Multipurpose Identity-Based Signcryption
A Swiss Army Knife for Identity-Based Cryptography

Xavier Boyen

IdentiCrypt, 420 Florence, Palo Alto, California — crypto@boyen.org

Abstract. Identity-Based (IB) cryptography is a rapidly emerging approach to public-key cryptography that does not require principals to pre-compute key pairs and obtain certificates for their public keys—instead, public keys can be arbitrary identifiers such as email addresses, while private keys are derived at any time by a trusted private key generator upon request by the designated principals. Despite the flurry of recent results on IB encryption and signature, some questions regarding the security and efficiency of practicing IB encryption (IBE) and signature (IBS) as a joint IB signature/encryption (IBSE) scheme with a common set of parameters and keys, remain unanswered.

We first propose a stringent security model for IBSE schemes. We require the usual strong security properties of: (for confidentiality) indistinguishability against adaptive chosen-ciphertext attacks, and (for non-repudiation) existential unforgeability against chosen-message insider attacks. In addition, to ensure as strong as possible ciphertext armoring, we also ask (for anonymity) that authorship not be transmitted in the clear, and (for unlinkability) that it remain unverifiable by anyone except (for authentication) by the legitimate recipient alone.

We then present an efficient IBSE construction, based on bilinear pairings, that satisfies all these security requirements, and yet is as compact as pairing-based IBE and IBS in isolation. Our scheme is secure, compact, fast and practical, offers detachable signatures, and supports multi-recipient encryption with signature sharing for maximum scalability.

1 Introduction

Recently, Boneh and Franklin [5] observed that bilinear pairings on elliptic curves could be used to make identity-based encryption possible and practical. Following this seminal insight, the last couple of years have seen a flurry of results on a number of aspects of what has now become the nascent field of Identity-Based (IB) cryptography.

1.1 Identity-Based Cryptography

The distinguishing characteristic of IB cryptography is the ability to use any string as a public key; the corresponding private key can only be derived by a trusted Private Key Generator (PKG), custodian of a master secret. For encryption purposes, this allows Alice to securely send Bob an encrypted message, using as public key any unambiguous name identifying Bob, such as Bob's email address, possibly before Bob even knows his own private key. For signature purposes, Alice may sign her communications using a private key that corresponds to an unambiguous name of hers, so that anybody can verify the authenticity of the signature simply from the name, without the need for a certificate. Revocation issues are handled by using short-lived time-dependent identities [5].

An inherent limitation of IB cryptography is the trust requirement that is placed on the PKG, as an untrustworthy PKG will have the power to forge Alice's signature, and decrypt Bob's past and future private communications. This can be partially alleviated by splitting the master secret among several PKGs under the jurisdiction of no single entity, as explained in [5]. The window of vulnerability can also be reduced by periodically changing the public parameters, and purging any master secret beyond a certain age, effectively limiting the interval during which IB cryptograms can be decrypted. Traditional public-key cryptography is not completely immune to the problem, either: in a public key infrastructure, the certification authority has the power to issue fake certificates and impersonate any user for signature purposes; it can similarly spoof encryption public key certificates in order to decrypt future ciphertexts addressed to targeted users, albeit not in a manner not amenable to easy detection.

The idea of IB cryptography first emerged in 1984 [24], although only an IB signature (IBS) scheme was then suggested, based on conventional algebraic methods in $\mathbb{Z}_n$. Other IBS and identification schemes were quick to follow [13, 12]. However, it is only in 2001 that a practical IB encryption (IBE) mechanism was finally suggested [5], based on the much heavier machinery of bilinear pairings on elliptic curves, whose use in cryptography had slowly started to surface in the few years prior, *e.g.*, for key exchange [18] and IBS [23]. Interestingly, a more conventional approach to IBE was proposed shortly thereafter [10], albeit not as efficient as the pairing-based IBE scheme of Boneh and Franklin.

1.2 Motivation and Contribution

Following the original publication of the BF-IBE scheme [5], a number of authors have proposed various new applications of pairing-based IB cryptography. These include various IB signature schemes [22, 16, 8], key agreement [25], a 2-level hierarchical IB encryption [17], and a general hierarchical IB encryption that can also be used to produce signatures [15]. More specifically focusing on joint authentication and encryption, we note a repudiable authenticated IBE [20], an authenticated key agreement scheme [9], and a couple of IB signcryption schemes that efficiently combine signature and encryption [21, 19].

What the picture is currently missing is an algorithm that combines (existing or new) IBE and IBS in a practical and secure way. Indeed, it would be of great practical interest to be able to use the same IB infrastructure for signing and encrypting. A possibility is to combine some existing IBE and IBS using black-box composition techniques, such as [1]; this is however rather suboptimal.

A better approach would be to exploit the similarities between IBE and IBS, and elaborate a dual-purpose IB Encryption-Signature (IBSE) scheme based on a shared infrastructure, toward efficiency increases and security improvements. Doing so, we would have to ensure that no hidden weakness arises from the combination, which is always a risk if the same parameters and keys are used. The issues that arise from this approach are summarized as follows:

– Can IBE and IBS be practiced in conjunction, in a secure manner, sharing infrastructure, parameters, and keys, toward greater efficiency?
– What emerging security properties can be gained from such a combination?

Our contributions to answering these questions are twofold. We first specify a security model that a strong IBSE combination should satisfy. Our model specifies the IBSE version of the strongest notions of security usually considered in public-key cryptography. For confidentiality, we define a notion of ciphertext indistinguishability under adaptive chosen-ciphertext attacks. For non-repudiation, we define a notion of signature unforgeability under chosen-message attacks, in the stringent case of an 'insider' adversary, *i.e.*, with access to the decryption private key, as considered in [1]. We also specify the additional security features of ciphertext authentication, anonymity, and unlinkability, that, if less conventional, are highly desirable in practice: together, they convince the legitimate recipient of the ciphertext origin, and conceal it from anyone else.

We then propose a fast IBSE scheme satisfying our strong security requirements, which we prove in the random oracle model [3]. Our scheme uses the properties of bilinear pairings to achieve a two-layer sign-then-encrypt combination, featuring a detachable randomized signature, followed by anonymous deterministic encryption. The scheme is very efficient, more secure than what we call *monolithic* signcryption—in which a single operation is used for decryption and signature verification, as in the original signcryption model of [27]—and more compact than generic compositions of IBE and IBS. Our two-layer design is also readily adapted to provide multi-recipient encryption of the same message with a shared signature and a single bulk message encryption.

Performance-wise, our dual-purpose optimized IBSE scheme is as compact as most existing single-purpose IBE and IBS taken in isolation. It is also about as efficient as the monolithic IB signcryption schemes of [21] and [19], with the added flexibility and security benefits that separate anonymous decryption and signature verification layers can provide. A comparative summary of our scheme with competing approaches can be found in §6, Table 2.

1.3 Outline of the Paper

We start in §2 by laying out the abstract IBSE specifications. In §3, we formalize the various security properties sought from the cryptosystem. In §4, we review the principles of IB cryptography based on pairings. In §5, we describe an implementation of our scheme. In §6, we make detailed performance and security comparisons with the competition. In §7, we prove compliance of our implementation with the security model. In §8, we study a few extensions of practical significance. Finally, in §9, we draw some conclusions.

