Generalized Key Delegation for Hierarchical Identity-Based Encryption

Michel Abdalla¹, Eike Kiltz², and Gregory Neven^{1,3}

¹ Département d'Informatique, École normale supérieure

Michel.Abdalla@ens.fr

<http://www.di.ens.fr/~mabdalla>

² Cryptology and Information Security Research Theme, CWI Amsterdam

kiltz@cw.nl

<http://kiltz.net>

³ Department of Electrical Engineering, Katholieke Universiteit Leuven

Gregory.Neven@esat.kuleuven.be

<http://www.neven.org>

Abstract. In this paper, we introduce a new primitive called identity-based encryption with wildcard key derivation (WKD-IBE, or “wicked IBE”) that enhances the concept of hierarchical identity-based encryption (HIBE) by allowing more general key delegation patterns. A secret key is derived for a vector of identity strings, where entries can be left blank using a wildcard. This key can then be used to derive keys for any pattern that replaces wildcards with concrete identity strings. For example, one may want to allow the university’s head system administrator to derive secret keys (and hence the ability to decrypt) for all departmental sysadmin email addresses `sysadmin@*.univ.edu`, where `*` is a wildcard that can be replaced with any string. We provide appropriate security notions and provably secure instantiations with different tradeoffs in terms of ciphertext size and efficiency. We also present a generic construction of identity-based broadcast encryption (IBBE) from any WKD-IBE scheme. One of our instantiation yields an IBBE scheme with constant ciphertext size.

1 Introduction

IDENTITY-BASED ENCRYPTION. Securely linking users to their public keys is a notorious obstacle in the adoption of public-key encryption schemes in practice. Most commonly, it is overcome by means of a public key infrastructure (PKI) where a trusted authority certifies, by means of a digital signature, the relation between users and their public keys. The high cost of setting up and maintaining such a PKI can be prohibitive for many organizations however. In 1984, Shamir [20] proposed identity-based encryption (IBE) as a cheaper alternative to traditional PKIs. Here, the public key of a user *is* his identity (e.g. his name or email address), while the corresponding private key is handed to him by a trusted key distribution center. It lasted until 2000 however for the first practical IBE schemes [18,7] to be proposed based on bilinear maps.

Hierarchical identity-based encryption (HIBE) schemes [14,12] are the hierarchical extension of IBEs where user identities are vectors of bit strings. The root entity generates private keys for users at the first level; users at level ℓ can derive keys for their children at level $\ell + 1$. This prevents the distribution center from becoming a bottleneck in the system, and at the same time reflects the hierarchical structure of many organizations and user identities, in particular email addresses. For example, the head of the computer science department of a university could be given the key for identity `(edu,univ,cs)` allowing him to derive keys for identities `(edu,univ,cs,username)` corresponding to email addresses `username@cs.univ.edu`.

WILDCARD KEY DERIVATION. Hierarchical key derivation is a useful feature, but has its limitations. For example, it would be reasonable to prevent end-users from further deriving keys for identities below them. This feature was referred to before as *limited delegation* by Boneh-Boyen-Goh [6], who show a tweak to their HIBE scheme offering exactly this functionality—albeit without a formal security notion or proof for their approach. In some circumstances, it could also be useful to be able to deviate from the hierarchical structure. For example, one may want to allow the university’s head system administrator to derive keys for all departmental sysadmin email addresses `sysadmin@*.univ.edu`, where `*` is a wildcard that can be replaced with any string. As another example, it could be practical to provide a company like Google Inc. that registers its name at all top-level domains with a key for `*@google.*`.

These applications lead us to generalize the concept of HIBE schemes to *identity-based encryption with wildcard key derivation* (WKD-IBE), or more succinctly *wicked IBE*. After defining adequate security notions, we start looking for constructions. First observe that if a HIBE scheme allows a maximal hierarchy depth L to be fixed, then the limited-delegation property of [6] can be achieved generically by padding the identity vector with “dummy” strings at the unused lower levels. (But this may come at the cost of efficiency.) The more general functionality of wildcard key delegation cannot be achieved generically though. Nevertheless, we show that many of the existing HIBE schemes are amenable to a modification that enables wildcard key derivation, including the Gentry-Silverberg [12], Boneh-Boyen [5], Waters [21], and Boneh-Boyen-Goh [6] HIBE schemes. For the former three this may come as a bit of a surprise, because no limited-delegation tweaks were previously proposed for these schemes. We prove the security of the modified schemes under our new notions, thereby providing as a special case formal ground for the intuition of [6] regarding their limited-delegation tweak.

APPLICATION TO IDENTITY-BASED BROADCAST ENCRYPTION. Broadcast encryption [11] allows to encrypt a message to any subset $S \subseteq \{1, \dots, N\}$ of N users so that only users in S can decrypt the message. A trivial solution consists of concatenating encryptions of the message under the public key of each user in S separately, but this yields ciphertexts of size linear in $|S|$. The most efficient fully collusion-resistant (meaning where the adversary can corrupt all users outside of S) public-key broadcast encryption schemes are due to Boneh et al. [8],

who present a first construction with constant-size ciphertexts and private keys but with $O(N)$ -size public keys, and a second construction with $O(\sqrt{N})$ -size ciphertexts and public keys.

Identity-based broadcast encryption (IBBE) is the natural extension of broadcast encryption to the identity-based setting. It is particularly appealing as a primitive because the total number of users in the system N is limited only by the size of the identity space. We propose a generic construction of an IBBE schemes from any WKD-IBE scheme. The construction inflates the private key size by a factor L being the maximal number of identities in a recipient set, but otherwise shares the same cost as the underlying wicked IBE.

Of all the instantiations of wicked IBE that we propose, the most attractive resulting IBBE scheme is that obtained from the scheme based on [6], because it achieves constant-size ciphertexts. However, it has the disadvantage of having private keys of size $O(L^2)$, where L is the maximum number of recipients in a ciphertext. The other concrete instantiations are less attractive because they have ciphertext size $O(L)$, just like the trivial scheme that concatenates individual ciphertexts. Unlike most other broadcast schemes however, they do have the remarkable feature that knowledge of the recipient set is not required in order to decrypt the message.

WILDCARD SIGNATURES. Just like the key derivation of an IBE scheme automatically gives rise to a signature scheme [7], a WKD-IBE scheme gives rise to a new primitive that we call a *wildcard signature scheme*. It allows a signer to issue a signature on a message containing wildcards, which anyone can replace with concrete values at a later point without invalidating the signature. Our constructions of wicked identity-based encryption yield a number of wildcard signature schemes with different tradeoffs.

RELATED WORK. Wicked identity-based encryption can be seen as the dual notion of identity-based encryption with wildcards [1] (WIBE). There, one can use wildcards in the recipient identity to which a ciphertext is encrypted, so that all users whose identity matches the recipient pattern can decrypt it. In fact, the notions of WKD-IBE and WIBE could be combined into a universal primitive that allows wildcards to be used in both the encryption and key derivation algorithms. Instantiations of this primitive can be obtained from all WKD-IBE schemes presented in this work, except for the one based on Gentry-Silverberg's HIBE [12].

