, ta_i} may use the relation $rel_{ta_i, ta_j} \in \mathcal{R}_{ta_i}$ in Algorithm

TranslateKey to obtain usk_{id, ta_j} . This enables that user to decrypt ciphertexts computed using identifier id and the master public key of ta_j .

A. Security Notions

We introduce security models appropriate to the above cryptographic primitive. We define the mANO-IND-CCA security game to capture an indistinguishability and anonymity notion of security similar to that in a single-TA setting. The mANO-IND-CCA security game is defined in terms of the following game between an adversary \mathcal{A} and a challenger \mathcal{C} . \mathcal{C} takes as input the security parameter 1^k and an index $n = n(k)$, $n \geq 2$. It runs algorithm Setup and gives \mathcal{A} the master public keys for all TAs, $ta \in \mathcal{TA}$, along with all the data \mathcal{R}_{ta} . It keeps all the master secret keys to itself.

- **Phase 1:** \mathcal{A} issues a series of adaptively selected key derivation and decryption queries on TA/identifier (ta, id) and TA/identifier/ciphertext (ta, id, c) combinations of its choice. These are replied to by \mathcal{C} by using algorithms KeyDer and Dec and its knowledge of the master secret keys.
- **Challenge:** After \mathcal{A} decides to end Phase 1, it outputs the tuples $(m_0, id_0), (m_1, id_1)$ and a challenge TA, ta . Here m_0 and m_1 are equal length messages and id_0 and id_1 must not be the subject of any KeyDer query in Phase 1. \mathcal{C} selects $b \xleftarrow{\$} \{0, 1\}$ and sets $c^* = Enc(mpk_{ta}, id_b, m_b)$. \mathcal{C} gives c^* to \mathcal{A} .
- **Phase 2:** This phase proceeds as in phase 1, with the constraint that id_0 and id_1 are not the subject of any key derivation query and (ta, id_0, c^*) and (ta, id_1, c^*) are not the subject of any decryption query.
- **Guess:** \mathcal{A} outputs a bit b' and wins the game if $b' = b$.

The advantage of \mathcal{A} against the IBE scheme in the mANO-IND-CCA security game is defined to be:

$$Adv_{\mathcal{A}}^{mANO-IND-CCA}(k) = \left| \Pr[b = b'] - \frac{1}{2} \right|$$

where the probability is measured over the random choices of coins of \mathcal{A} and \mathcal{C} . An IBE scheme is said to be mANO-IND-CCA secure if the function $Adv_{\mathcal{A}}^{mANO-IND-CCA}(k)$ is negligible for all polynomial time adversaries \mathcal{A} .

Removing the adversary's access to the decryption oracle gives us the mANO-IND-CPA security notion.

Setting $id_0 = id_1$ gives the m-IND-CCA security notion. Setting $id_0 = id_1$ and removing the adversary's access to the decryption oracle gives the m-IND-CPA security notion.

Setting $m_0 = m_1$ gives the m-ANO-CCA security notion. Setting $m_0 = m_1$ and removing the adversary's access to the decryption oracle gives the m-ANO-CPA security notion.

The corresponding single-TA security notions can also be obtained by setting $n = 1$.

B. An instantiation based on the Boneh-Franklin scheme

We give an instantiation based on the BasicIdent scheme of [5].

- **Setup:** On input 1^k and $n = n(k)$, $n \geq 2$

- Let $\mathcal{TA} = \{ta_i : 1 \leq i \leq n\}$ denote the set of labels of all TAs.
- Runs `PairingGen` to obtain $(\mathbb{G}, \mathbb{G}_T, e, q, P)$.
- Chooses cryptographic hash functions $H_1 : \{0, 1\}^* \rightarrow \mathbb{G}^*$ and $H_2 : \mathbb{G}_T \rightarrow \{0, 1\}^l$.
- Sets $params = (\mathbb{G}, \mathbb{G}_T, e, q, P, H_1, H_2, l, k)$. The message space is $\text{MsgSp} = \{0, 1\}^{l-k}$ and k is the length of padding employed. The ciphertext space is $\text{CtSp} = \mathbb{G}^* \times \{0, 1\}^l$.
- For $1 \leq i \leq n$ sets $s_i \xleftarrow{\$} \mathbb{Z}_q^*$ and sets $msk_i = s_i$ and $mpku_i = s_i P$.
- For every $ta_i \in \mathcal{TA}$ generates the set

$$\mathcal{R}_i = \{rel_{ta_i, ta_j} : 1 \leq j \leq n\}$$

where $rel_{ta_i, ta_j} = s_j \cdot s_i^{-1} \bmod q$.

