Identity-Based Cryptosystem Privacy

Joseph A. Calandrino* and Alfred C. Weaver
University of Virginia
Department of Computer Science
{jac4dt, acw}@cs.virginia.edu

ABSTRACT

Identity-based cryptosystems (IBCs) rely on the use of a private key generator (PKG), which maintains a master secret. This master secret allows the PKG to produce the private key for any identity under its authority. Distrustful groups may wish to maintain distinct PKGs and master secrets to avoid reliance on external entities. IBC privacy is the property that an adversary can gain no more than a negligible advantage in distinguishing between encrypted messages destined for users under multiple distinct cryptosystems. This paper formalizes the notion of IBC privacy and, through modification of key-privacy proofs by Abdalla et al. [1], demonstrates that the Boneh-Franklin identity-based encryption scheme [9] offers the IBC privacy property given agreement on certain system parameters.

1. INTRODUCTION

In an identity-based cryptosystem (IBC), strings representing user identities serve as users' public keys. Unlike standard asymmetric cryptosystems, senders in IBCs have no need to retrieve recipients' public keys. These systems rely on a private key generator (PKG), which maintains a master secret. The master secret determines a master public key and allows the PKG to produce and distribute private keys for any users under its authority. To avoid reliance on external parties, distrustful entities may wish to maintain distinct cryptosystems, including PKGs and master secrets. In spite of any distrust, entities may have a common interest in obscuring which PKG can generate a private key to decrypt any given ciphertext. IBC privacy is the property that an adversary can gain no more than a negligible advantage in distinguishing between encrypted messages destined for users under multiple distinct cryptosystems. This paper formally defines IBC privacy and modifies proofs in [1] to show that, given agreement certain system parameters, the Boneh-Franklin identity-based cryptosystem [9] offers this property.

1.1 Motivation

Numerous scenarios exist in which entities may desire distinct, internally maintained IBCs. Suppose that Microsoft and IBM each use identity-based encryption for internal or external-to-internal communications. Each company is likely unwilling to give a third party the ability to distribute keys and, consequently, the power to decipher all company communications. One option is use of a distributed master secret with each company maintaining a portion of the secret to preserve some control [9]. Unfortunately, the companies would be dependent on each other, and an attack on one could disrupt the other's communications. The use of separate cryptosystems is therefore an attractive alternative.

Under such scenarios, IBC privacy offers substantial benefits for participating entities. Adversaries unrelated to those entities lose the ability to derive definitive knowledge of an entity's communications based transmitted ciphertext alone. If Microsoft's traffic doubles in the earlier example, adversaries outside of Microsoft and IBM would be unable to verify that Microsoft alone is responsible for the increase in overall traffic. Large coalitions of unrelated entities using IBCs with IBC privacy may be able to limit adversaries' abilities to perform even coarse analysis. Suppose government standards mandate the use of a given cryptosystem with IBC privacy in all government

* This research was performed while on appointment as a U.S. Department of Homeland Security (DHS) Fellow under the DHS Scholarship and Fellowship Program, a program administered by the Oak Ridge Institute for Science and Education (ORISE) for DHS through an interagency agreement with the U.S Department of Energy (DOE). ORISE is managed by Oak Ridge Associated Universities under DOE contract number DE-AC05-00OR22750. All opinions expressed in this paper are the author's and do not necessarily reflect the policies and views of DHS, DOE, or ORISE.

systems. Traffic from the National Institute of Health, the Internal Revenue Service, the Department of Education, etc. is unlikely to be related. The randomness of unrelated traffic from many sources can have a leveling effect that diminishes an adversary's ability to perceive peaks and valleys in entity communications.

1.2 Paper Overview

The remainder of this paper is organized as follows. Section 2 describes work related to the topic of IBC privacy, including a relevant proof by Abdalla et al. [1] of key-privacy in the Boneh-Franklin IBC. Section 3 formally defines IBC privacy. Using results from [1], section 4 details sufficient conditions for IBC privacy to exist. Section 5 demonstrates that the Boneh-Franklin IBC offers this property for certain common system parameters and specifies those system parameters. Finally, section 6 summarizes these results and describes future research directions.