Key-policy attribute-based encryption (KP-ABE) [13] associates to each decryption key an access structure consisting of a logical combination of attribute values using AND and OR gates. A ciphertext is encrypted under a set of descriptive attributes and can only be decrypted with a key whose access structure is satisfied by the set of attributes. As discussed in [13], HIBE schemes can be seen as a special case of KP-ABE schemes by mapping the identity vector (`edu`, `univ`, `cs`, `sysadmin`) to the access structure ($1||\text{edu} \wedge 2||\text{univ} \wedge 3||\text{cs} \wedge 4||\text{sysadmin}$). Likewise, wicked IBE can be seen as a special case of KP-ABE by letting the key for identity (`edu`, `*`, `*`, `sysadmin`) be given by the key for ($1||\text{edu} \wedge 4||\text{sysadmin}$). The wicked IBE scheme obtained through the first

construction of [13] has the disadvantage of having public keys linear in the size of the attribute universe. The instantiation obtained from their second, large-universe construction is quite similar to the scheme that we derive from the Boneh-Boyen HIBE scheme [5]. None of the schemes derived from [13] achieve constant ciphertext size though, like our wicked IBE construction based on [6].

The use of HIBE schemes in the design of broadcast encryption schemes was first considered by Dodis and Fazio [10]. Chatterjee and Sarkar [9] gave a direct construction of an IBBE scheme that is closely related to the instantiation of our generic construction with the WKD-IBE scheme based on [6]. Our generic construction provides insight into the design of their scheme, but their construction contains some interesting efficiency-improving tweaks. The schemes are compared in more detail in Section 4.3.

In independent work, Shacham [19] formalizes the concept of *limited delegation* for HIBE schemes and proves this feature for the HIBE scheme of [6]. As we pointed out above, limited delegation for HIBEs can be seen as a special case of WKD-IBE where wildcards can only appear at the end of the identity vector. Our WKD-IBE scheme based on [6] can therefore be seen as a generalization of the result of [19].

2 Basic Definitions

In this section, we introduce some notation and computational problems that we will use throughout the rest of the paper. In doing so, we adopt the same notation and definition style used in [1].

NOTATION. Let $\mathbb{N} = \{0, 1, \dots\}$ be the set of natural numbers. Let ε be the empty string. If $n \in \mathbb{N}$, then $\{0, 1\}^n$ denotes the set of n -bit strings, and $\{0, 1\}^*$ is the set of all bit strings. More generally, if S is a set, then S^n is the set of n -tuples of elements of S , $S^{\leq n}$ is the set of tuples of length at most n . If S is finite, then $x \xleftarrow{\$} S$ denotes the assignment to x of an element chosen uniformly at random from S . If A is an algorithm, then $y \leftarrow A(x)$ denotes the assignment to y of the output of A on input x , and if A is randomized, then $y \xleftarrow{\$} A(x)$ denotes that the output of an execution of $A(x)$ with fresh coins is assigned to y .

THE DECISIONAL BILINEAR DIFFIE-HELLMAN ASSUMPTION [7]. Let $\mathbb{G}, \mathbb{G}_T$ be multiplicative groups of prime order p with an admissible map $\hat{e} : \mathbb{G} \times \mathbb{G} \rightarrow \mathbb{G}_T$. By admissible we mean that the map is bilinear, non-degenerate and efficiently computable. Bilinearity means that for all $a, b \in \mathbb{Z}_p$ and all $g \in \mathbb{G}$ we have $\hat{e}(g^a, g^b) = \hat{e}(g, g)^{ab}$. By non-degenerate we mean that $\hat{e}(g, g) = 1$ if and only if $g = 1$. Let $g \in \mathbb{G}$ be a generator. In such a setting, the bilinear decisional Diffie-Hellman (BDDH) problem is to determine, given g , $A = g^a$, $B = g^b$, $C = g^c$, and $Z = \hat{e}(g, g)^z$, whether $Z = \hat{e}(g, g)^{abc}$ for hidden values of a, b, c and z . More formally, let $\mathcal{A}$ be an adversary for the BDDH problem. Such an adversary has advantage ϵ in solving the BDDH problem if $|\Pr[\mathcal{A}(g, A, B, C, \hat{e}(g, g)^{abc}) = 1] - \Pr[\mathcal{A}(g, A, B, C, \hat{e}(g, g)^z) = 1]| \geq \epsilon$, where the probabilities are over the choice of a, b, c, z and over the random coins consumed by $\mathcal{A}$.

Definition 1. *The (t, ϵ) -BDDH assumption holds if no t -time adversary has at least ϵ advantage in the above game.*

We note that throughout this paper we will assume that the time t of an adversary includes its code size, in order to exclude trivial “lookup” adversaries.

THE BDHE ASSUMPTION. The ℓ -decisional bilinear Diffie-Hellman exponent (ℓ -BDHE) problem [6] in $\mathbb{G}$ is: given g, h and $g^{(\alpha^i)} \in \mathbb{G}$, for $i = 1, \dots, \ell-1, \ell+1, \dots, 2\ell$ as input, output $\hat{e}(g, h)^{(\alpha^\ell)} \in \mathbb{G}_T$. Boneh, Boyen and Goh, conjectured that the ℓ -BDHE is a hard problem, meaning with this that no polynomially bounded adversary $\mathcal{A}$ can solve it with more than negligible probability, over the random choices of $g, h \in \mathbb{G}$, the choice of $\alpha \in \mathbb{Z}_p$, and the random coin tosses of $\mathcal{A}$.

The decisional version of the problem can be defined in the usual manner. Let $\mathbf{y} = (g^\alpha, g^{(\alpha^2)}, \dots, g^{(\alpha^{\ell-1})}, g^{(\alpha^{\ell+1})}, \dots, g^{(\alpha^{2\ell})})$. An algorithm $\mathcal{B}$ that outputs a bit b , has advantage ϵ in solving the decisional ℓ -BDHE problem in $\mathbb{G}$ if $\left| \Pr [\mathcal{B}(g, h, \mathbf{y}, \hat{e}(g, h)^{(\alpha^\ell)}) = 1] - \Pr [\mathcal{B}(g, h, \mathbf{y}, T) = 1] \right| \geq \epsilon$, where the probabilities are taken over the random choices of $g, h \in \mathbb{G}$, the random choice of $\alpha \in \mathbb{Z}_p$, the random choice of $T \in \mathbb{G}_T$, and the internal coin tosses of $\mathcal{B}$.