2 Specification of the Cryptosystem

An *Identity-Based Signature/Encryption* scheme, or IBSE, consists of a suite of six algorithms: Setup, $\mathsf{Extract}$, Sign, $\mathsf{Encrypt}$, $\mathsf{Decrypt}$, and Verify. In essence, Setup generates random instances of the common public parameters and master secret; $\mathsf{Extract}$ computes the private key corresponding to a given public identity string; Sign produces a signature for a given message and private key; $\mathsf{Encrypt}$ encrypts a signed plaintext for a given identity; $\mathsf{Decrypt}$ decrypts a ciphertext using a given private key; Verify checks the validity of a given signature for a given message and identity. Messages are arbitrary strings in $\{0,1\}^*$.

The functions that compose a generic IBSE are thus specified as follows.

Setup On input 1^n, produces a pair $\langle \sigma, \pi \rangle$ (where σ is a randomly generated master secret and π the corresponding common public parameters, for the security meta-parameter n).

$\mathsf{Extract}_{\pi,\sigma}$ On input id, computes a private key pvk (corresponding to the identity id under $\langle \sigma, \pi \rangle$).

Sign_π On input $\langle \mathsf{pvk}_A, \mathsf{id}_A, m \rangle$, outputs a signature s (for pvk_A, under π), and some ephemeral state data r.

$\mathsf{Encrypt}_\pi$ On input $\langle \mathsf{pvk}_A, \mathsf{id}_B, m, s, r \rangle$, outputs an anonymous ciphertext c (containing the signed message $\langle m, s \rangle$, encrypted for the identity id_B under π).

$\mathsf{Decrypt}_\pi$ On input $\langle \mathsf{pvk}_B, \hat{c} \rangle$, outputs a triple $\langle \hat{\mathsf{id}}_A, \hat{m}, \hat{s} \rangle$ (containing the purported sender identity and signed message obtained by decrypting $\hat{c}$ by the private key pvk_B under π).

Verify_π On input $\langle \hat{\mathsf{id}}_A, \hat{m}, \hat{s} \rangle$, outputs $\top$ 'true' or $\perp$ 'false' (indicating whether $\hat{s}$ is a valid signature for the message $\hat{m}$ by the identity $\hat{\mathsf{id}}_A$, under π).

Since we are concerned with sending messages that are simultaneously encrypted and signed, we allow the encryption function to make use of the private key of the sender. Accordingly, we assume that $\mathsf{Encrypt}$ is always used on an output from Sign, so that we may view the $\mathsf{Sign}/\mathsf{Encrypt}$ composition as a single 'signcryption' function; we keep them separate to facilitate the treatment of multi-recipient encryption with shared signature in §8.1. We also insist on the dichotomy $\mathsf{Decrypt}$ *vs.* Verify, to permit the decryption of anonymous ciphertexts, and to decouple signature verification from the data that is transmitted over the wire, neither of which would be feasible had we used a monolithic 'unsigncryption' function.

It is required that these algorithms jointly satisfy the following consistency constraints.

Definition 1. For all master secret and common parameters $\langle \sigma, \pi \rangle \leftarrow \mathsf{Setup}[1^n]$, any identities id_A and id_B, and matching private keys $\mathsf{pvk}_A = \mathsf{Extract}_{\pi,\sigma}[\mathsf{id}_A]$ and $\mathsf{pvk}_B = \mathsf{Extract}_{\pi,\sigma}[\mathsf{id}_B]$, we require for consistency that, $\forall m \in \{0,1\}^*$:

$$\left\{ \begin{array}{r} \langle s, r \rangle \leftarrow \mathsf{Sign}_\pi[\mathsf{pvk}_A, \mathsf{id}_A, m] \\ c \leftarrow \mathsf{Encrypt}_\pi[\mathsf{pvk}_A, \mathsf{id}_B, m, s, r] \\ \langle \hat{\mathsf{id}}_A, \hat{m}, \hat{s} \rangle \leftarrow \mathsf{Decrypt}_\pi[\mathsf{pvk}_B, c] \end{array} \right\} \Longrightarrow \left\{ \begin{array}{r} \hat{\mathsf{id}}_A = \mathsf{id}_A \\ \hat{m} = m \\ \mathsf{Verify}_\pi[\mathsf{id}_A, \hat{m}, \hat{s}] = \top \end{array} \right\}$$

In the sequel, we omit the subscripted parameters π and σ when understood from context.

3 Formal Security Model

Due to the identity-based nature of our scheme, and the combined requirements on confidentiality and non-repudiation, the security requirements are multi-faceted and quite stringent. For example, for confidentiality purposes, one should assume that the adversary may obtain any private key other than that of the targeted recipient, and has an oracle that decrypts any valid ciphertext other than the challenge. For non-repudiation purposes, we assume that the forger has access to any private key other than that of the signer, and can query an oracle that signs and encrypts any message but the challenge. These assumptions essentially amount to the 'insider' model in the terminology of [1].

We also consider the notions of *ciphertext unlinkability* and *ciphertext authentication*, which allow the legitimate recipient to privately verify—but not prove to others—that the ciphertext addressed to him and the signed message it contains were indeed produced by the same entity. We note that these properties are not jointly achieved by other schemes that combine confidentiality and non-repudiation, such as the signcryption of [21] and [19]. We also ask for *ciphertext anonymity*, which simply means that no third party should be able to discover whom a ciphertext originates from or is addressed to, if the sender and recipient wish to keep that a secret.

All these properties are recapitulated as follows.

1. *message confidentiality* (§3.1): allows the communicating parties to preserve the secrecy of their exchange, if they choose to.
2. *signature non-repudiation* (§3.2): makes it universally verifiable that a message speaks in the name of the signer (regardless of the ciphertext used to convey it, if any). This implies message authentication and integrity.
3. *ciphertext unlinkability* (§3.3): allows the sender to disavow creating a ciphertext for any given recipient, even though he or she remains bound to the valid signed message it contains.
4. *ciphertext authentication* (§3.4): allows the legitimate recipient, alone, to be convinced that the ciphertext and the signed message it contains were crafted by the same entity. This implies ciphertext integrity.
5. *ciphertext anonymity* (§3.5): makes the ciphertext appear anonymous (hiding both the sender and the recipient identities) to anyone who does not possess the recipient decryption key.

For simplicity of the subsequent analysis, we disallow messages from being addressed to the same identity as authored them—a requirement that we call the *irreflexivity assumption*. Remark that if such a mode of operation is nonetheless desired, it can easily be achieved, either, (1) by endowing each person with an additional 'self' identity, under which they can encrypt messages signed under their regular identity, or, (2) by splitting each identity into a 'sender' identity and a 'recipient' identity, to be respectively used for signature and encryption purposes. This can be done, *e.g.*, by prepending an indicator bit to all identity strings; each individual would then be given two private keys by the PKG.

For clarity, and regardless of which of the above convention is chosen, if any, we use the subscripts 'A' for Alice the sender and 'B' for Bob the recipient.

3.1 Message Confidentiality

Message confidentiality against adaptive chosen-ciphertext attacks is defined in terms of the following game, played between a challenger and an adversary. We combine signature and encryption into a dual-purpose oracle, to allow $\mathsf{Encrypt}$ to access the ephemeral random state data r from Sign.

Start The challenger runs the Setup procedure for a given value of the security parameter n, and provides the common public parameters π to the adversary, keeping the secret σ for itself.