2. RELATED WORK

2.1 Identity-Based Cryptosystems

In 1984, Shamir [15] first proposed the concept of identity-based encryption, in which any string uniquely representing a user's identity can serve as the user's public key. His vision relies on the use of a private key generator, which maintains a master secret allowing the production of any private key. The private key generator is responsible for dissemination of private keys, including authentication as necessary. Beyond their encryption and decryption processes, identity-based cryptosystems require a setup process and a private key extraction process. The setup procedure may accept a security parameter and must generate a master secret, a master public key, and any public parameters. Shamir assumes a secure means of initially distributing system parameters, such as the master public key. The private key extraction process requires the PKG to produce the private key for a given identity. Because an identity-based encryption scheme does not require parties to retrieve public keys, it mitigates the overhead and vulnerabilities, such as man-in-the-middle attacks [14], associated with key retrieval in traditional asymmetric cryptosystems. While [15] offers a signature scheme based on the difficulty of factoring products of large primes, Shamir leaves the existence of an identity-based cryptosystem as an open problem.

Boneh and Franklin [9] created the first practical implementation of an identity-based cryptosystem through the use of bilinear pairings, including Weil and Tate pairings, on elliptic curves. The security of their cryptosystem rests on the Bilinear Diffie-Hellman Assumption, described in [9]. Boneh and Franklin offer two schemes: a straightforward design offering semantic security and a more complex design with additional chosen-ciphertext security. This paper restricts discussion to the former scheme, but trivial modification of the IBC privacy proofs extends the conclusions to the latter. The setup, private key extraction, encryption, and decryption processes for the Boneh-Franklin IBC operate as follows:

- Setup: For clarity in later proofs, this paper divides the setup process into two stages.
  1. Choose two groups, $\mathbb{G}_1$ and $\mathbb{G}_2$, of prime order q and a non-degenerate, efficiently computable bilinear map $\hat{e} : \mathbb{G}_1 \times \mathbb{G}_1 \rightarrow \mathbb{G}_2$. Pick two cryptographic hash functions, $H_1 : \{0,1\}^* \rightarrow \mathbb{G}_1^*$ and $H_2 : \mathbb{G}_2 \rightarrow \{0,1\}^n$ for some n where the message space is $\{0,1\}^n$. Select an arbitrary generator $P \in \mathbb{G}_1$.
  2. Choose a random value $s \in \mathbb{Z}_q^*$ to serve as the master secret. All parties using of this IBC must know the parameters $params = (\mathbb{G}_1, \mathbb{G}_2, \hat{e}, n, P, H_1, H_2)$ and public key $pk = sP$.
- Private Key Extraction: For an identity represented by $id \in \{0,1\}^*$, return $d_{id} = sH_1(id)$.
- Encryption: To encrypt $m \in \{0,1\}^n$ with public key u, select a random $r \in \mathbb{Z}_q^*$, and return $c = (rP, m \oplus H_2(\hat{e}(H_1(id), pk)^r))$.
- Decryption: To decrypt $c = (U, V)$ with private key d_{id}, compute $m = V \oplus H_2(\hat{e}(d_{id}, U))$.

Boneh and Franklin propose the ability to distribute a master secret using threshold cryptography [9]. Section 1.1 describes why master secret distribution is an insufficient solution under some scenarios where IBC privacy is desirable.

2.2 Notions of Privacy and Indistinguishability

Calandrino and Weaver [10] describe a property similar to IBC privacy in the context of identity-based signature schemes. Adversaries receive the signature of a secret random nonce and a set of signature scheme instantiations, one of which generated the private key that produced the signature. If an adversary can gain only a negligible advantage in determining the instantiation responsible for the signature, the signature scheme offers the desired privacy property. Calandrino and Weaver note that Shamir's original identity-based signature scheme [15] lacks this privacy due to the publicly known modulus, but they do not prove that any signature scheme possesses this property. Techniques useful for proving IBC privacy may be effective for proving signature scheme privacy and vise versa.