Definition 2. *The decisional (t, ϵ, ℓ) -BDHE assumption holds in $\mathbb{G}$ if no t -time (probabilistic) algorithm has advantage at least ϵ in solving the decisional ℓ -BDHE problem in $\mathbb{G}$.*

3 Wicked Identity-Based Encryption

SYNTAX. A wicked identity-based encryption scheme (WKD-IBE) is a generalization of a HIBE scheme which allows for more general key delegation patterns. In a WKD-IBE scheme, secret keys are associated with patterns rather than identity vectors. A pattern P is a vector $(P_1, \dots, P_\ell) \in (\{0, 1\}^* \cup \{\ast\})^\ell$ of length $\ell \leq L$, where $\ast$ is a special wildcard symbol and L is the maximal depth of the WKD-IBE scheme. That is, each component of a pattern P is either a specific identity string or a wildcard. The main idea behind the WKD-IBE notion is that a user in possession of the secret key for a given pattern P can generate secret keys for any pattern P' that matches P . We say that a pattern $P' = (P'_1, \dots, P'_{\ell'})$ matches P , denoted $P' \in_\ast P$, if and only if $\ell' \leq \ell$; $\forall i = 1 \dots \ell'$, $P'_i = P_i$ or $P_i = \ast$; and $\forall i = \ell' + 1 \dots \ell$, $P_i = \ast$.

More formally, a WKD-IBE scheme is a tuple of algorithms $\mathcal{WKD}\text{-IBE} = (\text{Setup}, \text{KeyDer}, \text{Enc}, \text{Dec})$ providing the following functionality. The root authority generates a master key pair $(mpk, msk) \xleftarrow{\$} \text{Setup}$. Via $sk_{P'} \xleftarrow{\$} \text{KeyDer}(sk_P, P')$, a user possessing the secret key sk_P for a pattern $P = (P_1, \dots, P_\ell)$ can derive a secret key for any pattern $P' \in_\ast P$. The secret key of the root identity is $msk = sk_{(\ast, \dots, \ast)}$.

To create a ciphertext of message $m \in \{0, 1\}^*$ intended for an identity $ID = (ID_1, \dots, ID_\ell)$, the sender computes $C \xleftarrow{\$} \text{Enc}(mpk, ID, m)$. Any user in possession of the secret key for a pattern P such that $ID \in_\ast P$ can decrypt the ciphertext using sk_P as $m \leftarrow \text{Dec}(sk_P, C, ID)$. Correctness requires that for all key pairs

(mpk, msk) output by **Setup**, all messages $m \in \{0, 1\}^*$, all $0 \leq \ell \leq L$, all patterns $P \in (\{0, 1\}^* \cup \{*\})^\ell$, and all identities $ID \in (\{0, 1\}^*)^{\ell'}$ such that $ID \in_* P$, $\text{Dec}(\text{KeyDer}(msk, P), \text{Enc}(mpk, ID, m), ID) = m$ with probability one.

SECURITY. We define the security of WKD-IBE schemes in a way that is very similar to the case of HIBE schemes, but where the adversary can query for the secret keys corresponding to arbitrary patterns, rather than specific identity vectors. Of course, the adversary is not allowed to query the key derivation oracle for any pattern matched by the challenge identity.

More specifically, security is defined through the following game with an adversary. In the first phase, the adversary is run on input of the master public key of a freshly generated key pair $(mpk, msk) \xleftarrow{\$} \text{Setup}$. In a chosen-plaintext attack (IND-WKID-CPA), the adversary is given access to a key derivation oracle that on input a pattern $P \in (\{0, 1\}^* \cup \{*\})^{\leq L}$ returns $sk_P \xleftarrow{\$} \text{KeyDer}(msk, P)$.

At the end of the first phase, the adversary outputs two equal-length challenge messages $m_0^*, m_1^* \in \{0, 1\}^*$ and a challenge identity $ID^* = (ID_1^*, \dots, ID_{\ell^*}^*)$ where $0 \leq \ell^* \leq L$. The adversary is given a challenge ciphertext $C^* \xleftarrow{\$} \text{Enc}(mpk, ID^*, m_b^*)$ for a randomly chosen bit b , and is given access to the same oracles as during the first phase of the attack. The second phase ends when the adversary outputs a bit b' . The adversary is said to win the IND-WKID-CPA game if $b' = b$ and if it never queried the key derivation oracle for the key of any pattern P such that $ID^* \in_* P$. If Succ is the event that the adversary wins the above game, then its advantage is defined as $\epsilon = 2 \cdot \Pr[\text{Succ}] - 1$.

Definition 3. A WKD-IBE scheme is (t, q_K, ϵ) IND-WKID-CPA-secure if all t -time adversaries making at most q_K queries to the key derivation oracle have at most advantage ϵ in the IND-WKID-CPA game described above.

SELECTIVE-IDENTITY SECURITY. As for the case of HIBEs, we also define the weaker selective-identity (sWKID) security notion IND-sWKID-CPA. The IND-sWKID-CPA definition is analogous to the IND-WKID-CPA one given above except that the adversary has to commit to the challenge identity at the beginning of the game, before the master public key is made available.

3.1 Constructions with Linear-Size Ciphertexts

A CONSTRUCTION FROM GENTRY-SILVERBERG'S HIBE SCHEME. In the following, we present a wicked IBE scheme based on the Gentry-Silverberg HIBE scheme [12]. The scheme uses L independent random oracles $H_i : \{0, 1\}^* \rightarrow \mathbb{G}$ for $1 \leq i \leq L$. These can be derived from a single random oracle via standard techniques [4].

We provide some intuition into our construction by taking a closer look at the key derivation of (a slight variant of) the original Gentry-Silverberg HIBE scheme. For master secret key $\alpha \xleftarrow{\$} \mathbb{Z}_p$ and master public key $g_1 \leftarrow g^\alpha$, the decryption key of an identity (ID_0) at the top level is given by $sk_{(ID_0)} \leftarrow H_0(ID_0)^\alpha$. The key for a lower-level identity $(ID_0, \dots, ID_\ell)$ is given by

$$sk_{(ID_0, \dots, ID_\ell)} \leftarrow (H_0(ID_0)^\alpha \cdot \prod_{i=1}^{\ell} H_i(ID_i)^{r_i}, g^{r_1}, \dots, g^{r_\ell})$$

for random $r_1, \dots, r_\ell \xleftarrow{\$} \mathbb{Z}_p$. One could “insert a wildcard” at level $1 \leq j \leq \ell$ by omitting the factor $H_j(ID_j)^{r_j}$ from the product in the first component and omitting the entry g^{r_j} in the vector; any value for ID_j can then be filled in later by choosing r_j , multiplying $H_j(ID_j)^{r_j}$ into the first component and inserting a new component g^{r_j} . Inserting a wildcard at the top level is not so easy though, as knowledge of the master key α is required to compute the factor $H_0(ID_0)^\alpha$. We therefore “disable” the top level by fixing it to identity $\perp$, or equivalently, by including $h_0 = H_0(\perp)$ in the public key. A similar fix can be used to prevent a user at level $\ell < L$ to further derive keys for users at levels $\ell + 1, \dots, L$. Namely, the key is computed as if it were for the identity at level L with the components at levels $\ell + 1, \dots, L$ fixed to $\perp$. Equivalently, one can include the elements $h_i = H_i(\perp)$ for $1 \leq i \leq L$ in the public key.