Phase 1 The adversary makes a number of queries to the challenger, in an adaptive fashion (*i.e.*, one at a time, with knowledge of the previous replies). The following queries are allowed:

- ***signature/encryption queries*** in which the adversary submits a message and two distinct identities, and obtains a ciphertext containing the message signed in the name of the first identity and encrypted for the second identity;
- ***decryption queries*** in which the adversary submits a ciphertext and an identity, and obtains the identity of the sender, the decrypted message, and a valid signature, provided that (1) the decrypted identity of the sender differs from that of the specified recipient, and (2) the signature verification condition $\mathsf{Verify} = \top$ is satisfied; otherwise, the oracle only indicates that the ciphertext is invalid for the specified recipient;
- ***private key extraction queries*** in which the adversary submits an identity, and obtains the corresponding private key;

Selection At some point, the adversary returns two distinct messages m_0 and m_1 (assumed of equal length), a signer identity id_A, and a recipient identity id_B, on which it wishes to be challenged. The adversary must have made no private key extraction query on id_B.

Challenge The challenger flips $b \in \{0,1\}$, computes $\mathsf{pvk}_A = \mathsf{Extract}[\mathsf{id}_A]$, $\langle s, r\rangle \leftarrow \mathsf{Sign}[\mathsf{pvk}_A, m_b]$, $c \leftarrow \mathsf{Encrypt}[\mathsf{pvk}_A, \mathsf{id}_B, m_b, s, r]$, and returns the ciphertext c as challenge to the adversary.

Phase 2 The adversary adaptatively issues a number of additional encryption, decryption, and extraction queries, under the additional constraint that it not ask for the private key of id_B or the decryption of c under id_B.

Response The adversary returns a guess $\hat{b} \in \{0,1\}$, and wins the game if $\hat{b} = b$.

It is emphasized that the adversary is allowed to know the private key pvk_A corresponding to the signing identity, which gives us *insider-security* for confidentiality [1]. On the one hand, this is necessary if confidentiality is to be preserved in case the sender's private key becomes compromised. On the other hand, this will come handy when we study a 'repudiable' IBSE variant in §8.2.

This game is very similar to the IND-ID-CCA attack in [5]; we call it an IND-IBSE-CCA attack.

Definition 2. An identity-based joint encryption and signature (IBSE) scheme is said to be semantically secure against adaptive chosen-ciphertext insider attacks, or *IND-IBSE-CCA secure*, if no randomized polynomial-time adversary has a non-negligible advantage in the above game. In other words, any randomized polynomial-time IND-IBSE-CCA adversary $\mathcal{A}$ has an advantage $\mathbf{Adv}_{\mathcal{A}}[n] = |\mathbf{P}[\hat{b} = b] - \frac{1}{2}|$ that is $o[1/\mathrm{poly}[n]]$ for any polynomial $\mathrm{poly}[n]$ in the security parameter.

Remark that we insist that the decryption oracle perform a validity check before returning a decryption result, even though $\mathsf{Decrypt}$ does not specify it. This requirement hardly weakens the model, and allows for stronger security results. We similarly ask that the oracles enforce the irreflexivity assumption, *e.g.*, by refusing to produce or decrypt non-compliant ciphertexts.

3.2 Signature Non-Repudiation

Signature non-repudiation is formally defined in terms of the following game, played between a challenger and an adversary.

Start The challenger runs the Setup procedure for a given value of the security parameter n, and provides the common public parameters π to the adversary, keeping the secret σ for itself.

Query The adversary makes a number of queries to the challenger. The attack may be conducted adaptively, and allows the same queries as in the Confidentiality game of §3.1, namely: signature/encryption queries, decryption queries, and private key extraction queries.

Forgery The adversary returns a recipient identity id_B and a ciphertext c.

Outcome The adversary wins the game if the ciphertext c decrypts, under the private key of id_B, to a signed message $\langle \mathsf{id}_A, \hat{m}, \hat{s} \rangle$ that satisfies $\mathsf{id}_A \neq \mathsf{id}_B$ and $\mathsf{Verify}[\mathsf{id}_A, \hat{m}, \hat{s}] = \top$, provided that (1) no private key extraction query was made on id_A, and (2) no signature/encryption query was made that involved $\hat{m}$, id_A, and some recipient $\mathsf{id}_{B'}$, and resulted in a ciphertext c' whose decryption under the private key of $\mathsf{id}_{B'}$ is the claimed forgery $\langle \mathsf{id}_A, \hat{m}, \hat{s} \rangle$.

Such a model is very similar to the usual notion of existential unforgeability against chosen-message attacks [11, 26]; we call it an EUF-IBSE-CMA attack.

Definition 3. An IBSE scheme is said to be existentially signature-unforgeable against chosen-message insider attacks, or *EUF-IBSE-CMA secure*, if no randomized polynomial-time adversary has a non-negligible advantage in the above game. In other words, any randomized polynomial-time EUF-IBSE-CMA adversary $\mathcal{A}$ has an advantage $\mathbf{Adv}_{\mathcal{A}}[n] = \mathbf{P}[\mathsf{Verify}[\mathsf{id}_A, \hat{m}, \hat{s}] = \top]$ that behaves as $o[1/\mathrm{poly}[n]]$ for any polynomial $\mathrm{poly}[n]$.

In the above experiment, the adversary is allowed to obtain the private key pvk_B for the forged message recipient id_B, which corresponds to the stringent requirements of *insider-security* for authentication [1]. There is one important difference, however: in [1], non-repudiation applies to the ciphertext itself, which is the only sensible thing to do in the context of a signcryption model with a monolithic 'unsigncryption' function. Here, given our two-step $\mathsf{Decrypt}/\mathsf{Verify}$ specification, we define non-repudiation with respect to the decrypted signature, which is more intuitive and does not preclude ciphertext unlinkability (see §3.3).

3.3 Ciphertext Unlinkability

Ciphertext unlinkability is the property that makes it possible for Alice to deny having sent a given ciphertext to Bob, even if the ciphertext decrypts (under Bob's private key) to a message bearing Alice's signature. In other words, the signature should only be a proof of authorship of the plaintext message, and not the ciphertext. (We shall make one exception to this requirement in §3.4, where we seek that the legitimate recipient be able privately authenticate the ciphertext, in order to be convinced that it is indeed addressed to him or her.)

Ciphertext unlinkability allows Alice, *e.g.*, as a news correspondent in a hostile area, to stand behind the content of her reporting, but conceal any detail regarding the particular channel, method, place, or time of communication, lest subsequent forensic investigations be damaging to her sources. When used in conjunction with the multi-recipient technique of §8.1, this property also allows her to deniably provide exact copies of her writings to additional recipients.

We do not present a formal experiment for this property. Suffice it to say that it is enough to ask that, given a plaintext message signed by Alice, Bob be able to create a valid ciphertext addressed to himself for that message, that is indistinguishable from a genuine ciphertext from Alice.

Definition 4. An IBSE scheme is said to be ciphertext-unlinkable if there exists a polynomial-time algorithm that, given an identified signed message $\langle \mathsf{id}_A, m, s\rangle$ such that $\mathsf{Verify}[\mathsf{id}_A, m, s] = \top$, and a private key $d_B = \mathsf{Extract}[\mathsf{id}_B]$, assembles a ciphertext c that is computationally indistinguishable from a genuine encryption of $\langle m, s\rangle$ by id_A for id_B.

As mentioned earlier, ciphertext unlinkability is the reason why we considered the notion of signature unforgeability in §3.2, instead of the usual notion of ciphertext unforgeability as studied in the signcryption model of [1]. Indeed, if a ciphertext were unforgeable, surely it would be undeniably linkable to its author.

Note also that ciphertext unlinkability only makes sense in a two-layer signcryption model like ours, as opposed to the monolithic model of [27] used in [21, 19]. Indeed, if part of the ciphertext itself is needed to verify the authenticity of the plaintext, ciphertext indistinguishability is lost as soon as the recipient is compelled to prove authenticity to a third party.