In [6], Bellare, Boldyreva, Desai, and Pointcheval formalize the notion of key-privacy, also known as anonymity, in public key encryption. Given an encryption scheme with this property, a polynomial-time adversary can gain no more than a negligible advantage from ciphertext alone in determining the public key that resulted in the ciphertext. Consequently, an adversary cannot infer the destination of the ciphertext. RSA encryption [13], for example, lacks key-privacy because users' publicly available modulii allow adversaries to gain an advantage [6]. Bellare et al. define two attack models: chosen-plaintext and chosen-ciphertext attack. Halevi [12] explores sufficient conditions for a cryptographic scheme to offer key-privacy. The remainder of this paper considers only chosen-plaintext attacks, such as IBC privacy under chosen-plaintext attack.

To prove the security of a searchable encryption scheme based on the Boneh-Franklin IBC, Boneh, Di Crescenzo, Ostrovsky, and Persiano [8] implicitly demonstrate that the Boneh-Franklin IBC possesses key-privacy under chosen-plaintext attack. Abdalla et al. [1] use the results in [12] to present a simplified formal proof of the anonymity of the Boneh-Franklin IBC. To prove that the Boneh-Franklin IBC possesses IBC privacy, this paper adapts the key-privacy proof in Abdalla et al. [1]. The following section examines that proof.

2.3 Key-Privacy in the Boneh-Franklin Identity-Based Cryptosystem

Because this paper's proof of IBC privacy is a modification of the key-privacy proof in [1], this section summarizes that proof. Recall that an identity-based cryptosystem provides four algorithms: *IBC* = (*Setup*, *KeyExtract*, *Encrypt*, *Decrypt*). *Setup* accepts a security parameter, *k*, and generates a set of public parameters, *params*; a random master secret, *s*; and the corresponding master public key, *pk*. *KeyExtract* accepts a master secret, public parameters, and an identity, *id*, for which the private key is to be extracted. *Encrypt* accepts a master public key, public parameters, an identity, and a message, *m*, and it produces a piece of ciphertext, *c*. *Decrypt* accepts a user's secret key, public parameters, and a piece of ciphertext, and it outputs the message. To prove that a polynomial-time adversary can gain no more than a negligible advantage in a key-privacy challenge for the Boneh-Franklin IBC, Abdalla et al. [1] require several steps. First, they formally define key-privacy in terms of IBCs. They then define semantic security and a weakened key-privacy property for IBCs and show that a cryptosystem with both of these properties possesses key-privacy. Finally, they note that Boneh-Franklin identity-based encryption offers semantic security and prove that it possesses weakened key-privacy.

The following several paragraphs define key-privacy (or anonymity), weakened anonymity, and semantic security in IBCs through experiments that test for these properties. In all cases, assume that the adversary has access to an oracle that can provide the private key associated with any identity. Also, note that the advantage of an adversary, *A*, in guessing the correct value of *b* in an experiment for the property IBC-PROP ($\mathbf{Exp}_{IBC,A}^{\text{ibc-prop-b}}(k)$) is:

$$\mathbf{Adv}_{IBC,A}^{\text{ibc-prop}}(k) = \Pr[\mathbf{Exp}_{IBC,A}^{\text{ibc-prop-1}}(k) = 1] - \Pr[\mathbf{Exp}_{IBC,A}^{\text{ibc-prop-0}}(k) = 1] .$$

IBC Key-Privacy. If an identity-based cryptosystem possesses key-privacy, an adversary is unable to gain more than a negligible advantage at guessing the value b in the following experiment ($\mathbf{Exp}_{IBC,A}^{\text{ibc-ano-cpa-b}}(k)$):

1. The challenger computes (*params, s*, *pk*) = *Setup*(*k*) and presents *params* and *pk* to the adversary.
2. The adversary, *A*, may use the oracle to derive secret keys for any identities. Eventually, *A* must choose two valid identities, id_0 and id_1, based on all known data. *A* must not have queried for secret keys corresponding to these identities. In addition, *A* must generate a message, *m*, within the message space. *A* returns id_0, id_1, and *m* to the challenger and may save any state information.
3. The challenger randomly selects a bit $b \in \{0,1\}$, computes $c = Encrypt(pk, params, id_b, m)$, and returns *c* to the adversary.
4. The adversary may use the oracle again but may not derive secret keys associated with id_0 or id_1. Eventually, the adversary must submit a guess, *b'*, for *b* based on all known information, including saved state data.