- Sets $mpk_i = (params, mpku_i)$.
- `KeyDer`: On input mpk_i , msk_i and an identifier $id \in \{0, 1\}^*$, sets the the private key corresponding to id for ta_i to be $usk_{id, ta_i} = s_i H_1(id)$.
- `Enc`: On input mpk_i for $ta_i \in \mathcal{TA}$, to encrypt a message $m \in \text{MsgSp}$, under $id \in \{0, 1\}^*$
 - Pads m with k zero bits.
 - Chooses $r \xleftarrow{\$} \mathbb{Z}_q^*$.
 - Sets the ciphertext to be

$$c = (U, V) = (rP, (m || 0^k) \oplus H_2(g_{id}^r))$$

where $g_{id} = e(H_1(id), mpku_i)$.

- `TranslateKey`: On input usk_{id, ta_i} , the private key corresponding to $id \in \{0, 1\}^*$ for $ta_i \in \mathcal{TA}$, and $rel_{ta_i, ta_j} \in \mathcal{R}_i$ returns $rel_{ta_i, ta_j} \cdot usk_{id, ta_i}$. Note that

$$\begin{aligned} rel_{ta_i, ta_j} \cdot usk_{id, ta_i} &= s_j \cdot s_i^{-1} \cdot s_i H_1(id) \\ &= s_j H_1(id) \\ &= usk_{id, ta_j}. \end{aligned}$$

so that algorithm `TranslateKey` does indeed convert private keys correctly.

- `Dec`: On input mpk_i , $c = (U, V)$ the recipient with private key usk_{id, ta_i} does the following:
 - Runs `TranslateKey` to obtain usk_{id, ta_j} for all $1 \leq j \leq n, j \neq i$. (Note that this step needs to be performed just once if the user can store all the keys usk_{id, ta_j} .)
 - Computes $m'_j = V \oplus H_2(e(U, usk_{id, ta_j}))$ for all $1 \leq j \leq n$.
 - For each m'_j checks if the last k bits are zero. (Parameter k is selected such that with overwhelming probability only the decryption of c using the correct private key yields this padding format.) If it is, it sets m to the left $l - k$ bits of m'_j and outputs m , else outputs \perp . We note that if we assume that ciphertexts contain a label indicating which TA's master public key was used to perform the encryption, then the decryption operation can be done in a single step using the correct private key.

Note that, for this particular scheme, the complete set of n^2 values rel_{ta_i, ta_j} can be computed upon receipt of a single broadcast containing the values rel_{ta_i, ta_j} for $j = 2, \dots, n$. This means that this specific scheme needs a broadcast whose size is linear in the number of coalition TAs, rather than quadratic.

Theorem 1: The multi-TA scheme based on `BasicIdent` is mANO-IND-CPA secure assuming the hardness of the BDH problem in groups generated by `PairingGen`.

Proof: Suppose there is an mANO-IND-CPA adversary \mathcal{A} against the multi-TA IBE scheme with advantage ε and running in time t . We show how to construct an algorithm \mathcal{B} that uses \mathcal{A} to break the ANO-IND-CPA property of the `BasicIdent` scheme.

\mathcal{B} 's inputs are the parameters $mpk = (\mathbb{G}, \mathbb{G}_T, e, q, P, P_{pub}, H_1, H_2, l)$ of the `BasicIdent` scheme. \mathcal{B} 's task is to break the ANO-IND-CPA security of the `BasicIdent` scheme and it does this by acting as a challenger for \mathcal{A} .

