Protecting Receivers' Identities in Secure Data Distribution

Nelly Fazio*† and Irippuge Milinda Perera†
* The City College of CUNY
fazio@cs.ccny.cuny.edu
† The Graduate Center of CUNY
{nfazio,iperera}@gc.cuny.edu

July 16, 2012

***Abstract*—In the standard setting of broadcast encryption, information about the receivers is transmitted as part of the ciphertext. In several broadcast scenarios, however, the identities of the users authorized to access the content are often as sensitive as the content itself. In this paper, we propose the first broadcast encryption scheme to attain meaningful guarantees of receiver anonymity with ciphertexts that are sublinear in the number of authorized users. We formalize the notion of *outsider-anonymous broadcast encryption*(oABE), and describe generic constructions in the standard model that achieve outsider-anonymity under adaptive corruptions in the chosen-plaintext and chosen-ciphertext settings. We also describe two constructions with enhanced decryption, one under the gap Diffie-Hellman assumption, in the random oracle model, and the other under the decisional Diffie-Hellman assumption, in the standard model.**

***Index Terms*—Anonymous IBE, broadcast encryption, recipient privacy, subset cover framework.**

I. Introduction

Conventional encryption provides the means for secret transmission of data in point-to-point communication. The setting of broadcast encryption [1], [2], instead, consists of a *sender*, an insecure unidirectional *broadcast channel*, and a universe of *receivers*. When the sender wants to transmit some digital content, it specifies the set of authorized receivers and creates an encrypted version of the content. A secure broadcast encryption scheme enables legitimate receivers to recover the original content, while ensuring that excluded users just obtain meaningless data, even in the face of collusions.

The intrinsic access control capabilities of broadcast encryption schemes make them a useful tool for many natural applications, spanning from protecting copyrighted content distributed as stored media [3], to managing digital subscriptions to satellite TV, to controlling access in encrypted file systems [4]. Thanks to its versatility, broadcast encryption has received a lot of attention from the crypto research community in recent years (see *e.g.*, [5]–[14]). The quest, however, has been for ever more efficient solutions in terms of broadcast communication, key storage and encryption/decryption running time. Little attention, instead, has been devoted to the exploration of refined security models that accurately account for the requirements inherent in multi-recipient communication. More specifically, the focus has been on providing assurance for sender-oriented properties, while overlooking the security and privacy concerns of the receivers.

One problem with the above (informal) definition of broadcast encryption is the implicit requirement that, whenever the digital content is encrypted and sent in broadcast, information about the set of authorized receivers is necessary to decrypt it correctly. Therefore, the set of authorized receivers is transmitted as part of the ciphertext. This in particular implies that an eavesdropper, even if unable to recover the message, can still easily discover the identities of the actual receivers of the content. A way to address the privacy implications that result from specifying explicitly the set of authorized receivers in the broadcast is to use ephemeral IDs and to keep secret the table that associates such IDs with the actual receivers. This simple solution, however, would at best result in a pseudonym system, in which it is still possible to link pseudonyms across transmissions and determine whether the same entity is an authorized receiver for two different broadcasts.

Breach of receivers' privacy is not tolerable in several broadcast scenarios. Consider, for example, the application of broadcast encryption schemes to the setting of pay-per-view TV. If the content of the encrypted broadcast is legal but controversial in nature, it is particularly important to protect the privacy of the receivers of the broadcast. A broadcast encryption scheme can also be used to realize efficient encrypted file systems. The identities of the users authorized to access a file, however, are often more sensitive than the content itself. A receiver-anonymous broadcast encryption scheme would prevent disclosure of the identities of the authorized users, not only from outsiders, but also from one another. This work is also relevant to TA5's goal of advancing networking technologies at support of military operations, in that it enables secure distribution of tactical data in missions with ad-hoc team formation, in which the identities of the operatives authorized to access the content need to be kept secret as well.

Anonymous Broadcast Encryption. An interesting variant of the broadcast encryption setting was proposed by

Barth *et al.* in [15]. Therein, the authors introduce the notion of *private* broadcast encryption scheme, explicitly aiming to protect the identities of the receivers. As a proof-of-concept, they also suggest both generic and number-theoretic public-key constructions that do not leak any information about the list of authorized receivers, and are secure in the standard model and in the random oracle model, respectively. The proposed schemes, however, have communication complexity linear in the number of recipients. In [16], Libert *et al.* recently suggested proof techniques to argue the security of (a variant of) the number-theoretic construction of [15] without reliance on random oracles, thus attaining anonymous broadcast encryption with efficient decryption in the standard model. Still, ciphertexts in the resulting construction have length linear in the number of recipients. In [17], Kiayias and Samari put forth lower bounds on the ciphertext size of private broadcast encryption schemes and show, among other results, that fully anonymous broadcast encryption schemes with a certain "atomicity" property (satisfied, *e.g.*, by the schemes of [15], [16]) must have $\Omega(s \cdot k)$ ciphertext size, where s is the number of authorized receivers and k is the security parameter.