Weakened Anonymity. If an IBC possesses weakened key-privacy, an adversary is unable to gain more than a negligible advantage at guessing the value b in the following experiment ($\mathbf{Exp}_{IBC,A}^{\text{ibc-ano-re-b}}(k)$):

1. The challenger computes (*params, s*, *pk*) = *Setup*(*k*) and presents *params* and *pk* to the adversary.
2. The adversary, *A*, may use the oracle to derive secret keys for any identities. Eventually, *A* must choose two valid identities, id_0 and id_1, based on all known data. *A* must not have queried for secret keys corresponding to these identities. In addition, *A* must generate a message, *m*, within the message space. *A* returns id_0, id_1, and *m* to the challenger and may save any state information.
3. The challenger randomly selects a message $m' \in \{0,1\}^{|m|}$ and a bit $b \in \{0,1\}$, computes $c = Encrypt(pk, params, id_b, m')$, and returns *c* to the adversary.
4. The adversary may use the oracle again but may not derive secret keys associated with id_0 or id_1. Eventually, the adversary must submit a guess, *b'*, for *b* based on all known information, including saved state data.

Semantic Security. If an identity-based cryptosystem possesses semantic security, an adversary is unable to gain more than a negligible advantage at guessing the value b in the following experiment ($\mathbf{Exp}_{IBC,A}^{\text{ibc-ind-cpa-b}}(k)$):

1. The challenger computes (*params, s*, *pk*) = *Setup*(*k*) and presents *params* and *pk* to the adversary.
2. The adversary, *A*, may use the oracle to derive secret keys for any identities. Eventually, *A* must choose a valid identity, *id*, based on all known data. *A* must not have queried for a secret key corresponding to this identity. In addition, *A* must generate two messages, m_0 and m_1, of equal length within the message space. *A* returns *id*, m_0, and m_1 to the challenger and may save any state information.
3. The challenger randomly selects a bit $b \in \{0,1\}$, computes $c = Encrypt(pk, P, id, m_b)$, and returns *c* to the adversary.
4. The adversary may use the oracle again but may not derive secret keys associated with *id*. Eventually, the adversary must submit a guess, *b'*, for *b* based on all known information, including saved state data.

Abdalla et al. then demonstrate that any identity-based cryptosystem possessing the latter two properties possesses the former. Suppose that an adversary, *A*, is participating in the experiment $\mathbf{Exp}_{IBC,A}^{\text{ibc-ano-cpa-b}}(k)$. By definition, we can create an adversary, A_1, such that:

$$\Pr[\mathbf{Exp}_{IBC,A}^{\text{ibc-ano-re-1}}(k)=1]-\Pr[\mathbf{Exp}_{IBC,A}^{\text{ibc-ano-re-0}}(k)=1]\le \mathbf{Adv}_{IBC,A_1}^{\text{ibc-ano-re}}(k).$$

In addition, because the difference between key-privacy and weakened anonymity is based on the difficulty of inferring the message behind given ciphertext, two adversaries, A_2 and A_3, exist such that:

$$\Pr[\mathbf{Exp}_{IBC,A}^{\text{ibc-ano-cpa-1}}(k)=1]-\Pr[\mathbf{Exp}_{IBC,A}^{\text{ibc-ano-re-1}}(k)=1]\le \mathbf{Adv}_{IBC,A_2}^{\text{ibc-ind-cpa}}(k)$$
$$\Pr[\mathbf{Exp}_{IBC,A}^{\text{ibc-ano-re-0}}(k)=1]-\Pr[\mathbf{Exp}_{IBC,A}^{\text{ibc-ano-cpa-0}}(k)=1]\le \mathbf{Adv}_{IBC,A_3}^{\text{ibc-ind-cpa}}(k)$$