\mathcal{B} generates the parameters of the multi-TA IBE scheme. It sets $params = (\mathbb{G}, \mathbb{G}_T, e, q, P, H_1, H_2, l, k)$ where k is the length of padding employed. Let $\mathcal{TA} = \{ta_i : 1 \leq i \leq n\}$ denote the set of labels of all TAs where n is the number of TAs. Then \mathcal{B} sets $mpku_I = P_{pub}$ for $ta_I \in \mathcal{TA}$ with $I \xleftarrow{\$} \{1 \dots n\}$. Note that $P_{pub} = sP$ for $s \in \mathbb{Z}_q^*$ but the value of s is not known to \mathcal{B} . It sets $mpku_j = \lambda_j \cdot P_{pub}$ for ta_j , for each $1 \leq j \leq n, j \neq I$, where each λ_j is drawn uniformly at random from \mathbb{Z}_q^* . It also sets $\lambda_I = 1$. The master secret key for ta_j is $msk_j = \lambda_j \cdot s$ which, again, \mathcal{B} does not know. However, with the knowledge of the λ_j values it is able to generate, for every $ta_i \in \mathcal{TA}$, the appropriate set $\mathcal{R}_i = \{rel_{ta_i, ta_j} : 1 \leq j \leq n\}$. \mathcal{B} does this by setting $rel_{ta_i, ta_j} = \lambda_j \cdot \lambda_i^{-1} \bmod q$. \mathcal{B} then sets $mpk_i = (params, mpku_i)$ for each $ta_i \in \mathcal{TA}$ and gives these public keys to \mathcal{A} , along with the sets of relations $\mathcal{R}_i, 1 \leq j \leq n$.

\mathcal{A} makes a series of queries which \mathcal{B} answers as follows.

- **Phase 1:** In Phase 1 of the attack \mathcal{A} makes a series of key derivation queries on (ta, id) combinations. \mathcal{B} asks the corresponding key derivation query on id to its challenger which responds with the private key $usk_{id, ta}$. If $ta = ta_i$, \mathcal{B} relays the response to \mathcal{A} . Otherwise, \mathcal{B} uses the appropriate λ_j value to compute $usk_{id, ta_j} = \lambda_j \cdot usk_{id, ta_i}$ which corresponds to the private key of id for ta_j and sends this to \mathcal{A} .
- **Challenge:** When \mathcal{A} decides to end Phase 1 of the attack it outputs two messages/identifier tuples (m_0, id_0) and (m_1, id_1) and the TA, ta , on which it wishes to be challenged, subject to the condition that no key derivation query was asked on id_0 or id_1 in Phase 1. \mathcal{B} pads the two messages m_0 and m_1 with k zero bits to obtain $m'_0 = m_0 || 0^k$ and $m'_1 = m_1 || 0^k$ and sends the tuples (m'_0, id_0) and (m'_1, id_1) to its challenger and receives the `BasicIdent` encryption $c^* = (U, V)$ corresponding to m'_b and id_b for ta_i , for a bit b chosen uniformly at random. If $ta = ta_i$ then \mathcal{B} relays $c^{*'} = c^*$ to \mathcal{A} . Otherwise, \mathcal{B} sets $c^{*'} = (\lambda_j^{-1} \cdot U, V)$ which corresponds to the encryption of m'_b to identifier id_b for ta_j and sends this to \mathcal{A} .

- Phase 2: Phase 2 of the attack proceeds as in Phase 1 with the restriction that no key derivation query is allowed on id_0 or id_1 .

This completes our description of \mathcal{B} 's simulation. Note that \mathcal{A} 's view of the simulation is identical to its view in a real attack. All the queries are responded to correctly. When \mathcal{A} terminates by outputting a bit b' , \mathcal{B} simply relays this bit to its challenger. Clearly \mathcal{B} 's advantage in breaking the ANO-IND-CPA security of `BasicIdent` is equal to \mathcal{A} 's advantage against the mANO-IND-CPA security of the multi-TA IBE scheme. We know that \mathcal{B} 's advantage against `BasicIdent` is negligible and hence \mathcal{A} must have negligible advantage against the multi-TA scheme. ■

IV. OPEN PROBLEMS

We pose a number of open problems as a consequence of this work. It will be interesting to produce schemes where the relations can be derived without the TAs having to share their master secret keys, or to prove that this is impossible. Schemes secure in stronger models, for example where a subset of TAs can be corrupted, or which tolerate the corruption of private keys corresponding to an identifier for a subset of TAs pose additional challenges. Finally, schemes which are CCA secure and/or which do not make use of random oracles in the security analysis merit investigation.