Before presenting the scheme, we first need to introduce some additional notation. If $P = (P_1, \dots, P_\ell)$ is a pattern, then let $|P| = \ell$ be the length of P , let $W(P)$ be the set containing all wildcard indices in P , i.e. the indices $1 \leq i \leq \ell$ such that $P_i = *$, and let $\overline{W}(P)$ be the complementary set containing all non-wildcard indices. Clearly, $W(P) \cap \overline{W}(P) = \emptyset$ and $W(P) \cup \overline{W}(P) = \{1, \dots, \ell\}$. We also extend the notations $P|_{\leq i}$, $P|_{> i}$ and $P|_I$ that we introduced for identity vectors to patterns in the natural way. We are now ready to present the $\mathcal{GS}\text{-}\mathcal{WKD}\text{-}\mathcal{IBE}$ scheme in full details:

Setup. The root identity chooses random generators $g, h_0, \dots, h_L \xleftarrow{\$} \mathbb{G}^*$. It chooses $\alpha \xleftarrow{\$} \mathbb{Z}_p$ and computes $g_1 \leftarrow g^\alpha$. It publishes $mpk \leftarrow (g, g_1, h_0, \dots, h_L)$ as the master public key and keeps $msk \leftarrow h_0^\alpha$ secret.

Key Derivation. To compute a secret key for a pattern $P = (P_1, \dots, P_\ell)$ directly from the master secret key, the root proceeds as follows. Let $I = \overline{W}(P) \cup \{\ell + 1, \dots, L\}$. For all $i \in I$ the root chooses $r_i \xleftarrow{\$} \mathbb{Z}_p$ and lets $b_i \leftarrow g^{r_i}$. It then computes $a \leftarrow msk \cdot \prod_{i \in \overline{W}(P)} H_i(P_i)^{r_i} \cdot \prod_{i=\ell+1, \dots, L} h_i^{r_i}$. The secret key for pattern P is $sk_P \leftarrow (a, (b_i)_{i \in I})$.

Anyone knowing this secret key can generate a key for a pattern $P' = (P'_1, \dots, P'_{\ell'}) \in_* P$ as follows. Let $I' = \overline{W}(P') \cup \{\ell' + 1, \dots, L\}$. Note that $P' \in_* P$ implies that $I \subseteq I'$. For all $i \in I$, choose $r_i \xleftarrow{\$} \mathbb{Z}_p$ and compute $b'_i \leftarrow b_i \cdot g^{r_i}$; for all $i \in I' \setminus I$, choose $r_i \xleftarrow{\$} \mathbb{Z}_p$ and compute $b'_i \leftarrow g^{r_i}$. Finally, compute $a' \leftarrow a \cdot \prod_{i \in \overline{W}(P')} H_i(P'_i)^{r_i} \cdot \prod_{i=\ell'+1}^L h_i^{r_i}$ and return the secret key $sk_{P'} \leftarrow (a', (b'_i)_{i \in I'})$.

Encryption. To encrypt a message $m \in \mathbb{G}_T$ to identity $ID = (ID_1, \dots, ID_\ell)$ under $mpk = (g, g_1, h_0, \dots, h_L)$, the sender chooses $t \xleftarrow{\$} \mathbb{Z}_p$, computes

$$C_0 \leftarrow g^t$$

$$C_i \leftarrow H_i(ID_i)^t \quad \text{for } i = 1, \dots, \ell$$

$$C_i \leftarrow h_i^t \quad \text{for } i = \ell + 1, \dots, L$$

$$C_{L+1} \leftarrow \hat{e}(g_1, h_0)^t \cdot m$$

and outputs the ciphertext $C = (C_0, \dots, C_{L+1})$.

Decryption. A recipient knowing the secret key sk_P for a pattern $P = (P_1, \dots, P_\ell)$ can decrypt a ciphertext $(C_0, \dots, C_{L+1})$ intended to any identity $ID \in_\star P$ as follows. Let $I = \overline{W}(P) \cup \{\ell + 1, \dots, L\}$ and let $a_P = (a, (b_i)_{i \in I})$. The recipient recovers the plaintext as

$$m \leftarrow C_{L+1} \cdot \frac{\prod_{i \in I} \hat{e}(b_i, C_i)}{\hat{e}(C_0, a)}.$$

Note that the recipient need not even know the exact identity under which the message was encrypted.

The fact that decryption works can be seen as follows. Let $P = (P_1, \dots, P_\ell)$ be a pattern, let $I = \overline{W}(P) \cup \{\ell + 1, \dots, L\}$ and let $sk_P = (a, (b_i)_{i \in I})$ be a secret key for P . For all $i \in I$, let r_i be the discrete logarithm of b_i with respect to g , i.e. $b_i = g^{r_i}$. From the key derivation algorithm one can see that $a = h_0^\alpha \cdot \prod_{i \in \overline{W}(P)} H_i(ID_i)^{r_i} \cdot \prod_{i=\ell+1}^L h_i^{r_i}$. When $(C_0, \dots, C_{L+1})$ is a ciphertext intended for $ID = (ID_1, \dots, ID_{\ell'}) \in_\star P$, we have that

$$\begin{aligned} \hat{e}(C_0, a) &= \hat{e}\left(g^t, h_0^\alpha \cdot \prod_{i \in \overline{W}(P)} H_i(P_i)^{r_i} \cdot \prod_{i=\ell+1}^L h_i^{r_i}\right) \\ &= \hat{e}(g^t, h_0^\alpha) \cdot \prod_{i \in \overline{W}(P)} \hat{e}(g^{r_i}, H_i(P_i)^t) \cdot \prod_{i=\ell+1}^L \hat{e}(g^{r_i}, h_i^t) \\ &= \hat{e}(g_1, h_0)^t \cdot \prod_{i \in I} \hat{e}(b_i, C_i), \end{aligned}$$

where the last equality holds because $P_i = ID_i$ for all $i \in \overline{W}(P)$ if $ID \in_\star P$. Hence, the value of K at decryption is exactly the argument of H_2 at encryption, and the correct message is recovered.

The following theorem states the security of the above scheme in the selective-identity notion under the BDDH assumption in the random oracle model; the proof is given in the full version [2]. Security in the full-identity notion can be obtained at the cost of losing a factor $O(q_H^L)$ in the reduction.

