

Identity-Based Proxy Re-encryption

Matthew Green and Giuseppe Ateniese

The Johns Hopkins University
Department of Computer Science
3400 N. Charles Street; Baltimore, MD 21218, USA
{mgreen,ateniese}@cs.jhu.edu

Abstract. In a proxy re-encryption scheme a semi-trusted proxy converts a ciphertext for Alice into a ciphertext for Bob without seeing the underlying plaintext. A number of solutions have been proposed in the public-key setting. In this paper, we address the problem of Identity-Based proxy re-encryption, where ciphertexts are transformed from one *identity* to another. Our schemes are compatible with current IBE deployments and do not require any extra work from the IBE trusted-party key generator. In addition, they are non-interactive and one of them permits multiple re-encryptions. Their security is based on a standard assumption (DBDH) in the random oracle model.

Keywords: proxy re-encryption, identity-based encryption, bilinear maps.

1 Introduction

In a proxy re-encryption scheme, a proxy can convert an encryption computed under Alice's public-key into an encryption intended for Bob. Such a scheme can be used by Alice to temporarily forward encrypted messages to Bob without giving him her secret key. The fundamental property of proxy re-encryption schemes is that the proxy is not fully trusted, i.e., it does not know the secret keys of Alice or Bob and does not learn the plaintext during the conversion. The proxy and Bob, however, are not allowed to collude, thus it is usually assumed that at least one of the two is honest or that their collusion is preventable or detectable via other means.

A number of proxy re-encryption protocols have been proposed in the context of public-key encryption [1,2,3,4,5]. In this work we extend the notion of proxy re-encryption to the area of Identity-Based Encryption (IBE), in which senders encrypt messages using the recipient's identity (a string) as the public key. For example, Charles could encrypt a message for Alice by just using her email address. First introduced by Shamir in 1984 and then realized by Boneh-Franklin [6] and by Cocks [7] several years later, identity-based encryption has proven useful in solving many key-distribution issues, and has facilitated the development of a variety of novel cryptographic protocols, *e.g.*, secret handshakes [8], public-key searchable encryption [9,10], CCA2-secure public-key encryption [11], and digital signatures [12]. The Boneh-Franklin scheme is particularly efficient, and has been commercially deployed [13].

Our identity-based proxy re-encryption (IB-PRE) schemes allow a proxy to translate a ciphertext encrypted under Alice’s identity into one computed under Bob’s identity. To permit this translation, Alice generates and provisions the proxy with a *delegation key* (or “*re-encryption key*”), that the proxy uses to perform the re-encryption. No information about the secret keys of Alice or Bob can be deduced from this value, nor does the proxy learn anything about the underlying plaintext of the messages it processes. Our constructions are compatible with existing Boneh-Franklin IBE deployments, and can be implemented using existing secrets and parameters.

Users in an Identity-Based Encryption scheme request keys from a trusted party known as a Private Key Generator (PKG). Thus, in principle, it is possible that delegation keys could be generated by the PKG directly, rather than by individual scheme users. However, we categorically exclude this possibility and we focus only on *non-interactive* schemes where individual users delegate their own decryption rights without the involvement of the Private Key Generator. This is for theoretical and practical reasons: (1) From a theoretical point of view, having the PKG, or any other trusted party, generating the proxy keys makes the problem of finding IB-PRE schemes quite unchallenging given prior art, (2) from a practical point of view, it is clearly undesirable to have the PKG involved in the generation of proxy keys. It would constitute a considerable bottleneck in many applications, it would force the PKG to be *on-line* and available even during the generation of proxy keys (other than IBE keys), and, in certain applications, it would make the PKG liable for creating (potentially unwanted) decryption rights.

Previous Work. Mambo and Okamoto proposed a technique for delegating decryption rights in [1]. Blaze, Bleumer and Strauss [2] later presented the first secure “atomic” re-encryption primitive: an Elgamal-based scheme in which the proxy could not learn the message being processed. Unfortunately, the BBS approach was inherently bidirectional: a corrupted proxy could re-encrypt ciphertexts not only from Alice to Bob, but also from Bob to Alice. Jakobsson [4], and Zhou, *et. al.* [14] addressed this collusion problem via quorum-based protocols which divided the proxy into many distinct components.

More recent works have focused on *unidirectional* schemes, where collusion between a delegator and the proxy does not compromise the delegatee. Dodis and Ivan [5] realized a form of unidirectional proxy *encryption* by doubly-encrypting messages under two separate keys (or by splitting a single decryption key into two parts). Their approach permits a form of proxy re-encryption when parties pre-distribute shared secrets. Ateniese, Fu, Green and Hohenberger [3] proposed several non-interactive unidirectional proxy re-encryption schemes that eliminated the need for pre-shared keys and permitted arbitrary delegations. That work left an interesting open problem, which we address in this paper: namely, to construct chosen-ciphertext secure (CCA) proxy re-encryption schemes. Canetti and Hohenberger [15] also addressed this problem in the public key setting, proposing a CCA-secure bidirectional proxy re-encryption scheme. Though the

constructions differ from ours, their security definition is compatible and we adopt some aspects of their presentation for consistency.

Finally, Boneh, Goh and Matsuo [16] presented a “hybrid” form of proxy re-encryption based on IBE. In such schemes, the PKG performs all delegations; thus users are unable to perform offline (“non-interactive”) delegations and each delegation requires an online request to the PKG. Furthermore, the Boneh-Goh-Matsuo approach specifies a new private-key generation algorithm and it seems therefore incompatible with existing IBE deployments.

Paper Outline. The outline of the rest of this paper is as follows. In section 3 we present definitions for Identity-Based Proxy Re-encryption and for the hardness assumptions used in our proofs. In section 4 we introduce our constructions. In section 5 we discuss several applications for the new primitives. Finally, section 6 lists open research problems and provides our conclusions.

2 Properties of Our Schemes

Ateniese *et. al.* [3] proposed a series of properties by which to evaluate proxy re-encryption schemes. We briefly reiterate some of these properties, in particular those that our scheme provides and that, we believe, are relevant for practical instantiations of Identity-Based Proxy Re-encryption.

- *Unidirectionality.* A unidirectional scheme permits user A to delegate to user B , without permitting A to decrypt user B 's ciphertexts.
- *Non-Interactivity.* Non-interactive schemes permit user A to construct a re-encryption key $rk_{id_A \rightarrow id_B}$ while offline, (*i.e.*, without the participation of B or the Private Key Generator).
- *Multiple-use capability.* A multi-use scheme permits the proxy (or proxies) to perform multiple consecutive re-encryptions on a ciphertext, *e.g.*, re-encrypt from id_A to id_B , then re-encrypt the result from id_B to id_C and so on.
- *Space-optimality.* Many existing schemes (*e.g.*, [5,16,3]) incur additional communication costs in order to support re-encryption. This inefficiency takes several common forms, including: (*a*) ciphertext expansion upon re-encryption (see the practical implementations of [3]), (*b*) a required pre-distribution stage in which secrets are shared with delegates (as in [5]), or (*c*) the inclusion of ciphertext material that is discarded during re-encryption (see [16]).

In this paper we focus on unidirectional schemes only. Notice that a bidirectional scheme can always be achieved by running a unidirectional one in both directions, *i.e.*, from Alice to Bob and vice versa. Thus, a unidirectional IB-PRE is clearly a more powerful primitive than a bidirectional one but also harder to devise.

In addition, we believe that non-interactivity is a fundamental property and our schemes provide it. In a non-interactive scheme, Alice can generate the re-encryption key from Bob's identity, without ever involving Bob. In the identity-based setting, this property provides an interesting twist: Alice can delegate

decryption rights to delegates that do not exist yet or will join the system later. Moreover, as noted by Boneh and Franklin [6], identities can be seen as credentials and express conditions. For instance, an encryption under “*Alice || security-clearance || time period*” can be opened by Alice only if she has security clearance and within the time period specified in the string. Analogously, in our schemes, Alice can specify the conditions under which the delegation of decryption rights has to happen. We will explore applications of this feature in section 5.