3.4 Ciphertext Authentication

Ciphertext authentication is, in a sense, the complement to unlinkability. Authentication requires that the legitimate recipient be able to ascertain that the ciphertext did indeed come from the same person who signed the message it contains. (Naturally, he or she cannot prove this to anyone else, per the unlinkability property.)

We define ciphertext authentication in terms of the following game.

Start The challenger runs the Setup procedure for a given value of the security parameter n, and provides the common public parameters π to the adversary, keeping the secret σ for itself.

Query The adversary makes a number of queries to the challenger, as in the Confidentiality game of §3.1 and the Non-repudiation game of §3.2.

Forgery The adversary returns a recipient identity id_B and and a ciphertext c.

Outcome The adversary wins the game if c decrypts, under the private key of id_B, to a signed message $\langle \mathsf{id}_A, \hat{m}, \hat{s} \rangle$ such that $\mathsf{id}_A \neq \mathsf{id}_B$ and that satisfies $\mathsf{Verify}[\mathsf{id}_A, \hat{m}, \hat{s}] = \top$, provided that (1) no private key extraction query was made on either id_A or id_B, and (2) c did not result from a signature/encryption query with sender and recipient identities id_A and id_B.

We contrast the above experiment, which is a case of 'outsider' security for authentication on the whole ciphertext, with the scenario for signature non-repudiation, which required insider security on the signed plaintext only. We call the above experiment an AUTH-IBSE-CMA attack.

Definition 5. An IBSE scheme is said to be existentially ciphertext-unforgeable against chosen-message outsider attacks, or *AUTH-IBSE-CMA secure*, if no randomized polynomial-time adversary has a non-negligible advantage in the above game. In other words, any randomized polynomial-time EUF-IBSE-CMA adversary $\mathcal{A}$ has an advantage $\mathbf{Adv}_{\mathcal{A}}[n] = \mathbf{P}[\mathsf{Verify}[\mathsf{id}_A, \hat{m}, \hat{s}] = \top]$ that behaves as $o[1/\mathrm{poly}[n]]$ for any polynomial $\mathrm{poly}[n]$.

3.5 Ciphertext Anonymity

Finally, we require ciphertext anonymity, which is to say that the ciphertext must contain no information in the clear that identifies the author or recipient of the message (and yet be decipherable by the intended recipient without that information).

Ciphertext anonymity against adaptive chosen-ciphertext attacks is defined as follows.

Start The challenger runs the Setup procedure for a given value of the security parameter n, and provides the common public parameters π to the adversary, keeping the secret σ for itself.

Phase 1 The adversary is allowed to make adaptive queries of the same types as in the Confidentiality game of §3.1, *i.e.*: signature/encryption queries, decryption queries, and private key extraction queries.

Selection At some point, the adversary returns a message m, two sender identities id_{A_0} and id_{A_1}, and two recipient identities id_{B_0} and id_{B_1}, on which it wishes to be challenged. The adversary must have made no private key extraction query on either id_{B_0} or id_{B_1}.

Challenge The challenger flips two random coins $b', b'' \in \{0, 1\}$, computes $\mathsf{pvk} = \mathsf{Extract}[\mathsf{id}_{A_{b'}}]$, $\langle s, r \rangle \leftarrow \mathsf{Sign}[\mathsf{pvk}, m]$, $c \leftarrow \mathsf{Encrypt}[\mathsf{pvk}, \mathsf{id}_{B_{b''}}, m, s, r]$, and gives the ciphertext c to the adversary.

Phase 2 The adversary adaptatively issues a number of additional encryption, decryption, and extraction queries, under the additional constraint that it not ask for the private key of either id_{B_0} or id_{B_1}, or the decryption of c under id_{B_0} or id_{B_1}.

Response The adversary returns two guesses $\hat{b}', \hat{b}'' \in \{0,1\}$, and wins the game if $\langle \hat{b}', \hat{b}'' \rangle = \langle b', b'' \rangle$.

This game is the same as for confidentiality, except that the adversary is challenged on the identities instead of the message; it is an insider attack. We call it an ANON-IBSE-CCA attack.

Definition 6. An IBSE is said to be ciphertext-anonymous against adaptive chosen-ciphertext insider attacks, or *ANON-IBSE-CCA secure*, if no randomized polynomial-time adversary has a non-negligible advantage in the above game. In other words, any randomized polynomial-time ANON-IBSE-CCA adversary $\mathcal{A}$ has an advantage $\mathbf{Adv}_{\mathcal{A}}[n] = |\mathbf{P}[\hat{b} = b] - \frac{1}{4}|$ that is $o[1/\text{poly}[n]]$ for any polynomial $\text{poly}[n]$ in the security parameter.

We emphasize that anonymity only applies to the ciphertext, against non-recipients, and is thus consistent with both non-repudiation (§3.2) and authentication (§3.4). To illustrate the difference bewteen unlinkability and anonymity, note that the authenticated IBE scheme of [20] is unlinkable but not anonymous, since the sender identity must be known prior to decryption.

4 Review of IB Cryptography from Pairings

We now give a brief summary of the Boneh-Franklin algorithm for identity-based cryptography based on bilinear pairings on elliptic curves.

Let $\mathbb{G}_1$ and $\mathbb{G}_2$ be two cyclic groups of prime order p, writing the group action multiplicatively (in both cases using 1 to denote the neutral element).

Definition 7. An (efficiently computable, non-degenerate) map $\mathbf{e} : \mathbb{G}_1 \times \mathbb{G}_1 \to \mathbb{G}_2$ is called a *bilinear pairing* if, for all $x, y \in \mathbb{G}_1$ and all $a, b \in \mathbb{Z}$, we have $\mathbf{e}[x^a, y^b] = \mathbf{e}[x, y]^{a\,b}$.

Definition 8. The (computational) *bilinear Diffie-Hellman problem* for a bilinear pairing as above is described as follows: given $g, g^a, g^b, g^c \in \mathbb{G}_1$, where g is a generator and $a, b, c \in \mathbb{F}_p^\star$ are chosen at random, compute $\mathbf{e}[g, g]^{a\,b\,c}$. The advantage of an algorithm $\mathcal{B}$ at solving the BDH problem is defined as $\mathbf{Adv}_{\mathcal{B}}[\mathbf{e}] = \mathbf{P}[\mathcal{B}[g, g^a, g^b, g^c] = \mathbf{e}[g, g]^{a\,b\,c}]$.

Definition 9. Let $\mathcal{G}$ be a polynomial-time randomized function that, on input 1^n, returns the description of a bilinear pairing $\mathbf{e} : \mathbb{G}_1 \times \mathbb{G}_1 \to \mathbb{G}_2$ between two groups $\mathbb{G}_1$ and $\mathbb{G}_2$ of prime order p. A BDH parameter generator $\mathcal{G}$ satisfies the *bilinear Diffie-Hellman assumption* if there is no randomized algorithm $\mathcal{B}$ that solves the BDH problem in time $\mathcal{O}[\text{poly}[n]]$ with advantage $\Omega[1/\text{poly}[n]]$. The probability space is that of the randomly generated parameters $\langle \mathbb{G}_1, \mathbb{G}_2, p, \mathbf{e} \rangle$, the BDH instances $\langle g, g^a, g^b, g^c \rangle$, and the randomized executions of $\mathcal{B}$.

The Boneh-Franklin system provides a concrete realization of the above definitions. It is based on an elliptic-curve implementation of the BDH parameter generator $\mathcal{G}$, which we describe following [2, 14] as recently generalized in [6].