Krzywiecki *et al.* presented a private public-key broadcast encryption scheme with communication complexity proportional to the number of revoked users [18]. The security analysis of the proposed solution is rather informal, however, so the security guarantees are at best heuristic.

In [19], Yu *et al.* presented the first *secret-key* multicast scheme with membership anonymity and communication complexity independent of the number of receivers. The proposed scheme not only hides the *identities* of the receives, but also *the number* of users allowed to receive the content. A shortcoming is that only a single user can be revoked for each broadcast.

A promising research line toward practical receiver-anonymous broadcast encryption has recently been started by Jarecki and Liu [20]. The authors propose the first construction of an efficient unlinkable secret handshake scheme, which is an authenticated key exchange protocol providing *affiliation/policy hiding* (*i.e.*, the transmission hides the affiliation and the identities of all parties) and *unlinkability* (*i.e.*, it is impossible to link any two instances of the secret handshake protocol). The proposed construction can be seen as a *stateful* version of a public-key broadcast encryption scheme, with the additional property of protecting the receivers' identities. Statefulness, however, implies that the key used to encrypt the broadcasts changes for each transmission, and receivers need to keep track of the changes to be able to recover the content.

An interesting trait of the of construction of [20] is that it trades some degree of anonymity for better efficiency: while the receiver's identities are hidden from outsiders, the scheme still allows authorized users to learn information about other members of the receiver set.

OUR CONTRIBUTIONS. In this paper we propose the first broadcast encryption scheme with sublinear ciphertexts to achieve meaningful guarantees of receiver anonymity. In particular, we formalize the notion of *outsider-anonymous broadcast encryption*(oABE), and describe a generic construction based on any anonymous identity-based encryption scheme (AIBE). Compared with the work of [20], our construction has the advantage of being *stateless*, and with constant public key size.

Additionally, by adapting the techniques of [15], we also obtain an efficient construction with enhanced decryption, where for a given oABE ciphertext, the decryption algorithm executes a single AIBE decryption operation. As outlined in Table I, by relaxing the anonymity guarantees, our constructions achieve sublinear ciphertexts size and constant public key size.

ORGANIZATION. Section II provides a brief review of the Subset Cover Framework [6] and of Anonymous Identity-Based Encryption [21], [22]. The setting of outsider-anonymous broadcast encryption is introduced in Sect. III. In Sect. IV we first present generic constructions in the standard model that achieve outsider-anonymity under adaptive corruptions in the chosen-plaintext (Sect. IV-A) and chosen-ciphertext (Sect. IV-B) settings. Next, we describe a CCA-secure construction with enhanced decryption (Sect. IV-C) under the gap Diffie-Hellman assumption in the random oracle model, and outline how to extend it to the standard model, using the twin-DH-based techniques of [16]. Finally, we outline an optimization for the symmetric-key setting to accommodate storage-sensitive systems and attain constant key storage at the Center, while maintaining efficient decryption and logarithmic storage at the receivers (Sect. IV-D).

II. BACKGROUND

A. The Subset Cover Framework

The *Subset Cover Framework* proposed by Naor *et al.* [6] is an environment for defining and analyzing the security of revocation schemes in the symmetric key setting, where only the Center can broadcast. The main idea of this framework is to define a collection $\mathcal{S}$ of subsets of the universe of users $\mathcal{U} = \{1, \ldots, N\}$ in the system, and assign each subset $S_j \in \mathcal{S}$ a long-lived key, which is also provided to the users belonging to S_j. When broadcasting a message m, first the Center determines the set of revoked users $\mathcal{R}$, then it finds a set of disjoint subsets $\mathcal{C}$ from the collection $\mathcal{S}$ that "covers" the set $\mathcal{U} \backslash \mathcal{R}$ of receivers, and finally it encrypts the short-lived session key used to encrypt m under all the long-lived keys associated with each subset in $\mathcal{C}$.

In [6], the authors also provide two instantiations of revocation schemes in the Subset Cover framework namely, the Complete Subtree (CS) method and the Subset Difference (SD) method. In the CS method, the key assignment is information-theoretic but the ciphertext is $\mathcal{O}\left(r \log\left(\frac{N}{r}\right)\right)$ long, whereas in the SD method, the ciphertext length is $\mathcal{O}\left(2r - 1\right)$ but the key assignment is computational, where r is the number of revoked users. Although the ciphertext length of the CS method is asymptotically bigger than that of the SD method, we are still interested in the CS method due to its information-theoretic key assignment nature, which seems to be crucial for efficiently preserving the anonymity of the receivers.

TABLE I
COMPARISON OF THE MAIN EFFICIENCY PARAMETERS OF OUR oABE SCHEME WITH [15] AND [16].