Theorem 1. *Under the (t', ϵ') BDDH assumption, the $\mathcal{GS}\text{-}\mathcal{WKD}\text{-}\mathcal{IBE}$ scheme described above is (t, q_K, q_H, ϵ) IND-sWKID-CPA-secure in the random oracle model for $\epsilon \geq 2\epsilon'$ and $t \leq t' - (q_K + (q_K + 3)L)t_{\text{exp}}$, where t_{exp} is the time required to perform an exponentiation in $\mathbb{G}$.*

CONSTRUCTIONS FROM BONEH-BOYEN'S AND WATERS' HIBE SCHEMES. The attentive reader will have noticed the resemblance of the above scheme with the HIBE schemes of Boneh-Boyen [5] and Waters [21]. Indeed, if identity strings are elements of $\mathbb{Z}_p^*$, then one can obtain a wicked IBE variant of [5] by setting $H_i(ID_i) = h_{i,0} h_{i,1}^{ID_i}$, where $h_{i,0}, h_{i,1}$ are random elements of $\mathbb{G}$ that are fixed in the master public key. This scheme can be proved IND-sWKID-CPA secure under the BDDH assumption in the standard (i.e., non-random oracle) model

using a proof quite similar to the above analysis. Likewise, one can obtain a variant based on Waters' HIBE scheme when identities are n -bit strings by setting $H_i(ID_i = ID_{i,1} \dots ID_{i,n}) = h_{i,0} \prod_{ID_{i,j}=1} h_{i,j}$. An analysis similar to the one in [21] can be used to prove this scheme IND-WKID-CPA secure under the BDDH assumption in the standard model at the cost of losing a factor $O((nq_K)^L)$ in the reduction.

3.2 Constructions with Constant-Size Ciphertexts

In this section, we describe efficient wicked IBE schemes with constant-size ciphertexts based on the Boneh-Boyen-Goh [6] and Waters [21] HIBE schemes. We build the wicked IBE scheme $\mathcal{BBG}\text{-}\mathcal{WKID}\text{-}\mathcal{IBE} = (\text{Setup}, \text{KeyDer}, \text{Enc}, \text{Dec})$ described as follows:

Setup. The trusted authority chooses random generators g from $\mathbb{G}$, a random $\alpha \in \mathbb{Z}_p$ and sets $g_1 \leftarrow g^\alpha$. Next, it picks random elements $g_2, g_3, h_1, \dots, h_L$ from $\mathbb{G}$ and sets $g_4 \leftarrow g_2^\alpha$. The master public key is $mpk = (g, g_1, g_2, g_3, h_1, \dots, h_L)$. The corresponding master secret key is $msk = g_4$.

Key Derivation. Let $P' = (P'_1, \dots, P'_\ell) \in (\mathbb{Z}_p^* \cup \{*\})^{\leq L}$ be the pattern for which a secret key needs to be generated. To compute the secret key for P' from the master secret key, first a random $r \xleftarrow{\$} \mathbb{Z}_p$ is chosen, then the secret key $sk_{P'} = (a'_1, a'_2, b')$ for P' is constructed as

$$a'_1 = g_4 \cdot \left(g_3 \prod_{i \in \overline{W}(P')} h_i^{P'_i} \right)^r ; \quad a'_2 = g^r ; \quad b' = (b_i = h_i^r)_{i \in W(P')} .$$

In order to generate the secret key $sk_{P'}$ for pattern P' from the secret key $sk_P = (a_1, a_2, b)$ for pattern P such that $P' \in_* P$, one simply chooses a random $r' \xleftarrow{\$} \mathbb{Z}_p$ and outputs $sk_{P'} = (a'_1, a'_2, b')$, where

$$\begin{aligned} a'_1 &= a_1 \cdot \left(g_3 \prod_{i \in \overline{W}(P')} h_i^{P'_i} \right)^{r'} \cdot \left(\prod_{i \in \overline{W}(P') \cap W(P)} b_i^{P'_i} \right) \\ a'_2 &= a_2 \cdot g^{r'} \\ b' &= (b'_i = b_i \cdot h_i^{r'})_{i \in W(P')} \end{aligned}$$

Encryption. To encrypt a message $m \in \mathbb{G}_T$ for an identity $ID = (ID_1, \dots, ID_\ell)$, the sender first chooses $t \xleftarrow{\$} \mathbb{Z}_p$ and outputs the ciphertext $C = (C_1, C_2, C_3) \in \mathbb{G} \times \mathbb{G} \times \mathbb{G}_T$, where

$$C_1 = g^t ; \quad C_2 = \left(g_3 \prod_{i=1}^{\ell} h_i^{ID_i} \right)^t ; \quad C_3 = m \cdot \hat{e}(g_1, g_2)^t .$$

Decryption. Let be the $C = (C_1, C_2, C_3)$ and $ID = (ID_1, \dots, ID_\ell)$ be the identity to which the ciphertext was created. If the receiver is the root authority

holding the master key msk , then he can recover the message by computing $C_3/\hat{e}(C_1, msk)$. Any other receiver holding a secret key for pattern P such that $ID \in_* P$ can decrypt the ciphertext as follows. Let $sk_P = (a_1, a_2, b)$ be the decryption key for the receiver. He can recover the message by computing

$$a'_1 \leftarrow a_1 \cdot \left(\prod_{i \in \overline{W}(P)|_{\leq \ell}} b_i^{ID_i} \right) \quad \text{and} \quad m \leftarrow C_3 \cdot \frac{\hat{e}(a_2, C_2)}{\hat{e}(C_1, a'_1)}.$$

The fact that decryption works can be seen as follows. Since $ID \in_* P$, we have that $P_i = ID_i$ for all $i \in \overline{W}(P)|_{\leq \ell}$. Thus the quantity $\frac{\hat{e}(a_2, C_2)}{\hat{e}(C_1, a'_1)}$ becomes:

$$\begin{aligned} \frac{\hat{e}(a_2, C_2)}{\hat{e}(C_1, a_1 \prod_{i \in \overline{W}(P)|_{\leq \ell}} b_i^{ID_i})} &= \frac{\hat{e}(g^r, (g_3 \prod_{i=1}^{\ell} h_i^{ID_i})^r)}{\hat{e}(g^t, g_4 \cdot (g_3 \prod_{i \in \overline{W}(P)} h_i^{P_i})^r \cdot \prod_{i \in \overline{W}(P)|_{\leq \ell}} b_i^{ID_i})} \\ &= \frac{\hat{e}(g^r, (g_3 \prod_{i=1}^{\ell} h_i^{ID_i})^t)}{\hat{e}(g^t, g_4) \hat{e}(g^t, (g_3 \prod_{i=1}^{\ell} h_i^{ID_i})^r)} = \frac{1}{\hat{e}(g^t, g_4)} = \frac{1}{\hat{e}(g_1, g_2)^t} \end{aligned}$$

The following theorem states the security of the above scheme in the selective-identity notion under the ℓ -BDHE assumption in the standard model. The proof is given in the full version [2]. We remark that, interestingly, we can only prove security of the scheme based on the ℓ -BDHE assumption, whereas the weaker ℓ -BDHI assumption was sufficient for the security proof of the HIBE scheme [6].