Let $E/\mathbb{F}_q$ be an elliptic curve defined over some ground field $\mathbb{F}_q$ of prime characteristic χ. For any extension degree $r \geq 1$, let $E(\mathbb{F}_{q^r})$ be the group of points in $\{\langle x, y \rangle \in (\mathbb{F}_{q^r})^2\} \cup \{\infty\}$ that satisfy the curve equation over $\mathbb{F}_{q^r}$. Let $\nu = \#E(\mathbb{F}_q)$, the number of points on the curve including ∞. Let p be a prime $\neq \chi$ and $\nmid \chi - 1$, such that $p \mid \nu$ and $p^2 \nmid \nu$. Thus, there exists a subgroup $\mathbb{G}'_1$ of order p in $E(\mathbb{F}_q)$. Let κ be the *embedding degree* of $\mathbb{G}'_1$ in $E(\mathbb{F}_q)$, *i.e.*, the smallest integer ≥ 1 such that $p \mid q^\kappa - 1$, but $p \nmid q^r - 1$ for $1 \leq r \leq \kappa$. Under those conditions, there exist a subgroup $\mathbb{G}''_1$ of order p in $E(\mathbb{F}_{q^\kappa})$, and a subgroup $\mathbb{G}_2$ of order p in the multiplicative group $\mathbb{F}^\star_{q^\kappa}$. For appropriately chosen curves, one can then construct a non-degenerate bilinear map $\mathbf{e} : \mathbb{G}_1 \times \mathbb{G}_1 \to \mathbb{G}_2$ believed to satisfy the BDH assumption, where $\mathbb{G}_1$ is either $\mathbb{G}'_1$ or $\mathbb{G}''_1$.

Specifically, [7] show how to obtain a non-degenerate pairing $\bar{\mathbf{e}} : \mathbb{G}'_1 \times \mathbb{G}''_1 \to \mathbb{G}_2$, based on the Tate or the Weil pairing, which can then be combined with a computable isomorphism $\psi : \mathbb{G}''_1 \to \mathbb{G}'_1$, called the *trace map*, to obtain a suitable bilinear map $\mathbf{e} : \mathbb{G}_1 \times \mathbb{G}_1 \to \mathbb{G}_2$ with $\mathbb{G}_1 = \mathbb{G}''_1$. Alternatively, selected curves afford efficiently computable isomorphisms $\phi : \mathbb{G}'_1 \to \mathbb{G}''_1$, called *distortion maps*, which can be combined with $\bar{\mathbf{e}}$ to yield pairings of the form $\mathbf{e} : \mathbb{G}_1 \times \mathbb{G}_1 \to \mathbb{G}_2$ with $\mathbb{G}_1 = \mathbb{G}'_1$. The benefit of the latter construction is that the elements of $\mathbb{G}'_1$ have more compact representations than those of $\mathbb{G}''_1$.

It is desired that p and q^κ be large enough for the discrete logarithm to be intractable in generic groups of size p and in the multiplicative group $\mathbb{F}^\star_{q^\kappa}$. Most commonly, q is a large prime or power of 2 or 3, and $\log p \geq 160$, $\log q^\kappa \geq 1000$. We refer the reader to [4] for background information, and to [5] and [14] for details on the concrete implementation.

In the sequel, we treat the above notions as abstract mathematical objects satisfying the properties summarized in Definitions 7, 8, and 9.

Based on this setup, the Boneh-Franklin system defines four operations, the first two for setup and key extraction purposes by the PKG, the last two for encryption and decryption purposes. The two PKG functions are recalled below.

bfSetup On input a security parameter $n \in \mathbb{N}$: obtain $\langle \mathbb{G}_1, \mathbb{G}_2, p, \mathbf{e} \rangle \leftarrow \mathcal{G}[1^n]$ from the BDH parameter generator; pick two random elements $g \in \mathbb{G}^\star_1$ and $\sigma \in \mathbb{F}^\star_p$, set $g^\sigma = (g)^\sigma \in \mathbb{G}^\star_1$; and construct the hash function $H_0 : \{0,1\}^* \to \mathbb{G}^\star_1$. Finally, output the common public parameters $\pi = \langle \mathbb{G}_1, \mathbb{G}_2, p, \mathsf{e}, g, g^\sigma, H_0 \rangle$ and the master secret $\sigma = \sigma$.

bfExtract On input $\mathsf{id} \in \{0,1\}^*$: hash the given identity into a public element $i_{\mathsf{id}} = H_0[\mathsf{id}] \in \mathbb{G}^\star_1$, and output $d_{\mathsf{id}} = (i_{\mathsf{id}})^\sigma \in \mathbb{G}^\star_1$ as the private key $\mathsf{pvk}_{\mathsf{id}}$.

5 Encryption-Signature Scheme

We now present an efficient realization of the abstract IBSE specifications of §2.

Table 1 details the six algorithms of our scheme. The Setup and Extract functions are essentially the same as in the original Boneh-Franklin system [5].

Table 1. The IBSE algorithms. The hash functions are modeled as random oracles. The output of H_4 is viewed as a stream that is truncated as dictated by context, *viz.*, $H_4[\mathsf{key}] \oplus \mathsf{data}$ perfoms a length-preserving "one-time pad" encryption or decryption.

Setup On input a security parameter $n \in \mathbb{N}$: establish the Boneh-Franklin parameters $\mathbb{G}_1, \mathbb{G}_2, p, \mathbf{e}, g, g^\sigma, \sigma$ as in bfSetup, and select five hash functions $H_0 : \{0,1\}^* \to \mathbb{G}_1^\star$, $H_1 : \mathbb{G}_1^\star \times \{0,1\}^* \to \mathbb{F}_p^\star$, $H_2 : \mathbb{G}_2^\star \to \{0,1\}^{\lceil \log p \rceil}$, $H_3 : \mathbb{G}_2^\star \to \mathbb{F}_p^\star$, $H_4 : \mathbb{G}_1 \to \{0,1\}^*$; then, output the common public parameters $\langle \mathbb{G}_1, \mathbb{G}_2, p, \mathbf{e}, g, g^\sigma, H_0, H_1, H_2, H_3, H_4 \rangle$ and the master secret σ.

Extract On input $\mathsf{id} \in \{0,1\}^*$: proceed as in bfExtract.

Sign On input the private key d_A of some sender identity id_A, and a message m:
derive $i_A = H_0[\mathsf{id}_A]$ (so $d_A = (i_A)^\sigma$),
pick a random $r \in \mathbb{F}_p^\star$,
let $j = (i_A)^r \in \mathbb{G}_1^\star$,
let $h = H_1[j, m] \in \mathbb{F}_p^\star$,
let $v = (d_A)^{r+h} \in \mathbb{G}_1$,
then, output the signature $\langle j, v \rangle$; also forward $\langle m, r, \mathsf{id}_A, i_A, d_A \rangle$ for further use by Encrypt.

Encrypt On input a recipient identity id_B, and $\langle j, v, m, r, \mathsf{id}_A, i_A, d_A \rangle$ from Sign as above:
derive $i_B = H_0[\mathsf{id}_B]$,
compute $u = \mathbf{e}[d_A, i_B] \in \mathbb{G}_2^\star$,
let $k = H_3[u] \in \mathbb{F}_p^\star$;
set $x = j^k \in \mathbb{G}_1^\star$,
let $w = u^{k\,r} \in \mathbb{G}_2^\star$,
set $y = H_2[w] \oplus v$,
set $z = H_4[v] \oplus \langle \mathsf{id}_A, m \rangle$;
then, output the ciphertext $\langle x, y, z \rangle$.