In section 4.2 we discuss an optimization that provides for space-optimal proxy re-encryption in some circumstances. Finally, one of our schemes is *multi-use* in the sense that once a re-encryption from Alice to Bob is computed, the resulting ciphertext can be re-encrypted again from Bob to Charles, etc., multiple times. Finding a unidirectional and multi-use scheme was left as an open problem in prior art for the public-key case. We show how to achieve this property for our IB-PRE but at the cost of allowing the ciphertext to expand linearly with respect to the number of re-encryptions (however, this appears to be inevitable for a non-interactive scheme).

3 Definitions

We begin by describing the setting and computational problems used within this work. We then formally define an Identity-Based Proxy Re-encryption scheme and propose a new, generalized security definition.

Definition 1 (Bilinear Map). We say a map $e : \mathbb{G} \times \mathbb{G} \rightarrow \mathbb{G}_T$ is a *bilinear map* if:

1. \mathbb{G}, \mathbb{G}_T are groups of the same prime order q .
2. For all $a, b \in \mathbb{Z}_q^*$, $g \in \mathbb{G}$, $e(g^a, g^b) = e(g, g)^{ab}$.
3. The map is non-degenerate (i.e., if $\mathbb{G} = \langle g \rangle$, then $\mathbb{G}_T = \langle e(g, g) \rangle$).
4. e is efficiently computable.

For simplicity our constructions are defined in the *symmetric* setting as above. However they also work in the *asymmetric* setting with a bilinear map of the form: $\hat{e} : \mathbb{G}_1 \times \mathbb{G}_2 \rightarrow \mathbb{G}_T$.

Definition 2 (Decisional Bilinear Diffie Hellman Assumption (DBDH)). Our schemes are based on the assumed intractability of the Decisional Bilinear Diffie-Hellman problem (DBDH) in \mathbb{G}, \mathbb{G}_T . This assumption is believed to hold in certain groups, and used as the basis of several Identity-Based Encryption schemes, e.g., [17,18].

We define the DBDH problem as follows: Let $(\mathbb{G}, \mathbb{G}_T)$ be a pair of bilinear groups with an efficiently computable pairing $e : \mathbb{G} \times \mathbb{G} \rightarrow \mathbb{G}_T$, and let g be a random generator of \mathbb{G} . The DBDH problem is to decide, given a tuple of values $(g, g^a, g^b, g^c, T) \in \mathbb{G}^4 \times \mathbb{G}_T$ (where $a, b, c \in_R \mathbb{Z}_q^*$), whether $T = e(g, g)^{abc}$ or if T is a random element of \mathbb{G}_T .

Let k be a security parameter of sufficient size. Formally, we say that the DBDH assumption holds in $(\mathbb{G}, \mathbb{G}_T)$ if for all probabilistic polynomial time algorithms \mathcal{A} , the following condition is true:

$$\left| \Pr \left[a, b, c \xleftarrow{\$} \mathbb{Z}_q^*; 1 \leftarrow \mathcal{A}(g, g^a, g^b, g^c, e(g, g)^{abc}). \right] - \Pr \left[a, b, c \xleftarrow{\$} \mathbb{Z}_q^*; T \xleftarrow{\$} \mathbb{G}_T; 1 \leftarrow \mathcal{A}(g, g^a, g^b, g^c, T). \right] \right| \leq \nu(k)$$

Where $\nu(\cdot)$ is defined as a *negligible* function, *i.e.*, for all polynomial functions $p(\cdot)$, $\nu(k) < 1/p(k)$.

3.1 Identity-Based Proxy Re-encryption

An Identity-Based Proxy Re-encryption (IB-PRE) scheme is an extended Identity Based Encryption scheme. The first extension is an algorithm that generates *re-encryption keys* that can be given to the proxy. The proxy uses the second algorithm to apply these re-encryption keys to ciphertexts and “atomically” re-encrypt them from one identity to another. In a *non-interactive* scheme, re-encryption keys may be generated by the delegator using only her IBE secret key—the IBE master secret is not required.

Encryption Levels. Our definitions refer to the notion of an “encryption level” as an implicit property of a ciphertext. A ciphertext generated directly using the `Encrypt` algorithm is termed a “level-1” ciphertext. Applying the re-encryption algorithm to a level- ℓ ciphertext results in a level- $(\ell + 1)$ ciphertext. Specific constructions may optionally place bounds on the number of consecutive re-encryptions; for instance, non-“multi-use” schemes such as [5,16,3] are limited to a single re-encryption. In our definitions below, we define `MaxLevels` as the highest-possible encryption level (for a single-use scheme, this value is 2).

Definition 3 (Non-interactive Identity-Based Proxy Re-encryption (IB-PRE)). A *non-interactive identity-based proxy re-encryption scheme* is tuple of algorithms (`Setup`, `KeyGen`, `Encrypt`, `Decrypt`, `RKGen`, `Reencrypt`):

- `Setup`($1^k, \text{MaxLevels}$) accepts a security parameter and optionally a value indicating the maximum number of consecutive re-encryptions permitted by the scheme. The algorithm outputs both the master public parameters (`params`) which are distributed to users, and the master secret key (`msk`) which is kept private.
- `KeyGen`(`params`, `msk`, id) on input an identity $id \in \{0, 1\}^*$ and the master secret key, outputs a decryption key sk_{id} corresponding to that identity.
- `Encrypt`(`params`, id, m) on input a set of public parameters, an identity $id \in \{0, 1\}^*$, and a plaintext $m \in \mathcal{M}$, output c_{id} , the encryption of m under the specified identity.
- `RKGen`(`params`, sk_{id_1}, id_1, id_2) on input a secret key sk_{id_1} (derived via the `KeyGen` algorithm) and identities $(id_1, id_2) \in \{0, 1\}^*$, outputs a *re-encryption key* $rk_{id_1 \rightarrow id_2}$.

- $\text{Reencrypt}(\text{params}, rk_{id_1 \rightarrow id_2}, c_{id_1})$ on input a ciphertext c_{id_1} under identity id_1 , and a re-encryption key $rk_{id_1 \rightarrow id_2}$ (generated by the RGen routine), outputs a “re-encrypted” ciphertext c_{id_2} .
- $\text{Decrypt}(\text{params}, sk_{id}, c_{id})$ decrypts the ciphertext c_{id} using the secret key sk_{id} , and outputs a plaintext or the distinguished symbol \perp .

Correctness. Intuitively, an IB-PRE scheme is *correct* if the Decrypt algorithm always outputs the expected decryption of a properly-generated ciphertext (when supplied with the appropriate decryption key). We define “proper generation” as the process of (1) encrypting a plaintext using Encrypt, and subsequently (2) iteratively applying the Reencrypt algorithm up to $\text{MaxLevels} - 1$ times using valid re-encryption keys.

Slightly more formally, let $c_{id_1} \leftarrow \text{Reencrypt}^n(\dots, \text{Encrypt}(\text{params}, \cdot, m))$ be a properly-generated ciphertext. Then $\forall m \in \mathcal{M}, \forall id_1, id_2 \in \{0, 1\}^*, \forall n < \text{MaxLevels} - 1$, where $sk_{id_1} = \text{KeyGen}(\text{msk}, id_1)$, $sk_{id_2} = \text{KeyGen}(\text{msk}, id_2)$, $rk_{id_1 \rightarrow id_2} \leftarrow \text{RGen}(\text{params}, sk_{id_1}, id_1, id_2)$, the following propositions hold:

- $\text{Decrypt}(\text{params}, sk_{id_1}, c_{id_1}) = m$
- $\text{Decrypt}(\text{params}, sk_{id_2}, \text{Reencrypt}(\text{params}, rk_{id_1 \rightarrow id_2}, c_{id_1})) = m$

Security. Security definitions for Identity-Based Encryption (see [6]) address the case where keyholders collude by combining secrets. Identity-Based Proxy re-encryption schemes require a further extension of this collusion guarantee, to model the presence of colluding *proxies* provisioned with re-encryption keys. Many existing security definitions (e.g., [2,5,3]) address the proxy via separate definitional games. We choose instead to incorporate all of these properties into a single game, by providing re-encryption keys to the adversary via an oracle.