Theorem 2. *Let $\mathcal{BBG}\text{-}\mathcal{WKD}\text{-IBE}$ be the WKD-IBE scheme as described above. Under the decisional (t, ϵ, ℓ) -BDHE assumption, the $\mathcal{BBG}\text{-}\mathcal{WKD}\text{-IBE}$ scheme of depth $L = \ell - 1$ is $(t', q_K, 2\epsilon)$ IND-sWKID-CPA-secure where $t' = t - O(Lq'_K) \cdot t_{\text{exp}}$ and t_{exp} is the time it takes to perform an exponentiation in $\mathbb{G}$.*

FULL SECURITY IN THE STANDARD MODEL. It is mentioned in [6] that using techniques from Waters [21] one can construct a variant of their HIBE scheme that achieves full security in the standard model. The same techniques can be also used to achieve full IND-WKID-CPA security in the standard model for the $\mathcal{BBG}\text{-}\mathcal{WKD}\text{-IBE}$ scheme, at the cost of increasing the master public key size to $(n + 1)L + 3$ group elements, where n is the length of an identity string.

3.3 Full Security in the Random Oracle Model

As in the case of IBE and HIBE schemes [5,6], any WKD-IBE scheme $\mathcal{WKD}\text{-IBE}$ that is IND-sWKID-CPA-secure can be transformed into a WKD-IBE scheme $\mathcal{WKD}\text{-IBE}'$ that is IND-WKID-CPA-secure in the random oracle model, by replacing every pattern (or identity) at key derivation or encryption with the hash of that pattern, if that pattern is not a wildcard. That is, any given pattern $P = (P_1, \dots, P_\ell)$ in $\mathcal{WKD}\text{-IBE}$ is mapped onto a pattern $P' = (P'_1, \dots, P'_\ell)$ in $\mathcal{WKD}\text{-IBE}'$, where $P'_i = H_i(P_i)$ if $P_i \neq *$ or $P'_i = *$ otherwise, and H_i , $1 \leq i \leq L$ are independent random oracles mapping arbitrary bit strings into an appropriate range $\mathcal{ID}$ corresponding to the identity space of $\mathcal{WKD}\text{-IBE}$. As in the cases of HIBE schemes, this transformation only works if the depth L is

logarithmic in the security parameter due to the loss of a factor $O(q_H^L)$ in the reduction. Moreover, ID needs to be sufficiently large to make the probability of collisions in the output of the hash function negligible.

4 Application to Identity-Based Broadcast Encryption

4.1 Definitions

An identity-based broadcast encryption (IBBE) scheme is a tuple of algorithms $IBBE = (\text{Setup}, \text{KeyDer}, \text{Enc}, \text{Dec})$ providing the following functionality. The trusted authority runs **Setup** to generate a master key pair (mpk, msk) . It publishes the master public key mpk and keeps the master secret key msk private. When a user with identity ID wishes to become part of the system, the trusted authority generates a user decryption key $sk_{ID} \xleftarrow{\$} \text{KeyDer}(msk, ID)$, and sends this key over a secure and authenticated channel to the user. To broadcast an encrypted message m to a set of users with identities $S = \{ID_1, \dots, ID_k\}$ of cardinality $k \leq L$, the sender computes the ciphertext $C \xleftarrow{\$} \text{Enc}(mpk, S, m)$, which can be decrypted by a user holding sk_{ID} for any $ID \in S$ as $m \leftarrow \text{Dec}(sk_{ID}, C, S)$. Here the value L is an upper bound on the maximal number of distinct receivers for a broadcast encryption.

The security of an IBBE scheme is defined through the following game. In a first phase, the adversary is given as input the master public key mpk of a freshly generated key pair $(mpk, msk) \xleftarrow{\$} \text{Setup}$. In a chosen-plaintext attack (IND-ID-CPA), the adversary is given access to a key derivation oracle that on input of an identity ID , returns the secret key $sk_{ID} \xleftarrow{\$} \text{KeyDer}(msk, ID)$ corresponding to identity ID . At the end of the first phase, the adversary outputs two equal-length challenge messages $m_0^*, m_1^* \in \{0, 1\}^*$ and a challenge set of identities $S^* = (ID_1^*, \dots, ID_{k^*}^*)$, where $0 \leq k^* \leq L$. The game chooses a random bit $b \xleftarrow{\$} \{0, 1\}^*$, generates a challenge ciphertext $C^* \xleftarrow{\$} \text{Enc}(mpk, S^*, m_b^*)$ and gives C^* as input to the adversary for the second phase, during which it gets access to the same oracles as during the first phase. Assume that during the attack the adversary made key derivation queries for identities $ID_1, \dots, ID_{q_K}$. The adversary wins the game if it outputs a bit $b' = b$ and $S^* \cap \{ID_1, \dots, ID_{q_K}\} = \emptyset$.

Definition 4. An IBBE scheme is (t, q_K, ϵ) -IND-ID-CPA-secure if all t -time adversaries making at most q_K queries to the key derivation oracle have at most advantage ϵ in winning the IND-ID-CPA game described above.

SELECTIVE-IDENTITY SECURITY. As for the previous primitives, we further define the weaker (sID) security notion IND-sID-CPA. The IND-sID-CPA definition is analogous to the IND-ID-CPA one except that the adversary has to commit to the challenge set of identities $S^* = (ID_1^*, \dots, ID_{k^*}^*)$ at the beginning of the game, before even seeing the public-key.

4.2 A Construction from Any Wicked IBE Scheme

First, observe that an IBBE scheme can be trivially constructed from any IBE scheme by concatenating ciphertext. Meaning, the IBBE encryption for the iden-

tity set $ID = \{ID_1, \dots, ID_k\}$ is simply the concatenation of k separate ciphertexts, one for each identity ID_i in the set ID . This leads to IBBE ciphertext sizes that are a factor of k longer than the original IBE ciphertexts.

We now present a generic construction from any WKD-IBE scheme that, depending on the instantiation, can offer advantages over the trivial one. To any WKD-IBE scheme $\mathcal{WKID-IBE} = (\text{Setup}, \text{KeyDer}, \text{Enc}, \text{Dec})$, we associate an IBBE scheme $\mathcal{IBBE} = (\text{Setup}, \text{KeyDer}', \text{Enc}', \text{Dec}')$. For an identity $ID \in \{0, 1\}^*$, define

$$P_i(ID) = (*, \dots, *, \underbrace{ID}_{i\text{th position}}, *, \dots, *)$$

as a pattern of length L that has ID at its i th position and the rest consists of wildcards.

Setup. Setup outputs whatever the WKD-IBE setup outputs.