Therefore, if adversaries can gain no more than a negligible advantage in successfully guessing b in $\mathbf{Exp}_{IBC,A}^{\text{ibc-ano-re-b}}(k)$ and $\mathbf{Exp}_{IBC,A}^{\text{ibc-ind-cpa-b}}(k)$, then summing the three inequalities shows that the advantage of A is bounded by a sum of three negligible functions.

In [9], Boneh and Franklin prove that their IBC possesses semantic security. By [1, 12], to demonstrate that this cryptosystem possesses weakened anonymity, one may simply show that the ciphertext, c, that the adversary receives in $\mathbf{Exp}_{IBC,A}^{\text{ibc-ano-re-b}}(k)$ has the same uniformly random distribution regardless whether $b = 0$ or $b = 1$. For the Boneh-Franklin IBC, $c=(rP, m'\oplus H_2(\hat{e}(H_1(id_b), pk)^r))$ in $\mathbf{Exp}_{IBC,A}^{\text{ibc-ano-re-b}}(k)$. Since the encrypting party chooses r at random from $\mathbb{Z}_q^*$, the value rP is chosen uniformly at random from $\mathbb{G}_1^*$ regardless of the value of b [1]. Similarly, m' is chosen uniformly at random from $\{0,1\}^{|m|}$ and is independent of the output of H_2, so $m'\oplus H_2(\hat{e}(H_1(id_b), pk)^r$ is chosen uniformly at random from $\{0,1\}^n$ regardless of the value of b [1]. Therefore, an adversary would be unable to gain an advantage at guessing the value of b in $\mathbf{Exp}_{IBC,A}^{\text{ibc-ano-re-b}}(k)$. Because the Boneh-Franklin IBC possesses semantic security and weakened key-privacy, it also possesses key-privacy.

2.4 Applicability of IBC Privacy to Existing Work

In [16], Waters, Balfanz, Durfee, and Smetters introduce a scheme for encrypted-yet-searchable audit logs. Their asymmetric scheme relies on the use of Boneh-Franklin identity-based encryption, including its property of anonymity. IBC privacy could allow entities with different IBCs and audit escrow agents to host encrypted audit records on the same servers to disguise the size of individual audit logs. Without an exploration of cross-domain optimizations, this change may result in dramatic performance penalties, however. Consequently, future work in this area would need to focus on optimizations.

An interesting potential use of IBC privacy appears in [2]. Adida, Hohenberger, and Rivest describe an architecture for the prevention of spoofed emails that maintains sender repudability. The architecture relies on identity-based signature schemes and can work with any such scheme that allows for separable identity-based ring signatures [4]. Adida et al. presume the use of separate PKGs for each email domain and use DNS servers to distribute domains' system parameters and master public keys [2]. Given the selection of certain signature schemes and cryptosystems, a single infrastructure can support both identity-based encryption and signatures. By specifying a signature scheme with infrastructure that also supports an IBC with IBC privacy, the Adida et al. architecture [2] could allow massive coalitions of IBCs such that adversaries would be unable to gather more information from transmitted ciphertext than they could from aggregated Internet traffic data alone. Note that [3] also presents extensions to the system in [2] that allow for encryption. While [3] does not consider IBC privacy, it does offer additional desirable security properties.