When the adversary possesses re-encryption keys, we must naturally restrict it in some ways to avoid a trivial condition, e.g., to prevent it from obtaining a set of re-encryption keys leading from the challenge identity id^* to some identity for which the adversary holds a decryption key. In the CCA case, these restrictions are more complex. To simplify the presentation, we adopt the notion of *derivative* ciphertexts introduced in [15].

Definition 4 (Security of Non-Interactive Identity Based Proxy Re-Encryption (IND-prID-CPA, IND-prID-CCA)). Let \mathcal{S} be an IB-PRE scheme defined as a tuple of algorithms (Setup, KeyGen, Encrypt, Decrypt, RGen, Reencrypt). Security is defined according to the following game $\text{Exp}^{\mathcal{A}, \text{IND-prID-ATK}}$, where $\text{ATK} \in \{\text{CPA}, \text{CCA}\}$.

1. **SETUP.** Run Setup(1^k) to get (params, msk), and give params to \mathcal{A} .
2. **FIND PHASE.** \mathcal{A} makes the queries (extract, rkextract, decrypt, reencrypt). If $\text{ATK} = \text{CPA}$, the queries (decrypt, reencrypt) are answered with \perp .
 - On (extract, id), return KeyGen(params, msk, id).
 - On (rkextract, id_1, id_2), extract the key $sk_{id_1} = \text{KeyGen}(\text{params}, \text{msk}, id_1)$ and return RGen(params, sk_{id_1}, id_1, id_2).

- On $(\text{decrypt}, id, c)$: Extract $sk_{id} = \text{KeyGen}(\text{params}, \text{msk}, id)$ and return $\text{Decrypt}(\text{params}, sk_{id}, c)$.
- On $(\text{reencrypt}, id_1, id_2, c)$: Extract $sk_{id} = \text{KeyGen}(\text{params}, \text{msk}, id)$, and derive a re-encryption key $rk_{id_1 \rightarrow id_2} = \text{RKGen}(\text{params}, sk_{id_1}, id_1, id_2)$. Return $\text{Reencrypt}(\text{params}, rk_{id_1 \rightarrow id_2}, id_1, id_2, c)$.

At the conclusion of this phase \mathcal{A} selects $id^* \in \{0, 1\}^*$ and $(m_0, m_1) \in \mathcal{M}^2$. \mathcal{A} is restricted to choices of id^* such that “trivial” decryption is not possible using keys extracted during this phase (e.g., by using re-encryption keys to translate from id^* to identity id' for which \mathcal{A} holds a decryption key).

3. CHALLENGE. When \mathcal{A} presents $(\text{choice}, id^*, m_0, m_1)$, select $i \xleftarrow{\$} \{0, 1\}$ and compute $c^* = \text{Encrypt}(\text{params}, id^*, m_i)$. Return c^* to \mathcal{A} .
4. GUESS STAGE. \mathcal{A} makes queries as in the FIND stage, with the following restrictions.
 - (a) \mathcal{A} is restricted from querying on $(\text{decrypt}, id, c)$ if $\langle id, c \rangle$ is a challenge derivative. This notion is defined inductively (as in [15]):
 - i. $\langle id^*, c^* \rangle$ is a challenge derivative.
 - ii. If $\langle id, c \rangle$ is a challenge derivative, and \mathcal{A} has issued the query $(\text{reencrypt}, id, id', c)$ to receive a value c' , then $\langle id', c' \rangle$ is a challenge derivative.
 - iii. If $\langle id, c \rangle$ is a challenge derivative, \mathcal{A} has issued query (rkgen, id, id') to receive $rk_{id \rightarrow id'}$, and $c' = \text{Reencrypt}(rk_{id_1 \rightarrow id_2}, id, id', c)$, then $\langle id', c' \rangle$ is a challenge derivative.
 - (b) \mathcal{A} is restricted from querying on $(\text{extract}, id)$ if there exists a challenge derivative $\langle id, c \rangle$.
 - (c) \mathcal{A} is restricted from querying on $(\text{rkextract}, id, id')$ if \mathcal{A} has previously issued the query $(\text{extract}, id')$ and there exists a challenge derivative $\langle id, c \rangle$.
 - (d) \mathcal{A} is restricted from querying on $(\text{reencrypt}, id, id', c)$ if the query would produce (perhaps implicitly) a challenge derivative $\langle id'', c'' \rangle$ and \mathcal{A} has previously issued the query $(\text{extract}, id'')$.

At the conclusion of this stage, \mathcal{A} outputs i' , where $i' \in \{0, 1\}$.

The outcome of the game is determined as follows: If $i' = i$ then \mathcal{A} wins the game. \mathcal{A} 's advantage in the above game, $Adv_{\mathcal{A}}^{\text{IND-prID-ATK}}$ is defined as $|Pr[i' = i] - 1/2|$. For $\text{ATK} \in (\text{CPA}, \text{CCA})$ we say that the Identity-Based Proxy Re-encryption scheme \mathcal{S} is IND-prID-ATK-secure if for all probabilistic polynomial time algorithms \mathcal{A} , $Adv_{\mathcal{A}}^{\text{IND-prID-ATK}} \leq \nu(k)$.

Bidirectional and PKG-based IBE Proxy Re-encryption. To underscore the importance of non-interactive unidirectional proxy re-encryption, we note that it is possible to construct a bidirectional or PKG-based scheme from any unidirectional scheme. In the bidirectional case, RKGen_{bi} is implemented via two separate calls to RKGen : derive $rk_{id_1 \rightarrow id_2}$ and $rk_{id_2 \rightarrow id_1}$, which are together functionally equivalent to $rk_{id_1 \leftrightarrow id_2}$. Similarly, RKGen_{pkg} can be implemented by deriving secret keys (sk_{id_1}, sk_{id_2}) at the PKG using the KeyGen routine and subsequently calling RKGen as in the bidirectional case. We leave further discussion and non-generic constructions of these alternative forms of proxy re-encryption for future work.

4 Non-interactive Unidirectional Proxy Re-encryption Schemes

The first schemes we present are based on Boneh and Franklin’s IBE scheme [6], and are secure under the Decisional Bilinear Diffie-Hellman Assumption (DBDH) in the random oracle model. While ciphertexts in the proposed schemes have a different form from those in the standard Boneh-Franklin scheme, the master parameters and secret keys remain unchanged. As a result, it is possible to implement proxy re-encryption within an existing Boneh-Franklin deployment (*i.e.*, using pre-existing parameters and keys).

4.1 A First Attempt (IBP1)

Consider a bilinear map $e : \mathbb{G} \times \mathbb{G} \rightarrow \mathbb{G}_T$, where $\mathbb{G} = \langle g \rangle$. Let \mathcal{H}_1 and \mathcal{H}_2 be two independent hash functions¹ such that: $\mathcal{H}_1 : \{0, 1\}^* \rightarrow \mathbb{G}$ and $\mathcal{H}_2 : \mathbb{G}_T \rightarrow \mathbb{G}$. Finally, let s and g^s be the master secret and public key of the PKG, respectively. For some $r \in_R \mathbb{Z}_q^*$, an encryption of $m \in \mathbb{G}_T$ under Alice’s identity can be computed as:

$$IBE_{Alice}(m) = (g^r, m \cdot e(g^s, \mathcal{H}_1(Alice)))^r$$

Suppose Alice wants to delegate her decryption rights to Bob. She must generate a re-encryption key to give to the proxy. Let $IBE_{Bob}(\cdot)$ be a standard identity-based encryption under Bob’s identity. Alice selects a random $X \xleftarrow{\$} \mathbb{G}_T$ and generates the re-encryption key as:

$$rk_{Alice \rightarrow Bob} = \mathcal{H}_1(Alice)^{-s} \cdot \text{fulldomainhash}_2(X), IBE_{Bob}(X),$$

Given an encryption for Alice, $IBE_{Alice}(m) = (c_1, c_2)$ the proxy can transform it into an encryption for Bob by releasing: $(c'_1 = c_1, c'_2 = c_2 \cdot e(g^r, rk_{Alice \rightarrow Bob}), c'_3 = IBE_{Bob}(X))$. Indeed, notice that:

$$\begin{aligned} c'_1 &= g^r \\ c'_2 &= m \cdot e(g^r, \mathcal{H}_2(X)), \\ c'_3 &= IBE_{Bob}(X). \end{aligned}$$

Bob can recover X from c'_3 and then m by computing $c'_2/e(c'_1, \mathcal{H}_2(X))$.