ACKNOWLEDGEMENTS

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

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

REFERENCES

- [1] M. Abdalla, M. Bellare, D. Catalano, E. Kiltz, T. Kohno, T. Lange, J. Malone-Lee, G. Neven, P. Paillier, and H. Shi. Searchable encryption revisited: Consistency properties, relation to anonymous IBE, and extensions. In V. Shoup, editor, *CRYPTO*, volume 3621 of *Lecture Notes in Computer Science*, pages 205–222. Springer, 2005.
- [2] S. Balfé, K. Boklan, Z. Klagsbrun, and K.G. Paterson. Key refreshing in identity-based cryptography and its applications in MANETs. In *IEEE Milcom 2007*.
- [3] M. Barbosa and P. Farshim. Efficient identity-based key encapsulation to multiple parties. In N. P. Smart, editor, *IMA Int. Conf.*, volume 3796 of *Lecture Notes in Computer Science*, pages 428–441. Springer, 2005.
- [4] M. Bellare and P. Rogaway. Random oracles are practical: A paradigm for designing efficient protocols. In *ACM Conference on Computer and Communications Security*, pages 62–73, 1993.
- [5] D. Boneh and M. K. Franklin. Identity-based encryption from the Weil pairing. In J. Kilian, editor, *CRYPTO*, volume 2139 of *Lecture Notes in Computer Science*, pages 213–229. Springer, 2001.
- [6] D.W. Carman. New directions in sensor network key management. *International Journal of Distributed Sensor Networks*, 1(1):3–15, 2004.
- [7] S.D. Galbraith, K.G. Paterson, and N.P. Smart. Pairings for cryptographers. *Discrete Applied Mathematics*, to appear, 2008.
- [8] C. Gentry and A. Silverberg. Hierarchical ID-based cryptography. In Y. Zheng, editor, *ASIACRYPT*, volume 2501 of *Lecture Notes in Computer Science*, pages 548–566. Springer, 2002.
- [9] K. Hoepfer and G. Gong. Key revocation for identity-based schemes in mobile ad hoc networks. In T. Kunz and S.S. Ravi (eds.), *Ad-Hoc, Mobile, and Wireless Networks, 5th International Conference*, volume 4104 of *Lecture Notes in Computer Science*, pages 224–237. Springer, 2006.
- [10] J. Horwitz and B. Lynn. Toward hierarchical identity-based encryption. In L. R. Knudsen, editor, *EUROCRYPT*, volume 2332 of *Lecture Notes in Computer Science*, pages 466–481. Springer, 2002.
- [11] A. Khalili, J. Katz, and W.A. Arbaugh. Toward secure key distribution in truly ad-hoc networks. In *SAINT-W '03: Proceedings of the 2003 Symposium on Applications and the Internet Workshops (SAINT'03 Workshops)*, pp. 342–346, Washington, DC, USA, 2003. IEEE Computer Society.
- [12] B.J. Matt. Toward hierarchical identity-based cryptography for tactical networks. In *MILCOM '03: Proceedings of the 2004 Military Communications Conference*. IEEE Computer Society, 2004.
- [13] K.G. Paterson and S. Srinivasan. Security and anonymity of identity-based encryption with multiple trusted authorities. In S.D. Galbraith and K.G. Paterson, editors, *Pairing 2008*, volume 5209 of *Lecture Notes in Computer Science*, pages 354–375. Springer, 2008.
- [14] D. Roberts, G. Lock and D.C. Verma. Holistan: A futuristic scenario for international coalition operations. In *Proceedings of the Fourth International Conference on Knowledge Systems for Coalition Operations (KSCO-2007)*, Waltham, MA, USA, 1-2 May 2007.
- [15] R. Sakai, K. Ohgishi, and M. Kasahara. Cryptosystems based on pairing. *The 2000 Symposium on Cryptography and Information Security, Okinawa, Japan, January*, pages 26–28, 2000.
- [16] A. Shamir. Identity-based cryptosystems and signature schemes. In *CRYPTO*, pages 47–53, 1984.
- [17] S. Wang and Z. Cao. Practical identity-based encryption (IBE) in multiple PKG environments and its applications. *Cryptology ePrint Archive*, Report 2007/100, 2007. <http://eprint.iacr.org/>.