3. DEFINITION OF IDENTITY-BASED CRYPTOSYSTEM PRIVACY

For the purpose of defining IBC privacy, the remainder of this paper treats the IBC setup process as two stages, $Setup_1$ and $Setup_2$. $Setup_1$ generates any necessary common system parameters. Rather than independently using $Setup_1$, this section assumes that cryptosystem operators agree in advance on all common system parameters, *commonParams*, for a given security parameter, k, which is common across cryptosystem instantiations. $Setup_2$ accepts a set of common parameters and a security parameter and generates a complete set of system parameters, a master secret, and the corresponding master public key. Following is the definition for IBC privacy. Assume that the adversary has access to an oracle that can provide the private key associated with any identity

If an identity-based cryptosystem possesses IBC privacy, an adversary is unable to gain more than a negligible advantage at guessing the value b in the following experiment ($\mathbf{Exp}_{IBC,A}^{\text{ibc-privacy-cpa-b}}(k)$):

1. The challenger computes ($params_0$, s_0, pk_0) = $Setup_2$(*commonParams*, k) and ($params_1$, s_1, pk_1) = $Setup_2$(*commonParams*, k) and presents $params_0$, pk_0, $params_1$, and pk_1 to the adversary.
2. The adversary, A, may use the oracle to derive secret keys for any identities under either instantiation. Eventually, A must choose two valid identities, id_0 and id_1, where id_0 corresponds to the cryptosystem instantiation with pk_0 and id_1 corresponds to the cryptosystem instantiation with pk_1. A must not have queried for the secret keys corresponding to these identities. In addition, A must generate a message, m, within the message space of both instantiations. A returns id_0, id_1, and m to the challenger and may save any state information.
3. The challenger randomly selects a bit $b \in \{0,1\}$, computes c = *Encrypt*(pk_b, $params_b$, id_b, m), and returns c to the adversary.
4. The adversary may make use of the oracle once more but may not derive secret keys associated with id_0 or id_1. Eventually, the adversary must submit a guess, *b'*, for b based on all known information, including saved state data.

4. A SUFFICIENT CONDITION FOR IBC PRIVACY

Analogous to key-privacy, possession of semantic security and a weakened notion of IBC privacy, defined below, is sufficient to show possession of IBC privacy.

Weakened IBC Privacy. If an adversary is unable to gain more than a negligible advantage at guessing the value b in the following experiment ($\mathbf{Exp}_{IBC,A}^{\text{ibc-privacy-re-b}}(k)$), the corresponding IBC possesses a weaker notion of key privacy:

1. The challenger computes ($params_0$, s_0, pk_0) = $Setup_2$(*commonParams*, k) and ($params_1$, s_1, pk_1) = $Setup_2$(*commonParams*, k) and presents $params_0$, pk_0, $params_1$, and pk_1 to the adversary.
2. The adversary, A, may use the oracle to derive secret keys for any identities under either instantiation. Eventually, A must choose two valid identities, id_0 and id_1, where id_0 corresponds to the cryptosystem instantiation with pk_0 and id_1 corresponds to the cryptosystem instantiation with pk_1. A must not have queried for the secret keys corresponding to these identities. In addition, A must generate a message, m, within the message space of both instantiations. A returns id_0, id_1, and m to the challenger and may save any state information.
3. The challenger randomly selects a message $m' \in \{0,1\}^{|m|}$ and a bit $b \in \{0,1\}$, computes c = *Encrypt*(pk_b, $params_b$, id_b, m'), and returns c to the adversary.
4. The adversary may make use of the oracle once more but may not derive secret keys associated with id_0 or id_1. Eventually, the adversary must submit a guess, *b'*, for b based on all known information, including saved state data.

Following the logic in section 2.3, adversaries A_1, A_2, and A_3 exist such that the following inequalities hold:

$$\Pr[\mathbf{Exp}_{IBC,A}^{\text{ibc-privacy-re-1}}(k)=1]-\Pr[\mathbf{Exp}_{IBC,A}^{\text{ibc-privacy-re-0}}(k)=1]\le \mathbf{Adv}_{IBC,A_1}^{\text{ibc-privacy-re}}(k)$$
$$\Pr[\mathbf{Exp}_{IBC,A}^{\text{ibc-privacy-cpa-1}}(k)=1]-\Pr[\mathbf{Exp}_{IBC,A}^{\text{ibc-privacy-re-1}}(k)=1]\le \mathbf{Adv}_{IBC,A_2}^{\text{ibc-ind-cpa}}(k)$$
$$\Pr[\mathbf{Exp}_{IBC,A}^{\text{ibc-privacy-re-0}}(k)=1]-\Pr[\mathbf{Exp}_{IBC,A}^{\text{ibc-privacy-cpa-0}}(k)=1]\le \mathbf{Adv}_{IBC,A_3}^{\text{ibc-ind-cpa}}(k)$$

Summing these inequalities demonstrates that possession of semantic security and weakened IBC privacy under certain common parameters is sufficient for an IBC to offer IBC privacy under those common parameters.