Key Derivation. Let ID be the identity for which the user secret key sk_{ID} needs to be generated. The user secret key is defined as the set of L distinct WKD-IBE user secret keys

$$sk_{ID} = \{sk_{P_1(ID)}, \dots, sk_{P_L(ID)}\},$$

where $sk_{P_i(ID)}$ can be computed by calling $\text{KeyDer}(msk, P_i(ID))$.

Encryption. Let m be the message and let $S = \{ID_1, \dots, ID_k\}$ be the set of broadcast recipients of cardinality $k \leq L$ that we assume to be ordered with respect to some unique standard ordering. The IBBE ciphertext is defined as the WKD-IBE encryption of message m and identity vector $ID = (ID_1, \dots, ID_k)$.

Decryption. Let $sk_{ID} = \{sk_{P_1(ID)}, \dots, sk_{P_L(ID)}\}$ be the user secret key of identity ID . Let $S = \{ID_1, \dots, ID_k\}$ be the set of $k \leq L$ recipients to whom the ciphertext C was encrypted, and let index $1 \leq j \leq k$ be such that $ID = ID_j \in S$. It is clear that $(ID_1, \dots, ID, \dots, ID_k) \in_* P_j(ID)$, and therefore that the ciphertext can be decrypted as $m \leftarrow \text{Dec}(sk_{P_j(ID)}, C, ID)$.

Theorem 3. *If $\mathcal{WKID-IBE}$ is a (t, q_K, ϵ) IND-sWKID-CPA-secure (resp. IND-WKID-CPA-secure) WKD-IBE scheme, then the IBBE scheme $\mathcal{IBBE}$ described above is (t, q_K, ϵ) -IND-sID-CPA-secure (resp. IND-ID-CPA-secure).*

The crucial observation is the following. Let $S^* = \{ID_1^*, \dots, ID_k^*\}$ be the set of challenge broadcast receivers and let $ID_1, \dots, ID_{q_K}$ be the identities an adversary attacking the IBBE scheme queries the user secret key for. The imposed requirement is that $S^* \cap \{ID_1, \dots, ID_{q_K}\} = \emptyset$. For $1 \leq i \leq q_K$ and $1 \leq j \leq L$ consider the user secret keys for the patterns $P_j(ID_i) = (*, \dots, *, ID_i, *, \dots, *)$ (i.e., ID_i is at the j th position) that are established by the transformation when simulating the IBBE key derivation oracle. For a successful simulation we have to show that $ID^* = (ID_1^*, \dots, ID_k^*) \notin_* P_j(ID_i)$. But this is the case since by $S^* \cap \{ID_1, \dots, ID_{q_K}\} = \emptyset$ and we can guarantee that $ID_i \neq ID_l^*$ for all $1 \leq i \leq q_K$ and $1 \leq l \leq k$.

The above construction allows for the following trade off between ciphertext size and key size. If $L = L'A$, then one can obtain an IBBE scheme with ciphertext size of A times that of the WKD-IBE scheme, while having a key length that is only L' times that of the WKD-IBE scheme. The new scheme creates master public keys to allow for broadcast encryption to sets of maximal cardinality L' . To encrypt a message to a set of broadcast identities $S = \{ID_1, \dots, ID_k\}$ of cardinality $k \leq L$ split the set S into A smaller sets $S_1, \dots, S_A$, each of cardinality $L/A \leq L'$ and define the new broadcast ciphertext to be $(C_1, \dots, C_A)$, where C_i is the encryption of the message m to the set S_i .

4.3 Instantiations

Among all the instantiations of IBBE schemes based on WKD-IBE schemes, the most attractive one is that obtained from the WKD-IBE scheme based on [6] because it achieves constant-size ciphertexts. However, it has the disadvantage of having private keys of size $O(L^2)$. Instantiations with any of the other WKD-IBE schemes that we proposed are less attractive because they have ciphertext size $O(L)$, just like the trivial ciphertext-concatenation scheme. Unlike most other (public-key) broadcast schemes however, these instantiations do have the remarkable advantage that knowledge of the set of recipients is not required in order to decrypt the message.

Chatterjee and Sarkar [9] recently proposed a direct IBBE scheme that is closely related to our generic construction when instantiated with the WKD-IBE scheme based on [6]. Their scheme does not impose an a priori maximum on the number of recipients ℓ , but makes clever use of a non-cryptographic hash function to achieve an average ciphertext size $O(\ell/L)$ and private key size $O(L)$, where the “average” is taken over the recipients’ identities. This means that when $\ell \leq L$, their scheme has constant ciphertext size on average. Worst-case however, their scheme has ciphertext size $O(\ell)$, which is worse than our construction.

5 Wicked and Wildcard Signatures

As observed by Naor [7], any IBE scheme automatically gives rise to a signature scheme by using as a signature on message m the decryption key for identity $ID = m$. Verification can be done by encrypting a random message to identity $ID = m$ and testing whether it decrypts correctly, but most concrete schemes have a more natural and efficient verification test. Likewise, one can construct an L -level hierarchical identity-based signature (HIBS) scheme from an $(L + 1)$ -level HIBE [12] by letting the signature on message m by identity $(ID_1, \dots, ID_\ell)$ be given by the decryption key for identity $(0 \| ID_1, \dots, 0 \| ID_\ell, 1 \| m)$. The same technique can be used to construct *wicked identity-based signatures* (WKD-IBS), the signing analogue to wicked IBE. Here, a root authority derives secret signing keys for identity patterns with wildcards, from which anyone can further derive signing keys for matching patterns. An L -level WKD-IBS is constructed from an $(L + 1)$ -level WKD-IBE by letting the signature on mes-

sage m by identity $(ID_1, \dots, ID_\ell)$ be given by the decryption key for identity $(0\|ID_1, \dots, 0\|ID_\ell, 1\|m)$.

Alternatively, and perhaps more interestingly, one could also use the wildcard functionality as a homomorphism on the message being signed, rather than for the signers' identities. This yields a new primitive that we call *wildcard signatures*, that allow to sign message patterns instead of simple messages, possibly containing wildcards at certain positions. Given such a signature, anyone can compute a valid signature for any message created by replacing wildcards with concrete values. This could be used for example to implement signed fill-out forms, where each input field is represented by a wildcard in the message.

The construction from a WKD-IBE scheme is straightforward: the key pair is given by the master key pair of the WKD-IBE scheme. The signature on a message pattern P is given by the decryption key for P . Deriving a valid signature for a message pattern $P' \in_* P$ can be done by deriving a decryption key for P' . Verification is done by filling up the remaining wildcards with random messages to create a vector of messages M , encrypting a random message under identity M , and checking whether decryption using the signature as secret key returns the correct message. In fact, one can easily see that the schemes discussed here allow for more efficient deterministic verification algorithms.