In practice, the scheme presented above can be seen as a variant of the efficient Dodis/Ivan [5] key-splitting approach applied to settings where the decryption process makes use of a bilinear map. Note that (1) The scheme is *unidirectional* since the key $rk_{Alice \rightarrow Bob}$ can be used to convert ciphertexts from Alice to Bob but not vice versa. (2) It is *non-interactive* since Bob is not involved during

¹ Both $\mathcal{H}_1(\cdot)$ and $\mathcal{H}_2(\cdot)$ are more properly “hash-and-encode” functions (see Boneh-Franklin [6] for a detailed definition). Each function consist of a standard hash function which maps inputs to elements of the finite field of order q and then uses an admissible encoding function, `MapToPoint`, to map those elements into points in \mathbb{G} .

the generation of the re-encryption key. (3) It provides *non-transitivity* since the proxy is not allowed to create new re-encryption keys from the existing ones. (4) Finally, we observe that the scheme is *multi-use* since the proxy can re-encrypt the result of a re-encryption and do it multiple times. To see this, consider the re-encryption ciphertext above: (c'_1, c'_2, c'_3) . Notice that c'_3 is just a standard IBE encryption for Bob! A proxy equipped with a re-encryption key $rk_{Bob \rightarrow Charles}$ could just apply the re-encryption algorithm *recursively* to c'_3 and allow Charles to recover X which in turn allows him to recover the original message m .

Scheme Description. We now provide a formal description of the scheme (IBP1).

- **Setup.** Let $e : \mathbb{G} \times \mathbb{G} \rightarrow \mathbb{G}_T$ be a bilinear map, where $\mathbb{G} = \langle g \rangle$ and \mathbb{G}_T have order q . Let $\mathcal{H}_1, \mathcal{H}_2$ be independent full-domain hash functions $\mathcal{H}_1 : \{0, 1\}^* \rightarrow \mathbb{G}$ and $\mathcal{H}_2 : \mathbb{G}_T \rightarrow \mathbb{G}$. To generate the scheme parameters, select $s \xleftarrow{\$} \mathbb{Z}_q^*$, and output $\text{params} = (\mathbb{G}, \mathcal{H}_1, \mathcal{H}_2, g, g^s)$, $\text{msk} = s$.
- **KeyGen**(params, msk, id). To extract a decryption key for identity $id \in \{0, 1\}^*$, return $sk_{id} = \mathcal{H}_1(id)^s$.
- **Encrypt**(params, id, m). To encrypt m under identity id , select $r \xleftarrow{\$} \mathbb{Z}_q^*$ and output $c_{id} = (g^r, m \cdot e(g^s, \mathcal{H}_1(id))^r)$.
- **RKGen**(params, sk_{id_1}, id_2). Select $X \xleftarrow{\$} \mathbb{G}_T$ and compute $\langle R_1, R_2 \rangle = \text{Encrypt}(\text{params}, id_2, X)$. Return $rk_{id_1 \rightarrow id_2} = \langle R_1, R_2, sk_{id_1}^{-1} \cdot \mathcal{H}_2(X) \rangle$.
- **Reencrypt**(params, $rk_{id_1 \rightarrow id_2}, c_{id_1}$). To re-encrypt a level- ℓ ciphertext from id_1 to id_2 , first parse c_{id_1} as $(C_1, \dots, C_{2\ell})$ and $rk_{id_1 \rightarrow id_2}$ as (R_1, R_2, R_3) . Next:
 1. If $\ell = 1$, output $c_{id_2} = \langle C_1, C_2 \cdot e(C_1, R_3), R_1, R_2 \rangle$.
 2. If $\ell > 1$, treat the elements $\langle C_{2\ell-1}, C_{2\ell} \rangle$ as a first-level ciphertext δ . Compute $\langle C'_1, C'_2, C'_3, C'_4 \rangle = \text{Reencrypt}(rk_{id_1 \rightarrow id_2}, \delta)$. Output the ciphertext $c_{id_2} = \langle C_1, \dots, C_{2\ell-2}, C'_1, C'_2, C'_3, C'_4 \rangle$.
- **Decrypt**(params, sk_{id}, c_{id}). Parse the level- ℓ ciphertext c_{id} as $(C_1, \dots, C_{2\ell})$. Next:
 1. If $\ell = 1$ output $m = C_2/e(C_1, sk_{id})$.
 2. If $\ell > 1$, treat the pair $\langle C_{2\ell-1}, C_{2\ell} \rangle$ as a first-level ciphertext c'_{id} , and compute $X_\ell = \text{Decrypt}(sk_{id}, c'_{id})$. For $i = (\ell-1)$ descending to 1, compute $X_i = C_{2i}/e(C_{2i-1}, \mathcal{H}_2(X_{i+1}))$. Finally, output X_1 as the plaintext.

Each level- ℓ ciphertext in the above scheme contain 2ℓ elements. In principle, the scheme permits an arbitrary number of re-encryptions on a ciphertext, with a two-element ciphertext expansion on each re-encryption.

Correctness. We first show correctness for first-level ciphertexts (*i.e.*, those produced by **Encrypt**). Let $c_{id_1} = (g^r, m \cdot e(g^s, \mathcal{H}_1(id_1))^r)$ be the first-level encryption of m under id_1 , and $sk_1 = \mathcal{H}_1(id_1)^s$ be the corresponding decryption key. The decryption process produces the following result:

$$(m \cdot e(g^s, \mathcal{H}_1(id_1))^r)/e(g^r, \mathcal{H}_1(id_1)^s) = m$$

The correctness under re-encryption is shown as follows. Given a first-level ciphertext $c_{id_1} = (g^r, C_2)$ and a correctly-formed re-encryption key $rk_{id_1 \rightarrow id_2} = ((R_1, R_2) = \text{Encrypt}(\text{params}, id_2, X), R_3)$, we obtain the “second-level” ciphertext $c_{id_2} = (g^r, C'_2 = C_2 \cdot e(g^r, R_3), R_1, R_2)$ where C'_2 is:

$$\begin{aligned} C'_2 &= C_2 \cdot e(g^r, R_3) \\ &= m \cdot e(g^s, \mathcal{H}_1(id_1))^r \cdot e(g^r, \mathcal{H}_1(id_1))^{-s} \cdot \mathcal{H}_2(X) \\ &= m \cdot e(g, \mathcal{H}_2(X))^r \end{aligned}$$

Given $sk_{id_2} = \mathcal{H}_1(id_2)^s$ we decrypt $c_{id_2} = (g^r, C'_2, R_1, R_2)$ as follows. Begin by decrypting the first-level ciphertext $\hat{c}_{id_2} = (R_1, R_2)$ under sk_{id_2} : $X = \text{Decrypt}(\text{params}, sk_{id_2}, \hat{c}_{id_2})$. Then compute $C'_2/e(g^r, \mathcal{H}_2(X))$ to obtain m . Having shown correctness for a single re-encryption, the correctness for multiple re-encryptions follows. Given level- ℓ ciphertext c_{id_i} and sk_{id_i} , strip the the final two elements and treat them as a first-level ciphertext under id_i , decrypting to reveal X_ℓ . Use the value X_ℓ as a decryption secret for the previous two elements, and repeat until the final two elements remain. The final value in this chain contains the original message m .