Decrypt On input a private key d_B for id_B, and an anonymous ciphertext $\langle \hat{x}, \hat{y}, \hat{z} \rangle$:
derive $i_B = H_0[\mathsf{id}_B]$,
compute $\hat{w} = \mathbf{e}[\hat{x}, d_B]$,
recover $\hat{v} = H_2[\hat{w}] \oplus \hat{y}$,
recover $\langle \hat{\mathsf{id}}_A, \hat{m} \rangle = H_4[\hat{v}] \oplus \hat{z}$,
derive $\hat{i}_A = H_0[\hat{\mathsf{id}}_A]$,
compute $\hat{u} = \mathbf{e}[\hat{i}_A, d_B]$,
let $\hat{k} = H_3[\hat{u}]$,
set $\hat{j} = \hat{x}^{\hat{k}^{-1}}$;
then, output the decrypted message $\hat{m}$, the signature $\langle \hat{j}, \hat{v} \rangle$, and the purported identity of the originator $\hat{\mathsf{id}}_A$.

Verify On input a signed message $\langle \hat{m}, \hat{j}, \hat{v} \rangle$ by purported sender identity $\hat{\mathsf{id}}_A$:
derive $\hat{i}_A = H_0[\hat{\mathsf{id}}_A]$,
let $\hat{h} = H_1[\hat{j}, m]$,
check whether $\mathbf{e}[g, \hat{v}] \stackrel{?}{=} \mathbf{e}[g^\sigma, (\hat{i}_A)^{\hat{h}}\, \hat{j}]$;
then, output $\top$ if the equality holds, output $\perp$ otherwise.

Sign and Encrypt implement the IBS of [8], although other randomized signature schemes could be substituted for it. Encrypt and Decrypt are less conventional.

Intuitively, Sign implements a randomized IBS whose signatures comprise a commitment j to some random r chosen by the sender, and a closing v that depends on r and the message m. Encrypt superposes two layers of (expansionless) deterministic encryption. The inner layer encrypts j into x using a minimalist authenticated IBE built from zero-round pairing-based key agreement. The outer layer concurrently determines the value w that encrypts to the same x under a kind of anonymous IBE, derandomized to rely on the entropy already present in x. Then, w is hashed into a one-time pad to encrypt the second half of the signature v, which in turn seeds a one-time pad for the bulk encryption of m.

It is helpful to observe that the exponentiations $\star^r$ and $\star^k$ used in Sign for commitment and in Encrypt for authenticated encryption, as well as the key extraction $\star^\sigma$, and the bilinear pairing $\mathbf{e}[\star, i_B]$ that intervenes in the determination of w, all commute. The legitimate recipient derives its ability to decrypt x from the capacity to perform all of the above operations (either explicitly or implicitly)—but it can only do so in a specific order, different than the sender.

The results of §7 show the scheme is secure. We now prove its consistency.

Theorem 10. *The IBSE scheme of Table 1 is consistent.*

Proof. For decryption, if $\langle \hat{x}, \hat{y}, \hat{z} \rangle = \langle x, y, z \rangle$, it follows that $\hat{w} = \mathbf{e}[{i_A}^{r\,k}, {i_B}^\sigma] = \mathbf{e}[{i_A}^\sigma, i_B]^{r\,k} = w$ (in $\mathbb{G}_2^\star$), and thus $\hat{v} = v$ and $\langle \hat{\mathsf{id}}_A, \hat{m} \rangle = \langle \mathsf{id}_A, m \rangle$; we also have $\hat{u} = \mathbf{e}[\hat{i}_A, i_B]^\sigma = u$ (in $\mathbb{G}_2^\star$), hence $\hat{k} = k$ (in $\mathbb{F}_p^\star$), and thus $\hat{j} = (j^k)^{\hat{k}^{-1}} = j$ (in $\mathbb{G}_1^\star$). For verification, if $\langle \hat{m}, \hat{\mathsf{id}}_A, \hat{j}, \hat{v} \rangle = \langle m, \mathsf{id}_A, j, v \rangle$, we have $\mathbf{e}[g, \hat{v}] = \mathbf{e}[g, i_A]^{\sigma\,(r+h)} = \mathbf{e}[g^\sigma, (\hat{i}_A)^h\,(\hat{i}_A)^r] = \mathbf{e}[g^\sigma, (\hat{i}_A)^h\,\hat{j}]$ (in $\mathbb{G}_2$), as required. □

6 Competitive Performance

Table 2 gives a comparison between various IB encryption and signature schemes, in terms of size, performance, and security properties.

Our comparisons include most relevant pairing-based IB schemes for encryption, authenticated encryption, signature, and signcryption. We also include a suite of hybrid schemes, obtained by combining IBS [8] with either IBE [5] or AuthIBE [20]; each pair is composed in three different ways depending on the order of application of the primitives: encrypt-then-sign ($\mathcal{E}t\mathcal{S}$), sign-then-encrypt ($\mathcal{S}t\mathcal{E}$), and commit-then-parallel-encrypt-and-sign ($\mathcal{C}t\mathcal{S}\&\mathcal{E}$), as per [1]. Roughly speaking, in $\mathcal{C}t\mathcal{S}\&\mathcal{E}$, the plaintext m is reversibly transformed into a redundant pair $\langle a, b \rangle$, where a is a commitment to m that reveals "no information" about m; then, a is signed and b encrypted using the given primitives, in parallel.

For fairness, the size comparison factors out the overhead of explicitly including the sender identity to the signed plaintext prior to encryption; our scheme does this to avoid sending the identity in the clear. Note that all authenticated communication schemes require the recipient to get hold of that information, but most simply assume that it is conveyed using a different channel.

Evidently, the proposed scheme offers an interesting solution to the problem of identity-based signed encryption: it offers an unmatched combination of security features that not only provide the usual confidentiality/non-repudiation requirements, but also guarantee authentication, anonymity, and unlinkability of the ciphertext. Our scheme achieves all this at a cost comparable to that of monolithic IB signcryption, and in a significantly tighter package than any generic combination of existing IB encryption and signature algorithms.

By comparison, the two listed signcryption schemes have comparable spatial and computational overheads but, by the very nature of monolithic signcryption, cannot offer ciphertext anonymity. As for the suite of generic compositions, they have a slight advantage in terms of cost, but incur a large size penalty, and require us to choose between ciphertext authentication and anonymity.

We also note that, in the original Boneh-Franklin setup, the IBSE ciphertexts and signed plaintexts are essentially as compact as that of IBE or IBS taken in

Table 2. Comparison between various IB encryption, signature, signcryption, and multipurpose schemes. Times are expressed as triples $\langle \#b, \#m, \#e \rangle$, where $\#b$ is the number of bilinear pairings, $\#m$ is the number of $\mathbb{G}_1$ exponentiations, and $\#e$ is the number of $\mathbb{G}_2$ or $\mathbb{F}_p$ exponentiations (simple group operations in $\mathbb{G}_1$ and multiplications and inversions in $\mathbb{F}_p$ or $\mathbb{G}_2$ are omitted). Sizes are reported as pairs $\langle \#p, \#q \rangle$, where $\#p$ is the number of $\mathbb{G}_1$ elements, and $\#q$ is the number of $\mathbb{F}_p$ or $\mathbb{G}_2$ elements, in excess of the original unsigned message size $\|m\|$ taken as baseline (treating the sender identity as part of m, if included); the 'cipher' size is the ciphertext overhead, $\|c\| - \|m\|$, while the 'plain' size is the signature overhead after decryption, or $\|\langle m, s \rangle\| - \|m\|$. Security is indicated as follows: message Confidentiality, signature Non-repudiation, and ciphertext Authentication, Unlinkability, and anOnymity; for non-IBSE schemes, an uppercase denotes an analogous security notion, a lowercase a weaker notion.