Wildcard signatures can be seen as a special instance of homomorphic signatures [17,15,3,16]. Their relation to wicked IBE is particularly reminiscent of the relation between HIBS schemes and append-only signatures [16]. They can also be seen as the dual of redactable signatures [15] that allow anyone to erase parts of a signed message without invalidating the signature.

A fairly simple, generic construction from standard signatures also exists. Namely, for each wildcard in the message the signer generates a fresh key pair, and then signs the message together with all generated public keys. The overall signature also contains the public and secret keys corresponding to all wildcards. To replace a wildcard at position i with a concrete value, the i -th secret key is replaced with a signature on the new value under the i -th public key. The disadvantage of this generic construction is that signature length and verification time are both linear in the number of *original* wildcards in the message, even after these wildcards have been replaced with original values. The signature length and verification time of the scheme derived from the *BBG-WKD-IBE* scheme on the other hand is only linear in the number of wildcards that are still present in the message. Also, signatures generated by the generic construction are linkable in the sense that one can check whether a given signature was derived from a second one by filling in wildcards. The decryption keys of the *BBG-WKD-IBE* scheme, and therefore the signatures of the associated wildcard signature scheme, can be re-randomized to prevent this type of linkability.

Finally, one could even imagine *wicked wildcard signatures* that allow for wildcards in *both* the signers' identities and the messages being signed. Such schemes are easily constructed from a WKD-IBE scheme by using a different encoding for identity strings and messages, as was done in the construction of WKD-IBS schemes above.

Acknowledgements

The authors have been supported in part by the European Commission through the IST Program under Contract IST-2002-507932 ECRYPT. The second author was supported in part by research program Sentinels (<http://www.sentinels.nl>). Sentinels is being financed by Technology Foundation STW, the Netherlands Organization for Scientific Research (NWO), and the Dutch Ministry of Economic Affairs. The third author is a Postdoctoral Fellow of the Research Foundation – Flanders (FWO-Vlaanderen), and was supported in part by the European Commission through IST Program under Contract IST-2006-034238 SPEED, and in part by the IAP Program P6/26 BCRYPT of the Belgian State (Belgian Science Policy).

References

1. Abdalla, M., Catalano, D., Dent, A., Malone-Lee, J., Neven, G., Smart, N.: Identity-based encryption gone wild. In: Bugliesi, M., Preneel, B., Sassone, V., Wegener, I. (eds.) ICALP 2006. LNCS, vol. 4052, pp. 300–311. Springer, Heidelberg (2006)
2. Abdalla, M., Kiltz, E., Neven, G.: Generalized Key Delegation for Hierarchical Identity-Based Encryption. Cryptology ePrint Archive, Report 2007/221 (2007), <http://eprint.iacr.org/>
3. Bellare, M., Neven, G.: Transitive signatures: new schemes and proofs. IEEE Transactions on Information Theory 51(6), 2133–2151 (2005)
4. Bellare, M., Rogaway, P.: Random oracles are practical: A paradigm for designing efficient protocols. In: ACM CCS '93, pp. 62–73. ACM Press, New York (1993)
5. Boneh, D., Boyen, X.: Efficient selective-ID secure identity based encryption without random oracles. In: Cachin, C., Camenisch, J.L. (eds.) EUROCRYPT 2004. LNCS, vol. 3027, pp. 223–238. Springer, Heidelberg (2004)
6. Boneh, D., Boyen, X., Goh, E.-J.: Hierarchical identity based encryption with constant size ciphertext. In: Cramer, R.J.F. (ed.) EUROCRYPT 2005. LNCS, vol. 3494, pp. 440–456. Springer, Heidelberg (2005)
7. Boneh, D., Franklin, M.K.: Identity based encryption from the Weil pairing. SIAM Journal on Computing 32(3), 586–615 (2003)
8. Boneh, D., Gentry, C., Waters, B.: Collusion resistant broadcast encryption with short ciphertexts and private keys. In: Shoup, V. (ed.) CRYPTO 2005. LNCS, vol. 3621, pp. 258–275. Springer, Heidelberg (2005)
9. Chatterjee, S., Sarkar, P.: Multi-receiver identity-based key encapsulation with shortened ciphertext. In: Barua, R., Lange, T. (eds.) INDOCRYPT 2006. LNCS, vol. 4329, pp. 394–408. Springer, Heidelberg (2006)
10. Dodis, Y., Fazio, N.: Public key broadcast encryption for stateless receivers. In: Feigenbaum, J. (ed.) DRM 2002. LNCS, vol. 2696, pp. 61–80. Springer, Heidelberg (2003)
11. Fiat, A., Naor, M.: Broadcast encryption. In: Stinson, D.R. (ed.) CRYPTO 1993. LNCS, vol. 773, pp. 480–491. Springer, Heidelberg (1994)
12. Gentry, C., Silverberg, A.: Hierarchical ID-based cryptography. In: Zheng, Y. (ed.) ASIACRYPT 2002. LNCS, vol. 2501, pp. 548–566. Springer, Heidelberg (2002)

13. Goyal, V., Pandey, O., Sahai, A., Waters, B.: Attribute-based encryption for fine-grained access control of encrypted data. In: ACM CCS 06, pp. 89–98. ACM Press, New York (2006) (available as Cryptology ePrint Archive Report 2006/309)
14. Horwitz, J., Lynn, B.: Toward hierarchical identity-based encryption. In: Knudsen, L.R. (ed.) EUROCRYPT 2002. LNCS, vol. 2332, pp. 466–481. Springer, Heidelberg (2002)
15. Johnson, R., Molnar, D., Song, D.X., Wagner, D.: Homomorphic signature schemes. In: Preneel, B. (ed.) CT-RSA 2002. LNCS, vol. 2271, pp. 244–262. Springer, Heidelberg (2002)
16. Kiltz, E., Mityagin, A., Panjwani, S., Raghavan, B.: Append-only signatures. In: Caires, L., Italiano, G.F., Monteiro, L., Palamidessi, C., Yung, M. (eds.) ICALP 2005. LNCS, vol. 3580, pp. 434–445. Springer, Heidelberg (2005)
17. Rivest, R.: Two signature schemes. Slides from talk given at Cambridge University (October 2000)
18. Sakai, R., Ohgishi, K., Kasahara, M.: Cryptosystems based on pairing. In: SCIS 2000, Okinawa, Japan (January 2000)
19. Shacham, H.: The BBG HIBE has limited delegation. Cryptology ePrint Archive, Report 2007/201 (2007), <http://eprint.iacr.org/>
20. Shamir, A.: Identity-based cryptosystems and signature schemes. In: Blakely, G.R., Chaum, D. (eds.) CRYPTO 1984. LNCS, vol. 196, pp. 47–53. Springer, Heidelberg (1985)
21. Waters, B.R.: Efficient identity-based encryption without random oracles. In: Cramer, R.J.F. (ed.) EUROCRYPT 2005. LNCS, vol. 3494, pp. 114–127. Springer, Heidelberg (2005)