Security. We next show that IBP1 scheme defined above meets the IND-prID-CPA definition if the Decisional Bilinear Diffie-Hellman assumption holds in $(\mathbb{G}, \mathbb{G}_T)$. Our proof is in the random oracle model, and is an extension of the original proof of Boneh/Franklin [6].

Theorem 1. *If there exists a p.p.t. adversary \mathcal{A} that wins the IND-prID-CPA game on IBP1 with non-negligible advantage, then there exists an adversary \mathcal{B} that solves the DBDH problem over \mathbb{G}, \mathbb{G}_T with non-negligible advantage.*

A proof sketch of Theorem 1 is presented in Appendix A.

4.2 An Optimization

Ciphertexts in scheme section 4.1 expand upon re-encryption. This is caused by the inclusion within the re-encrypted ciphertext of a portion of the re-encryption key. There are scenarios where Bob knows that the original ciphertext was intended for Alice (this information can even be appended to the ciphertext) and there is no need for multiple re-encryptions. In such cases we can simplify our construction by using a result of Sakai *et. al.* [19]. Specifically, in the Boneh-Franklin IBE symmetric setting, Alice and Bob inherently share a secret key $K_{AB} = e(\mathcal{H}_1(\text{Alice}), \mathcal{H}_1(\text{Bob}))^s$. Alice can use this value to compute the re-encryption key as follows:

$$rk_{\text{Alice} \rightarrow \text{Bob}} = \mathcal{H}_1(\text{Alice})^{-s} \cdot \mathcal{H}_3(\{K_{AB}\} || \text{Alice} \rightarrow \text{Bob}).$$

Where $\{K_{AB}\}$ denotes binary representation, and $\mathcal{H}_3 : \{0, 1\}^* \rightarrow \mathbb{G}$ is an independent full domain hash function. The string “*Alice* \rightarrow *Bob*” is added to ensure that a re-encryption key from Bob to Alice is computed under a distinct secret (bidirectional re-encryption). Note that the resulting scheme permits only a single re-encryption for each ciphertext. A primary advantage of this construction is the absence of ciphertext expansion during re-encryption.

4.3 A Chosen Ciphertext Secure Scheme (IBP2)

The scheme presented above is secure under chosen plaintext attack. While this is the level of security provided by many IBE and proxy re-encryption schemes (e.g., [3,18] and the practical proxy encryption constructions of Dodis/Ivan [5]), it is important to consider stronger definitions such as security under adaptive chosen ciphertext attack.

Background. A common approach to building CCA-secure Identity-Based *Encryption* schemes in the random oracle model is to begin with a CPA-secure construction, and then apply the generic Fujisaki-Okamoto conversion [20] (see e.g., [6,21]). It is tempting to believe that this approach is by itself sufficient to construct CCA-secure IB-PRE schemes. Unfortunately, this does not appear to be the case. Notice that a re-encryption proxy grants adversaries an alternative means by which adversaries may decrypt ciphertexts: a malicious delegatee B may decrypt A 's ciphertexts by first using the proxy to re-encrypt from $id_A \rightarrow id_B$, and then decrypting the result under his own secret key. When a malicious delegatee uses the proxy to “alternatively decrypt” in this manner, he need not follow the specified F-O decryption algorithm, and can ignore the critical ciphertext validity checks. Unfortunately, the validity checks of the F-O approach cannot be moved into the re-encryption process, as they fundamentally require access to the decryption secret.

Intuition. In order to surmount the issues raised above, we propose an approach that provides the proxy with the means to verify ciphertext validity and reject improperly-formed ciphertexts. As a result of this check, a malicious delegatee no longer gains any advantage by using the re-encryption proxy as an oracle. The building block of our construction is a Hierarchical Identity-Based Proxy Re-encryption scheme, which we implement using a modified form of the Gentry-Silverberg HIBE [22] (this scheme is in turn based on the Boneh/Franklin scheme). To achieve IND-prID-CCA-secure IB-PRE, we make use of the Canetti, Halevi and Katz (CHK) [11] technique, which allows us to transform a HIBE into a CCA-secure IBE scheme with a type of *publicly-verifiable* ciphertext validity check. In order to present a more efficient construction, we re-use randomness and implement the CHK transform using a Boneh/Lynn/Shacham short signature [23].

The Construction. We now present a *single-use*, non-interactive CCA-secure IB-PRE construction (IBP2).

- **Setup.** Let $n(\cdot)$ be a polynomial function of the security parameter k . Let $e : \mathbb{G} \times \mathbb{G} \rightarrow \mathbb{G}_T$ be a bilinear map, where \mathbb{G}, \mathbb{G}_T have order q and $\mathbb{G} = \langle g \rangle$. To generate the scheme parameters, select $s \xleftarrow{\$} \mathbb{Z}_q^*$ and output $\text{params} = (\mathcal{H}_1, \mathcal{H}_2, \mathcal{H}_3, \mathcal{H}_4, \mathcal{H}_5, g, g^s)$, $\text{msk} = s$, with independent hash functions \mathcal{H}_{1-6} defined as below:

$$\begin{aligned} \mathcal{H}_1 : \{0, 1\}^* &\rightarrow \mathbb{G}, \mathcal{H}_2 : \{0, 1\}^* \rightarrow \mathbb{G} \\ \mathcal{H}_3 : \{0, 1\}^* &\rightarrow \mathbb{G}, \mathcal{H}_4 : \mathbb{G}_T \times \{0, 1\}^n \rightarrow \mathbb{Z}_q^* \\ \mathcal{H}_5 : \mathbb{G}_T &\rightarrow \{0, 1\}^n \end{aligned}$$

- **KeyGen**(params, msk, id). To extract a decryption key for identity $id \in \{0, 1\}^*$, return $sk_{id} = \mathcal{H}_1(id)^s$.
- **Encrypt**(params, $id, m \in \{0, 1\}^n$). To encrypt m under identity $id \in \{0, 1\}^*$, first:
 1. Select $\sigma \xleftarrow{\$} \mathbb{G}_T$, and set $r = \mathcal{H}_4(\sigma, m)$.
 2. Compute $\langle A, B, C \rangle = (g^r, \sigma \cdot e(g^s, \mathcal{H}_1(id))^r), m \oplus \mathcal{H}_5(\sigma)$.
 3. Compute $S = \mathcal{H}_3(id || \langle A, B, C \rangle)^r$.
 4. Output the ciphertext $c = \langle S, A, B, C \rangle$.
- **RKGen**(params, sk_{id_1}, id_1, id_2). To compute a re-encryption key from $id_1 \rightarrow id_2$:
 1. Select $N \xleftarrow{\$} \{0, 1\}^{n(k)}$, and compute $K = e(sk_{id_1}, \mathcal{H}_1(id_2))$.
 2. Output $rk_{id_1 \rightarrow id_2} = \langle N, \mathcal{H}_2(K || id_1 || id_2 || N) \cdot sk_{id_1} \rangle$.
- **Reencrypt**(params, $rk_{id_1 \rightarrow id_2}, c_{id_1}$). To re-encrypt a first-level ciphertext, first parse c_{id_1} as $\langle S, A, B, C \rangle$, and parse $rk_{id_1 \rightarrow id_2}$ as $\langle N, R \rangle$. Next:
 1. Let $h = \mathcal{H}_3(id_1 || \langle A, B, C \rangle)$.
 2. Check if $e(g, S) = e(h, A)$. If not, return \perp .
 3. Otherwise, select $t \xleftarrow{\$} \mathbb{Z}_q^*$ and compute $B' = B / \frac{e(A, R \cdot h^t)}{e(g^t, S)}$.
 4. Output the re-encrypted ciphertext $c_{id_2} = \langle A, B', C, id_1, N \rangle$.
- **Decrypt**(params, sk_{id}, c_{id}). To decrypt a first-level (non re-encrypted) ciphertext, first parse c_{id} as $\langle S, A, B, C \rangle$. Next:
 1. Let $h = \mathcal{H}_3(id, \langle A, B, C \rangle)$.
 2. Select $t \xleftarrow{\$} \mathbb{Z}_q^*$, and compute $\sigma' = B / \frac{e(A, sk_{id} \cdot h^t)}{e(g^t, S)}$.
 3. Compute $m' = C \oplus \mathcal{H}_5(\sigma')$, and $r' = \mathcal{H}_4(\sigma', m')$.
 4. Verify that $S = h^{r'}$ and $A = g^{r'}$. If either check fails, return \perp , otherwise output m' .