Scheme	**Security**: Conf, Nrep,Auth,Ulnk,anOn	**Size**:#el.$\mathbb{G}_1$,$\mathbb{G}_2$+$\mathbb{F}_p$ Cipher	Plain	**Time**:#pair.,exp.$\mathbb{G}_1$,$\mathbb{G}_2$+$\mathbb{F}_p$ Sign	Encrypt	Decrypt	Verify
IB Encryption [5]	C,–, –, U,O	$1,1$	—	—	$1,0,0$	$1,1,0$	—
IB Auth.Encr. [20]	C,–, A,U,–	$0,2$	—	—	$1,0,0$	$1,0,0$	—
IB Signature [8]	–, N,A,–, –	—	$2,0$	$0,2,0$	—	—	$2,1,0$
[a] IB Signature [22]	–, N,A,–, –	—	$2,0$	$0,4,0$	—	—	$2(3),0,2$
IB Sign. [16, #3]	–, N,A,–, –	—	$1,1$	$1,2,1$	—	—	$2,0,1$
IB Sign. [16, #4]	–, N,A,–, –	—	$2,0$	$0,2,0$	—	—	$2,0,1$
[b] IB SignCrypt. [21]	*, N,A,–, –	$2,0$	$2,0$	$\cdots$ $1,3,0$	$\cdots$	$\cdots$ $4,0,1$	$\cdots$
IB SignCrypt. [19]	C,N,A,–, –	$1,1$	$1,1$	$\cdots$ $2,2,2$	$\cdots$	$\cdots$ $4,0,2$	$\cdots$
[c] IB E–then–S	c, N,A,–, –	$3,1$	$2,0$	$0,2,0$	$1,0,0$	$1,1,0$	$2,1,0$
[c] IB S–then–E	C,n, –, U,O	$3,1$	$2,0$	$0,2,0$	$1,0,0$	$1,1,0$	$2,1,0$
[c] IB commit–E&S	C,N,A,U,–	$3,1,+$	$2,0,+$	$0,2,0$	$1,0,0$	$1,1,0$	$2,1,0$
[d] IB AE–then–S	c, N,A,U,–	$2,2$	$2,0$	$0,2,0$	$1,0,0$	$1,0,0$	$2,1,0$
[d] IB S–then–AE	C,n, A,U,–	$2,2$	$2,0$	$0,2,0$	$1,0,0$	$1,0,0$	$2,1,0$
[d] IB commit–AE&S	C,N,A,U,–	$2,2,+$	$2,0,+$	$0,2,0$	$1,0,0$	$1,0,0$	$2,1,0$
IBSE: this paper	C,N,A,U,O	$2,0$	$2,0$	$0,2,0$	$1,0,2$	$2,1,0$	$2,1,0$

[a] Signature verification in [22] requires 3 pairings, one of which may be precomputed.
[b] The signcryption scheme of [21] is not adaptive CCA-secure, see [19] for details.
[c] These are compositions of IBE [5] and IBS [8, 16] using $\mathcal{E}t\mathcal{S}$, $\mathcal{S}t\mathcal{E}$, $\mathcal{C}t\mathcal{S}\&\mathcal{E}$ from [1].
[d] These are compositions of AuthIBE [20] and IBS [8, 16] using $\mathcal{E}t\mathcal{S}$, $\mathcal{S}t\mathcal{E}$, $\mathcal{C}t\mathcal{S}\&\mathcal{E}$ [1].
$\mathcal{E}t\mathcal{S}$ and $\mathcal{S}t\mathcal{E}$ respectively degrade the CCA indistinguishability and CMA unforgeability of its constituents in the insider model; the '+' are a reminder that the more secure $\mathcal{C}t\mathcal{S}\&\mathcal{E}$ incurs extra overhead due to the commitment redundancy. See [1] for details.

isolation; this is generally true when $p \approx q$, and when $\mathbb{G}_1 = \mathbb{G}_1'$ so that its points can be represented as elements of $\mathbb{F}_q$ using point compression [4]. However, the schemes of [5], [19], and especially [20] have smaller ciphertexts and signatures, as the case may be, in generalized setups where $p \ll q$, or $\mathbb{G}_1 = \mathbb{G}_1''$.

7 Security Analysis

We now state our security results for the scheme of §5 in the models of §3.

Theorem 11. *Let $\mathcal{A}$ be a polynomial-time IND-IBSE-CCA attacker that has advantage $\geq \epsilon$, and makes $\leq \mu_i$ queries to the random oracles H_i, $i = 0, 1, 2, 3, 4$. Then, there exists a polynomial-time algorithm $\mathcal{B}$ that solves the bilinear Diffie-Hellman problem with advantage $\geq \epsilon/(\mu_0\, \mu_2)$.*

Theorem 12. *Let $\mathcal{A}$ be an EUF-IBSE-CCA attacker that makes $\leq \mu_i$ queries to the random oracles H_i, $i = 0, 1, 2, 3, 4$, and $\leq \mu_{\mathrm{se}}$ queries to the signature/encryption oracle. Assume that, within a time span $\leq \tau$, $\mathcal{A}$ produces a successful forgery with probability $\geq \epsilon = 10\,(\mu_{\mathrm{se}} + 1)\,(\mu_{\mathrm{se}} + \mu_1)/2^n$, for a security parameter n. Then, there exists an algorithm $\mathcal{B}$ that solves the bilinear Diffie-Hellman problem in expected time $\leq 120686\, \mu_0\, \mu_1\, \tau/\epsilon$.*

Theorem 13. *There exists a polynomial-time algorithm that, given an identifier id_A, a signed plaintext $\langle m, j, v\rangle$ from id_A, and a private key d_B, creates a ciphertext $\langle x, y, z\rangle$ that decrypts to $\langle m, j, v\rangle$ under d_B, with probability 1.*

Theorem 14. *Let $\mathcal{A}$ be a polynomial-time AUTH-IBSE-CMA attacker with advantage $\geq \epsilon$, that makes $\leq \mu_i$ queries to the random oracles H_i, $i = 0, 1, 2, 3, 4$. Then, there exists a polynomial-time algorithm $\mathcal{B}$ that solves the bilinear Diffie-Hellman problem with advantage $\geq 2\,\epsilon/(\mu_0\,(\mu_0 - 1)\,(\mu_1\, \mu_2 + \mu_3))$.*

Theorem 15. *Let $\mathcal{A}$ be a polynomial-time ANON-IBSE-CCA attacker that has advantage $\geq \epsilon$, and makes $\leq \mu_i$ queries to the random oracles H_i, $i = 0, 1, 2, 3, 4$. Then, there exists a polynomial-time algorithm $\mathcal{B}$ that solves the bilinear Diffie-Hellman problem with advantage $\geq 3\,\epsilon/(\mu_0\,(\mu_0 - 1)\,(\mu_1\, \mu_2 + 2\,\mu_2 + \mu_3))$.*

8 Practical Extensions

We now mention a few straightforward generalizations of practical interest.

8.1 Encrypting for Multiple Recipients

Encrypting the same message m for a set of n recipients $\mathsf{id}_{B_1}, ..., \mathsf{id}_{B_n}$ is easily achieved as follows. The Sign operation is carried out once (which establishes the randomization parameter r), then the $\mathsf{Encrypt}$ operation is performed independently for each recipient, based on the output from Sign.