5. IBC PRIVACY IN THE BONEH-FRANKLIN IDENTITY-BASED CRYPTOSYSTEM

Based on the results in section 4, proof that the Boneh-Franklin IBC possesses IBC privacy for certain common parameters requires only proof that the Boneh-Franklin IBC has semantically security and weakened IBC privacy under those common parameters. Recall that [9] proves that the Boneh-Franklin IBC is semantically secure. To prove possession of weakened IBC privacy, follow the same process for proving possession of key-privacy, and initially assume that all parameters are common. By [1, 12], to demonstrate that the cryptosystem possesses weakened IBC privacy, one may simply show that the ciphertext, c, that the adversary receives in $\mathbf{Exp}_{IBC,A}^{\text{ibc-privacy-re-b}}(k)$ has the same uniformly random distribution regardless whether $b = 0$ or $b = 1$. For the Boneh-Franklin IBC, $c=(rP, m'\oplus H_2(\hat{e}(H_1(id_b), pk_b)^r))$ in $\mathbf{Exp}_{IBC,A}^{\text{ibc-privacy-re-b}}(k)$. As section 2.3 explains, rP is chosen uniformly at random from $\mathbb{G}_1^*$ regardless of the value of b. Also, m' is chosen uniformly at random from $\{0,1\}^{|m|}$ and is independent of the output of H_2. Thus, $m'\oplus H_2(\hat{e}(H_1(id_b), pk_b)^r$ is chosen uniformly at random from $\{0,1\}^n$ regardless of the value of b. Therefore, an adversary would be unable to gain an advantage at guessing the value of b in $\mathbf{Exp}_{IBC,A}^{\text{ibc-privacy-re-b}}(k)$. Because the Boneh-Franklin IBC possesses semantic security and weakened IBC privacy when all parameters are common, it also possesses IBC privacy under those common parameters.

Cryptosystem operators may find it useful to know which parameters can vary across instantiations without negating IBC privacy in the Boneh-Franklin IBC. Recall that the Boneh-Franklin IBC parameters are ($\mathbb{G}_1$, $\mathbb{G}_2$, $\hat{e}$, n, P, H_1, H_2). Because rP is chosen uniformly at random from $\mathbb{G}_1^*$, any modification of $\mathbb{G}_1$ across instantiations would offer an adversary an advantage in distinguishing between those instantiations. Because the randomness of $m'\oplus H_2(\hat{e}(H_1(id_b), pk_b)^r$ depends only on m' and the fact that H_2 maps from $\mathbb{G}_2$ to $\{0,1\}^n$, $\mathbb{G}_2$, $\hat{e}$, H_1, and H_2 may differ across instantiations as long as H_2 maps from $\mathbb{G}_2$ to $\{0,1\}^n$. n may not differ across instantiations: $m'\oplus H_2(\hat{e}(H_1(id_b), pk_b)^r$ is chosen uniformly at random from $\{0,1\}^n$, so a change in n would offer adversaries a non-negligible advantage. Finally, as the randomness of rP depends on the value r and the randomness of $m'\oplus H_2(\hat{e}(H_1(id_b), pk_b)^r$ depends on the value m', P may differ across instantiations provided that it is a generator of $\mathbb{G}_1$. Therefore, though choices for certain parameters may constrain selection of other parameters, only $\mathbb{G}_1$ and n must be common across instantiations for the Boneh-Franklin IBC to offer IBC privacy.