To decrypt a second-level (re-encrypted) ciphertext, first parse c_{id} as $\langle A, B, C, id_{src}, N \rangle$. Next:

1. Compute $K = e(\mathcal{H}_1(id_{src}), sk_{id})$.
2. Compute $\sigma' = B \cdot e(A, \mathcal{H}_2(K || id_{src} || id || N))$.
3. Compute $m' = C \oplus \mathcal{H}_5(\sigma')$, and $r' = \mathcal{H}_4(\sigma', m')$.
4. Verify that $A = g^{r'}$. If this check fails, return \perp , otherwise output m' .

Correctness. We begin by showing correctness for first-level ciphertexts (*i.e.*, those produced by **Encrypt**). Let $c_{id_1} = \langle S, A, B, C \rangle$ be the first-level encryption of m under id_1 , with $h = \mathcal{H}_3(id_1 || \langle A, B, C \rangle)$.

$$c_{id_1} = \langle h^r, g^r, \sigma \cdot e(g^s, \mathcal{H}_1(id)^r), m \oplus \mathcal{H}_5(\sigma) \rangle$$

Let $sk_1 = \mathcal{H}_1(id_1)^s$ be the corresponding decryption key. For a random $t \in \mathbb{Z}_q^*$, the decryption process proceeds as follows:

$$\begin{aligned} (\sigma \cdot e(g^s, \mathcal{H}_1(id_1))^r) / \frac{e(g^r, \mathcal{H}_1(id_1)^s \cdot h^t)}{e(g^t, h^r)} &= \sigma \\ \mathcal{H}_5(\sigma) \oplus (m \oplus \mathcal{H}_5(\sigma)) &= m \\ g^{\mathcal{H}_4(\sigma, m)} &\stackrel{?}{=} g^r \\ h^{\mathcal{H}_4(\sigma, m)} &\stackrel{?}{=} h^r \end{aligned}$$

The correctness of re-encryption is shown as follows. Given the first-level ciphertext c_{id_1} presented above, and a correctly-formed re-encryption key $rk_{id_1 \rightarrow id_2} = \langle N, R \rangle$, the re-encryption process begins with a ciphertext validity check:

$$e(g, h^r) \stackrel{?}{=} e(h, g^r)$$

Recall that $R = sk_{id_1} \cdot W$ where $W = \mathcal{H}_2(e(\mathcal{H}_1(id_1)^s, \mathcal{H}_1(id_2)) || id_1 || id_2 || N)$. To generate the “second-level” ciphertext $c_{id_2} = (A, B', C, id_1, N)$, we choose $t \in_R \mathbb{Z}_q^*$ and obtain:

$$B' = (\sigma \cdot e(g^s, \mathcal{H}_1(id_1))^r) / \frac{e(g^r, R \cdot h^t)}{e(g^t, h^r)} = \sigma / e(g^r, W)$$

Finally, we decrypt the re-encrypted ciphertext $c_{id_2} = (A, B', C, id_1, N) = (g^r, \sigma / e(g^r, W), m \oplus \mathcal{H}_5(\sigma), id_1, N)$ as follows. Given $sk_{id_2} = \mathcal{H}_1(id_2)^s$:

$$\begin{aligned} \mathcal{H}_2(e(\mathcal{H}_1(id_1), \mathcal{H}_1(id_2)^s) || id_1 || id_2 || N) &= W \\ (\sigma / e(g^r, W)) \cdot e(g^r, W) &= \sigma \\ (m \oplus \mathcal{H}_5(\sigma)) \oplus \mathcal{H}_5(\sigma) &= m \\ g^{\mathcal{H}_4(\sigma, m)} &\stackrel{?}{=} g^r \end{aligned}$$

Security. We claim that IBP2 meets the IND-prID-CCA definition if the Decisional Bilinear Diffie-Hellman assumption holds in $(\mathbb{G}, \mathbb{G}_T)$. Our proof is in the random oracle model. Due to space limitations we are unable to include the proof here. However, it can be found in the full version of this paper [24].

5 Applications of Identity-Based Proxy Re-encryption

Proxy Re-encryption has a number of practical applications, which have been detailed in previous works. All of these applications translate directly to the Identity-Based setting but with some additional features.

Secure Email with IBE. The most natural application of proxy re-encryption is to allow Bob to read Alice’s encrypted emails while she is on vacation. Messages are encrypted under the email address “*alice@company.com*” and are translated by the proxy into encryptions under “*bob@company.com*”. The proxy does not learn the content of the messages being translated.

Attribute-based Delegations. As noted by Boneh and Franklin [6], identities can be created to include attributes or to express conditions. For instance, a message encrypted under “*alice || lawyer || from 01/01/2008*” can be read by Alice only if she is a lawyer and not before the beginning of year 2008. This idea applies directly to our IBE-PRE scheme and it allows Alice to specify under which conditions the proxy is allowed to translate her ciphertexts into Bob’s. For instance, consider the case of *temporary* delegations [3] where the time is

divided in time intervals t_1, t_2, \dots, t_k and Alice can specify that the proxy can translate her ciphertexts for Bob only during t_i . With our scheme, Alice could just create the proxy key:

$$rk_{Alice||t_i \rightarrow Bob},$$

so that any encryption under $Alice || t_i$ can be converted into an encryption for Bob but not during other time periods. This eliminates the need for designing a separate and specialized scheme as it was done in [3].

Even more interestingly, Alice could specify the conditions under which Bob can read her messages. For instance, a re-encryption key of this form:

$$rk_{Alice \rightarrow Bob || after \text{ Nov 2007} || security-clearance},$$

would specify that encryptions under Alice's identity can be converted into encryptions for Bob but that Bob can read the messages only in the future, after Nov 2007, and under the condition that he is able to obtain a security clearance.

Bridging IBE and PKE. *Hybrid* proxy re-encryption is a concept put forward by Boneh, Goh and Matsuo [16] to create a bridge between IBE and public-key based encryption (PKE). Our scheme can also be used to translate from IBE to PKE. Indeed, consider the ciphertext after the re-encryption, which has the form:

$$c'_1 = g^r, \quad c'_2 = m \cdot e(g^r, \mathcal{H}_2(K)), \quad c'_3 = IBE_{Bob}(K).$$

Notice that c'_3 is a standard (semantically-secure) id-based encryption of a key K . This encryption can be substituted with a public-key based one (or even a semantically-secure symmetric one). In this way, an encryption under Alice's identity is converted into an encryption under Bob's public-key. Our approach provides some advantages over the one in [16]. Indeed, no TTP is involved in creating re-encryption keys and parameters in our scheme are compatible with those of the standard Boneh-Franklin IBE.

Travel Key. Boneh and Franklin [6] suggested to use an IBE system to store temporary keys into the laptop during travel so that, if the laptop is lost or stolen, only those keys get exposed. The idea is to let Bob act as a PKG that generates his own master secret and public keys. Alice could use Bob's master public-key to encrypt messages for Bob under identities *day1*, *day2*, ... etc., for all days in which Bob is traveling. Bob can store into his laptop just the keys corresponding to those days while leaving his master secret key safely stored elsewhere.