Since the message m and the randomization parameter r are invariant for all the $\mathsf{Encrypt}$ instances, it is easy to see that the z component of the ciphertext also remains the same. Thus, the multi-recipient composite ciphertext is easily assembled from one instance of $\langle x_i, y_i\rangle \in \mathbb{G}_1^\star \times \mathbb{G}_1^\star$ for each recipient B_i, plus a single instance of $z \in \{0,1\}^*$ to be shared by all. Thus, a multi-recipient ciphertext is compactly encoded in the form $c = \langle\langle x_1, y_1\rangle, ..., \langle x_n, y_n\rangle, z\rangle$.

The security models of §3 have to support two additional types of queries: *multi-recipient signature/encryption queries*, in which a given message, sender, and list of recipients, are turned into a multi-recipient ciphertext, and *multi-recipient decryption queries*, in which the individual elements of a multi-recipient ciphertext are decrypted, under a given identity, and a valid plaintext is returned, if there is any. The modified security analysis is deferred to the full paper.

8.2 Integrity Without Non-Repudiation

The scheme of §5 is trivially modified to provide message integrity without non-repudiation or authentication. To do this, the sender merely substitutes the public parameters $\langle g, g^\sigma \rangle$ for $\langle i_A, d_A \rangle$, wherever the sender's key pair is used in the Sign and Encrypt operations. The sender also tags the message as 'anonymous', instead of specifying an identity. Similarly, the Decrypt and Verify operations are performed substituting g^σ for $\hat{i}_A$ wherever it appears as a function argument.

This is valid since the key pair relation $d_A = (i_A)^\sigma$ is paralleled by $g^\sigma = (g)^\sigma$, but authentication is meaningless since the signing 'private' key g^σ is public.

9 Conclusion

In this paper, we have proposed a comprehensive security model for multi-purpose identity-based encryption-signature cryptosystems. Our security model defines five core properties that we believe precisely capture what a consumer of cryptography intuitively expects when he or she wishes to engage in "secure signed communication" with a remote party. It bears repeating that these do not only include the standard confidentiality and non-repudiation requirements, but also the much less commonly offered features of ciphertext authentication, ciphertext deniability or unlinkability, and true ciphertext anonymity with respect to third parties. We have given precise definitions for all these properties in the context of identity-based cryptography.

As second contribution, we have presented a new cryptographic scheme that precisely implements all facets of the aforementioned notion of "secure signed communication", in the certificate-free world of identity-based cryptography. Our scheme offers efficient security bounds in all the above respects; it is fast, compact, scalable, and practical—as we have illustrated through detailed comparisons with most or all mainstream identity-based cryptosystems to date.

Acknowledgements

The author would like to thank Dan Boneh, Jonathan Katz, and the anonymous referees of Crypto 2003 for many helpful suggestions and comments. Credit goes to Guido Appenzeller for suggesting the "Swiss Army Knife" moniker.

References

1. J.H. An, Y. Dodis, and T. Rabin. On the security of joint signature and encryption. In *Proc. Eurocrypt '02, LNCS 2332*, 2002.
2. P.S.L.M. Barreto, H.Y. Kim, B. Lynn, and M. Scott. Efficient algorithms for pairing-based cryptosystems. In *Proc. Crypto '02, LNCS 2442*, 2002.
3. M. Bellare and P. Rogaway. Random oracles are practical: A paradigm for designing efficient protocols. In *Proc. Conf. Computer and Communication Security*, 1993.

4. I. Blake, G. Seroussi, and N. Smart. *Elliptic Curves in Cryptography*. Cambridge University Press, 1999.
5. D. Boneh and M. Franklin. Identity-based encryption from the Weil pairing. In *Proc. Crypto '01, LNCS 2139*, pages 213–229, 2001. See [6] for the full version.
6. D. Boneh and M. Franklin. Identity based encryption from the weil pairing. Cryptology ePrint Archive, Report 2001/090, 2001. http://eprint.iacr.org/.
7. D. Boneh, B. Lynn, and H. Shacham. Short signatures from the Weil pairing. In *Proc. Asiacrypt '01, LNCS 2248*, pages 514–532, 2001.
8. J.C. Cha and J.H. Cheon. An identity-based signature from gap Diffie-Hellman groups. Cryptology ePrint Archive, Report 2002/018, 2002. http://eprint.iacr.org/.
9. L. Chen and C. Kudla. Identity based authenticated key agreement from pairings. Cryptology ePrint Archive, Report 2002/184, 2002. http://eprint.iacr.org/.
10. C. Cocks. An identity based encryption scheme based on quadratic residues. In *Proc. 8th IMA Int. Conf. Cryptography and Coding*, pages 26–28, 2001.
11. U. Feige, A. Fiat, and A. Shamir. A digital signature scheme secure against adaptive chosen-message attacks. *SIAM J. Computing*, 17(2):281–308, 1988.
12. U. Feige, A. Fiat, and A. Shamir. Zero-knowledge proofs of identity. *J. Cryptology*, 1:77–94, 1988.
13. A. Fiat and A. Shamir. How to prove yourself: Practical solutions to identification and signature problems. In *Proc. Crypto '86, LNCS 263*, pages 186–194, 1984.
14. S.D. Galbraith, K. Harrison, and D. Soldera. Implementing the Tate pairing. Technical Report HPL-2002-23, HP Laboratories Bristol, 2002.
15. C. Gentry and A. Silverberg. Hierarchical ID-based cryptography. Cryptology ePrint Archive, Report 2002/056, 2002. http://eprint.iacr.org/.
16. F. Hess. Exponent group signature schemes and efficient identity based signature schemes based on pairings. Cryptology ePrint Archive, Report 2002/012, 2002. http://eprint.iacr.org/.
17. J. Horwitz and B. Lynn. Toward hierarchical identity-based encryption. In *Proc. Eurocrypt '02, LNCS 2332*, pages 466–481, 2002.
18. A. Joux. A one round protocol for tripartite Diffie-Hellman. In *Proc. 4th Alg. Numb. Th. Symp., LNCS 1838*, pages 385–294, 2000.
19. B. Libert and J.-J. Quisquater. New identity based signcryption schemes based on pairings. Cryptology ePrint Archive, Report 2003/023, 2003. http://eprint.iacr.org/.
20. B. Lynn. Authenticated identity-based encryption. Cryptology ePrint Archive, Report 2002/072, 2002. http://eprint.iacr.org/.
21. J. Malone-Lee. Identity-based signcryption. Cryptology ePrint Archive, Report 2002/098, 2002. http://eprint.iacr.org/.
22. K.G. Paterson. ID-based signatures from pairings on elliptic curves. Cryptology ePrint Archive, Report 2002/004, 2002. http://eprint.iacr.org/.
23. R. Sakai, K. Ohgishi, and M. Kasahara. Cryptosystems based on pairings. In *Proc. SCIS '00*, pages 26–28, Okinawa, Japan, 2000.
24. A. Shamir. Identity-based cryptosystems and signature schemes. In *Proc. Crypto '84, LNCS 196*, pages 47–53, 1984.
25. N.P. Smart. An identity based authenticated key agreement protocol based on the Weil pairing. Cryptology ePrint Archive, Report 2001/111, 2001. http://eprint.iacr.org/.
26. D. Pointcheval J. Stern. Security arguments for digital signatures and blind signatures. *J. Cryptology*, 13:361–396, 2000.
27. Y. Zheng. Digital signcryption or how to achieve $cost\,(signature\,\&\,encryption) \ll cost\,(signature) + cost\,(encryption)$. In *Proc. Crypto '97, LNCS 1294*, 1997.