6. CONCLUSION AND FUTURE WORK

This paper makes three key contributions. First, it defines IBC privacy and provides motivation for exploration of this property. Second, it demonstrates that existing research on the topic of key-privacy is applicable to research on IBC privacy. Lastly, it demonstrates that the Boneh-Franklin IBC offers IBC privacy given certain common system parameters and establishes which parameters must be common across instantiations. Future research may wish to consider IBC privacy in different cryptosystems, such as the Cocks system [11], the Waters system [17], or the Boneh-Boyen system [7]. Because these three systems do not possess key-privacy [1, 8], they may allow further

exploration of the relationship between the similar IBC privacy and key-privacy properties. Future work exploring IBC privacy between cryptosystem types, such as the Boneh-Franklin IBC and the Waters IBC, may also be desirable.

7. ACKNOWLEDGEMENTS

We would like to thank Mihir Bellare for directing us towards existing work that proved critical in our proofs. We also thank Brent Waters for his helpful intuition regarding which IBCs may possess IBC privacy.

REFERENCES

[1] M. Abdalla, M. Bellare, D. Catalano, E. Kiltz, T. Kohno, T. Lange, J. Malone-Lee, G. Neven, P. Paillier, H. Shi. Searchable encryption revisted: Consistency properties, relation to anonymous IBE, and extensions (Full version). Cryptology ePrint Archive, Report 2005/254, 2005. http://eprint.iacr.org/2005/254/.

[2] B. Adida, S. Hohenberger, R. L. Rivest. Fighting phishing attacks: A lightweight trust architecture for detecting spoofed emails.

[3] B. Adida, S. Hohenberger, R. L. Rivest. Lightweight encryption for email.

[4] B. Adida, S. Hohenberger, R. L. Rivest. Separable identity-based ring signatures: Theoretical foundations for fighting phishing attacks.

[5] J. Baek, J. Newmarch, R. Safavi-Naini, W. Susilo. A survey of identity-based cryptography.

[6] M. Bellare, A. Boldyreva, A. Desai, D. Pointcheval. Key-privacy in public-key encryption. In *Proc. of Advances in Cryptology – ASIACRYPT 2001*. Springer-Verlag, 2001. LNCS 2248.

[7] D. Boneh and X. Boyen. Secure identity based encryption without random oracles. In *Proc. of CRYPTO 2004*, pages 443-459. Springer-Verlag, 2004. LNCS 3152.

[8] D. Boneh, G. Di Crescenzo, R. Ostrovsky, G. Persiano. Public key encryption with keyword search. In *Proc. of EUROCRYPT 2004*, pages 506-522. Springer-Verlag, 2004. LNCS 3027.

[9] D. Boneh, M. Franklin. Identity-based encryption from the Weil pairing. In *Proc. of CRYPTO 2001*, pages 213-229. Springer-Verlag, 2001. LNCS 2248.

[10] J. A. Calandrino, A. C. Weaver. Private Resource Pairing.

[11] C. Cocks. An identity based encryption scheme based on quadratic residues.

[12] S. Halevi. A sufficient condition for key-privacy. Cryptology ePrint Archive, Report 2005/005, 2005. http://eprint.iacr.org/2005/005.

[13] R. L. Rivest, A. Shamir, L. M. Adleman. A method for obtaining digital signatures and public-key cryptosystems.

[14] B. Schneier. *Applied Cryptography: Protocols, Algorithms, and Source Code in C.* John Wiley & Sons, 1994.

[15] A. Shamir. Identity-based cryptosystems and signature schemes. In *Proc. of CRYPTO 1984*, pages 47-53. Springer-Verlag, 1985. LNCS 196.

[16] B. R. Waters, D. Balfanz, G. Durfee, D. K. Smetters. Building an encrypted and searchable audit log. In *Proc. of 11th Annual Network and Distributed System Security Symposium*. 2004.

[17] B. R. Waters. Efficient identity-based encryption without random oracles. In *Proc. of EUROCRYPT 2005*, pages 440-456. Springer-Verlag, 2005. LNCS.