This solution, however, requires Bob to inform of his travels any of his potential correspondents and have them act according to the encryption scheme (that is, they have to encrypt under *day1*, *day2*, etc.). An alternative solution is to set up a proxy (Bob's mail server, for instance) with a re-encryption key of the form:

$$rk_{Bob \rightarrow Bob's-Travel-Key}.$$

Every encryption intended for Bob will be encrypted under Bob's travel key, which is the only secret key stored into his laptop. Notice that the proxy does

not have to be trusted and can be set-up by a system administrator who won't be able to read Bob's messages.

Access Control in Networked File Storage. In [3], the authors describe an application of proxy re-encryption to the distribution of key material within a cryptographic filesystem. Each file stored on an untrusted file server is encrypted using a symmetric key; these keys are encrypted under a public *master key* which is stored alongside the encrypted material. When a user wishes to decrypt a file, the semi-trusted *keyserver* re-encrypts these encapsulated symmetric keys from the master key to the keys of individual users who can then decrypt. The key server provides access control for the encrypted material, but does not itself possess the ability to decrypt files.

This application translates naturally to the Identity Based setting with the additional benefit of allowing the holder of the master key to specify access control policies directly within the identity strings of the users. A re-encryption key can even be generated before an individual has joined the system.

6 Conclusions and Future Work

In this work we introduced new constructions enabling non-interactive, unidirectional proxy re-encryption in the IBE setting. Our schemes are very efficient and can be deployed within standard IBE frameworks. New compelling applications can be realized thanks to our schemes, most notably attribute-based delegation and access control.

An interesting open problem is to find efficient constructions for *multi-use* CCA-secure IBE-PRE schemes. Another important open problem is to find efficient IBE-PRE secure in the standard model (rather than in the RO model).

Acknowledgments. We would like to thank Susan Hohenberger for helpful discussion and suggestions.

References

1. Mambo, M., Okamoto, E.: Proxy Cryptosystems: Delegation of the Power to Decrypt Ciphertexts. IEICE Trans. Fund. Electronics Communications and Computer Science **E80-A/1** (1997) 54–63
2. Blaze, M., Bleumer, G., Strauss, M.: Divertible protocols and atomic proxy cryptography. In: Proceedings of Eurocrypt '98. Volume 1403. (1998) 127–144
3. Ateniese, G., Fu, K., Green, M., Hohenberger, S.: Improved Proxy Re-encryption Schemes with Applications to Secure Distributed Storage. In: the 12th Annual Network and Distributed System Security Symposium. (2005) 29–43 Full version available at <http://eprint.iacr.org/2005/028>.
4. Jakobsson, M.: On quorum controlled asymmetric proxy re-encryption. In: Proceedings of Public Key Cryptography. (1999) 112–121
5. Dodis, Y., Ivan, A.: Proxy cryptography revisited. In: Proceedings of the Tenth Network and Distributed System Security Symposium. (2003)

6. Boneh, D., Franklin, M.K.: Identity-based encryption from the Weil Pairing. In: Advances in Cryptology (CRYPTO 2001). Volume 2139 of Lecture Notes in Computer Science., Springer (2001) 213–229
7. Cocks, C.: An identity based encryption scheme based on Quadratic Residues. In: Cryptography and Coding, 8th IMA International Conference. Volume 2260 of Lecture Notes in Computer Science., Springer (2001) 360–363
8. Balfanz, D., Durfee, G., Shankar, N., Smetters, D., Staddon, J., Wong, H.C.: Secret handshakes from pairing-based key agreements. In: Proceedings of the 2003 IEEE Symposium on Security and Privacy (SP '03), Washington, DC, USA, IEEE Computer Society (2003) 180
9. Boneh, D., Crescenzo, G.D., Ostrovsky, R., Persiano, G.: Public key encryption with keyword search. In: Advances in Cryptology - EUROCRYPT 2004. Volume 3027 of Lecture Notes in Computer Science. (2004) 506–522
10. Waters, B.R., Balfanz, D., Durfee, G., Smetters, D.K.: Building an encrypted and searchable audit log. In: Proceedings of Network and Distributed System Security Symposium 2004 (NDSS'04), San Diego, CA (2004)
11. Canetti, R., Halevi, S., Katz, J.: Chosen-ciphertext security from Identity Based Encryption. In: Proceedings of Eurocrypt '04. Volume 3027 of Lecture Notes in Computer Science., Springer-Verlag (2004) 207–222
12. Boneh, D., Franklin, M.K.: Identity-based encryption from the Weil Pairing. SIAM Journal of Computing **32**(3) (2003) 586–615
13. Voltage Security, Inc. (<http://www.voltage.com>)
14. Zhou, L., Marsh, M.A., Schneider, F.B., Redz, A.: Distributed blinding for ElGamal re-encryption. Technical Report 2004–1924, Cornell Computer Science Department (2004)
15. Canetti, R., Hohenberger, S.: Chosen-ciphertext secure proxy re-encryption. (*Unpublished manuscript*)
16. Boneh, D., Goh, E.J., Matsuo, T.: Proposal for P1363.3 Proxy Re-encryption. (<http://grouper.ieee.org/groups/1363/IBC/submissions/NTTDataProposal-for-P1363.3-2006-09-01.pdf>)
17. Boneh, D., Boyen, X.: Efficient selective-id secure Identity-Based Encryption without random oracles. In: Proceedings of Eurocrypt '04. Volume 3027 of Lecture Notes in Computer Science., Springer (2004) 223–238
18. Waters, B.: Efficient Identity-Based Encryption without random oracles. In: Proceedings of Eurocrypt '05. Volume 3494 of Lecture Notes in Computer Science., Springer (2005) 114–127
19. Sakai, R., Ohgishi, K., Kasahara, M.: Cryptosystems based on pairing. In: Symposium on Cryptography and Information Security (SCIS 2000). (2000)
20. Fujisaki, E., Okamoto, T.: Secure integration of asymmetric and symmetric encryption schemes. In: Proceedings of Crypto '99. Volume 1666 of Lecture Notes in Computer Science., Springer (1999) 537–554
21. Yang, P., Kitagawa, T., Hanaoka, G., Zhang, R., Matsuura, K., Imai, H.: Applying Fujisaki-Okamoto to Identity-Based Encryption. In: Applied Algebra, Algebraic Algorithms and Error-Correcting Codes, 16th International Symposium (AAECC-16). Volume 3857 of Lecture Notes in Computer Science., Springer (2006) 183–192
22. Gentry, C., Silverberg, A.: Hierarchical ID-Based cryptography. In: Proceedings of Asiacrypt '02. Volume 2501 of Lecture Notes in Computer Science., Springer (2002) 548–566

23. Boneh, D., Lynn, B., Shacham, H.: Short signatures from the Weil Pairing. In: ASIACRYPT '01: Proceedings of the 7th International Conference on the Theory and Application of Cryptology and Information Security. Volume 2248 of Lecture Notes in Computer Science., London, UK, Springer-Verlag (2001) 514–532
24. Green, M., Ateniese, G.: Identity-based proxy re-encryption. Cryptology ePrint Archive, Report 2006/473 (2006) <http://eprint.iacr.org/2006/473.pdf>.

A Security Proof of IBP1

Proof sketch. Let \mathcal{A} be a p.p.t. algorithm that has non-negligible advantage ϵ in $\text{Exp}^{\mathcal{A}, \text{IND-prID-CPA}}$ against IBP1. We use \mathcal{A} in order to construct a second algorithm \mathcal{B} which has non-negligible advantage at solving the DBDH problem in \mathbb{G}, \mathbb{G}_T . Algorithm \mathcal{B} accepts as input an appropriately-distributed tuple $\langle \mathbb{G} = \langle g \rangle, g^a, g^b, g^c, T \rangle \in \mathbb{G}^4 \times \mathbb{G}_T$ and outputs 1 if $T = e(g, g)^{abc}$. We now describe the algorithm \mathcal{B} , which interacts with algorithm \mathcal{A} via the IND-prID-CPA interface.

Oracle Queries. \mathcal{B} simulates the random oracle $\mathcal{H}_1 : \{0, 1\}^* \rightarrow \mathbb{G}$ as follows: On receipt of a query for id (on which it has not previously been queried), select $z \xleftarrow{\$} \mathbb{Z}_q^*$ and randomly flip a weighted coin to set $\alpha \leftarrow 1$ with probability γ (see below), and $\alpha \leftarrow 0$ otherwise. If $\alpha = 0$ then compute $h \leftarrow (g^c)^z$, else compute $h \leftarrow g^z$. Record the tuple (id, h, z, α) . Finally, return h as the result of the query (if id has previously been queried, simply locate the existing tuple and return the previously-computed h). Note that the distribution of the values h returned by the simulated oracle is random, regardless of the choice of α . \mathcal{B} simulates (initial) queries to the random oracle $\mathcal{H}_2 : \mathbb{G}_T \rightarrow \mathbb{G}$ by simply returning elements $\in_R \mathbb{G}$.

Our simulation proceeds as follows:

1. **SETUP.** \mathcal{B} generates the scheme’s master parameters $\text{params} = (\mathbb{G}, \mathcal{H}_1, \mathcal{H}_2, g, g^a)$ and gives this tuple to \mathcal{A} .
2. **FIND.** When \mathcal{A} submits $(\text{extract}, id)$, \mathcal{B} evaluates $\mathcal{H}(id)$ as described above, to obtain (id, h, z, α) . \mathcal{B} outputs $sk_{id} = (g^a)^z$ to \mathcal{A} .

When \mathcal{A} submits $(\text{rkextract}, id_1, id_2)$, \mathcal{B} selects $r \xleftarrow{\$} \mathbb{Z}_q^*$, $x \xleftarrow{\$} \mathbb{G}$ and $X \xleftarrow{\$} \mathbb{G}_T$, then evaluates $\mathcal{H}_1(id_1)$ and $\mathcal{H}_1(id_2)$ to obtain the values $(\alpha_1, z_1), (\alpha_2, z_2)$ (for id_1, id_2 respectively). Now:

- (a) If $\alpha_1 = 0$ then \mathcal{B} returns $rk_{id_1 \rightarrow id_2} = ((g^b)^r, T^{rz_2} \cdot X, x)$ to \mathcal{A} (note that this key is incorrectly formed, see section below).
- (b) If $\alpha_1 = 1$ then \mathcal{B} returns the correctly-formed tuple $rk_{id_1 \rightarrow id_2} = (g^r, e(g^a, \mathcal{H}_1(id_2)^r) \cdot X, (g^a)^{-z_1} \cdot \mathcal{H}_2(X))$.

3. **CHALLENGE.** At the conclusion of the FIND phase, \mathcal{A} outputs (id^*, m_0, m_1) with the condition that \mathcal{A} ’s choice of id^* is not trivial.² \mathcal{B} selects $i \xleftarrow{\$} \{0, 1\}$,

² We reject a choice of id^* when \mathcal{A} has previously extracted a series of re-encryption keys, and a decryption key sk_{id^*} such that \mathcal{A} can consecutively re-encrypt ciphertexts from id^* to id' .

then evaluates $\mathcal{H}_1(id^*)$ to recover (id^*, h, z, α) from the \mathcal{H}_1 table. \mathcal{B} returns $c^* = \langle g^b, T^z \cdot m_i \rangle$ to \mathcal{A} .

4. GUESS. \mathcal{A} makes queries (`extract, ...`) and (`rkextract, ...`) as in the FIND stage, except that \mathcal{A} is restricted from making any query that would result in a trivial situation (a valid decryption path from id^* to an identity for which the adversary possesses a secret key). At the conclusion of this phase, \mathcal{A} outputs its guess $i' \in \{0, 1\}$.

Conditions for Abort. Let α_i represent the value α generated by $\mathcal{H}_1(id_i)$. Prior to outputting a value, \mathcal{B} verifies several conditions:

- (a) The value α corresponding to id^* is 0.
- (b) For each of \mathcal{A} 's queries (`extract, idi`), $\alpha_i = 1$.
- (c) For each of \mathcal{A} 's queries (`rkextract, idi, idj`), where $id_i \rightarrow id_j$ lies along a path leading from id^* , $\alpha_j = 0$.
- (d) For each of \mathcal{A} 's queries (`rkextract, idi, idj`), where $id_i \rightarrow id_j$ does not lie along a path leading from id^* , $\alpha_i = 1$.

If any of the above conditions are false, \mathcal{B} aborts the simulation. Otherwise, if $i' = i$, \mathcal{B} outputs 1, or 0 otherwise.

Claim. If $\langle g, g^a, g^b, g^c, T \rangle$ is a DBDH tuple and \mathcal{B} does not abort, then \mathcal{A} 's view is identical to the real attack— with the significant exception of re-encryption keys having the form $rk_{id^* \rightarrow \dots}$. We argue below that \mathcal{A} cannot detect these improperly-formed re-encryption keys, and thus cannot distinguish the simulation. Hence, when the input to \mathcal{B} is a DBDH tuple, then the challenge ciphertext c^* is a correct encryption of m_i under id^* and thus from definition of \mathcal{A} and the argument above it holds that $|\Pr[i = i'] - \frac{1}{2}| = \epsilon$ and thus \mathcal{B} outputs 1 with probability $\frac{1}{2}$ plus a non-negligible quantity. When the input to \mathcal{B} is random, c^* represents the encryption of a random element in \mathbb{G}_T and is independent of \mathcal{B} 's choice of i (and therefore $\Pr[i = i'] = \frac{1}{2}$). Thus, \mathcal{B} succeeds in distinguishing DBDH tuples with non-negligible advantage.

Invalid Re-encryption keys. In the simulation above, every re-encryption key that lies along a path from id^* is incorrectly formed. (At the same time, it is easy to see that all other re-encryption keys *are* correctly formed.) Unfortunately, this condition is unavoidable, as the simulator does not possess the knowledge required in order generate a valid re-encryption key from the challenge identity id^* . To complete our proof, therefore, we make a separate argument that no adversary \mathcal{A} can distinguish our simulation from a “real-world” interaction in which all values have the correct form. The heuristic argument for security is simple: each correctly-formed re-encryption key $rk_{id_1 \rightarrow id_2}$ consists of a semantically secure encryption (R_1, R_2) of some element $X \in_R \mathbb{G}_T$, along with the value $R_3 = (\mathcal{H}_1(id_2))^{-s} \cdot \mathcal{H}_2(X) \in \mathbb{G}$. An incorrectly-formed re-encryption key replaces R_3 with some value $x \in_R \mathbb{G}$; this x can naturally be expressed as $(\mathcal{H}_1(id_2))^{-s} \cdot y$ for some unknown $y \in \mathbb{G}$. Intuitively, an adversary who can distinguish malformed re-encryption keys in our simulation must therefore be able to determine that (R_1, R_2) do *not* encrypt some value $Y \in \mathbb{G}_T$ s.t. $\mathcal{H}_2(Y) = y$. We formalize the statement via the following Lemma.

Lemma 1 (Indistinguishability of simulations). If there exists a *p.p.t.* algorithm \mathcal{A}' with non-negligible advantage ϵ' at distinguishing the simulation above from a “correct” simulation (in which all values are correctly-formed), then we can construct an algorithm \mathcal{B}' that solves the DBDH problem in $(\mathbb{G}, \mathbb{G}_T)$ with non-negligible advantage.

Probability of abort. A variety of conditions in the above simulation can lead the simulator to abort. Boneh and Franklin [6] provide a technique for computing the value γ used in simulating the random oracle \mathcal{H}_1 , and for placing bounds on the abort probability. We refer the reader to this discussion, and provide a detailed argument in the full version. \square