	Scheme	PK Length	SK Length	CT Length	Decryption Attempts
Regular	BBW06 [15]	$\mathcal{O}(N)$	$\mathcal{O}(1)$	$\mathcal{O}(N-r)$	$\mathcal{O}(N-r)$
	LPQ12 [16]	$\mathcal{O}(N)$	$\mathcal{O}(1)$	$\mathcal{O}(N-r)$	$\mathcal{O}(N-r)$
	Ours (oABE)	$\mathcal{O}(1)$	$\mathcal{O}(\log N)$	$\mathcal{O}\left(r\log\left(\frac{N}{r}\right)\right)$	$\mathcal{O}\left(r\log\left(\frac{N}{r}\right)\log N\right)$
Enhanced	BBW06 [15]	$\mathcal{O}(N)$	$\mathcal{O}(1)$	$\mathcal{O}(N-r)$	1
	LPQ12 [16]	$\mathcal{O}(N)$	$\mathcal{O}(1)$	$\mathcal{O}(N-r)$	1
	Ours (oABE)	$\mathcal{O}(N)$	$\mathcal{O}(\log N)$	$\mathcal{O}\left(r\log\left(\frac{N}{r}\right)\right)$	1

1) Complete Subtree (CS) Method: In the Complete Subtree (CS) method as introduced in [6], the N users in the system are represented as the leaves of a full binary tree $\mathcal{T}$. Since this requires N to be a power of 2, dummy users are added to the system in case N is not a power of 2. The collection $\mathcal{S}$ contains all possible complete subtrees of $\mathcal{T}$. More precisely, $\mathcal{S}$ contains a subtree for every node $v_j \in \mathcal{T}$. Since there are $2N-1$ nodes in $\mathcal{T}$, $|\mathcal{S}| = 2N-1$.

As for key assignment, every subtree in $\mathcal{S}$ is assigned a long-lived symmetric key which is also made available to the users (leaves) of the given subtree. Since any user u_i, for $1 \leq i \leq N$, is a member of all the subtrees rooted at each node v_j, for $1 \leq j \leq \log N + 1$, in the path from the root of $\mathcal{T}$ down to u_i, the length of the user secret key is $\mathcal{O}(\log N)$.

The ciphertext length in the CS method is $\mathcal{O}\left(r\log\left(\frac{N}{r}\right)\right)$ due to the fact that a logarithmic number of subtrees is required to exclude each of the r revoked users (see [6] for further details).

2) Extension of the CS Method to the Public Key Setting: As mentioned earlier, the original CS method applies in the symmetric key setting. Thus, only the Center can broadcast since only it knows all the long-lived keys associated with each subtree in $\mathcal{S}$. In [8], Dodis and Fazio extended the original CS method to the public key setting by using a two step process.

The first step is a unique assignment of hierarchical identifiers (HID) to the nodes in $\mathcal{T}$ as follows. First, assign the root of $\mathcal{T}$ a special ID, which we refer to as Root. Then, assign each edge of $\mathcal{T}$ with ID 0 or 1 depending on whether the edge connects its parent node to the left or right child. Now, HID_j of any node $v_j \in \mathcal{T}$ can be computed by concatenating all the edge IDs starting from the root of $\mathcal{T}$ down to v_j and then pre-pending the root ID at the front. Since any prefix of HID_j of v_j represents the valid HID of a parent node of v_j, for the simplicity of notation, we denote by $\mathsf{HID}_{i|j}$ the prefix of the hierarchical identifier HID_i of length j.

The second step is to use Identity-Based Encryption (IBE), further explained in Sect. II-B, to encrypt the short-lived session key during broadcast, essentially porting the original CS method to the public key setting. This allows any user to broadcast a message since the tree structure of the users $\mathcal{T}$ and the HIDs of the roots of the subtrees of $\mathcal{T}$ are publicly known. In this setting, the Center acts as the trusted authority to provide each user with the $\log N + 1$ IBE secret keys of the HIDs of the roots of the subtrees that the user belongs to.

B. Anonymous Identity-Based Encryption (AIBE)

Identity-Based Encryption (IBE), originally proposed by Shamir in [23], is a public key encryption scheme in which the user public key is an arbitrary bit-string and the user secret key is generated by a trusted authority known as the *Private Key Generator* (PKG) using its master key. The first implementation of this scheme was given in [24] (further implementations can be found in [25]–[27] to name a few).

An IBE scheme is called anonymous, formally called Anonymous Identity-Based Encryption (AIBE), if an adversary cannot distinguish the public key under which a ciphertext is generated. This notion of anonymity was first introduced in [21]. Subsequent implementations can be found in [28] and [22]. Given below is the formal definition of an AIBE scheme. We refer the reader to [21] for further details including the formal definition of security.

Definition 1: An anonymous identity-based encryption (AIBE) scheme, associated with a message space $\mathcal{MSP}$, and a ciphertext space $\mathcal{CSP}$, is a tuple of probabilistic polynomial algorithms $(\mathsf{Init}, \mathsf{Ext}, \mathsf{Enc}, \mathsf{Dec})$ such that:

$(\mathsf{PK}, \mathsf{MSK}) \leftarrow \mathsf{Init}(1^\lambda)$:
The initialization algorithm Init takes as input the security parameter 1^λ, and outputs the public key PK and the master secret key MSK of the system.

$sk_{\mathsf{ID}} \leftarrow \mathsf{Ext}(\mathsf{PK}, \mathsf{MSK}, \mathsf{ID})$:
The key extraction algorithm Ext takes as input the public key PK, the master secret key MSK, and an identifier $\mathsf{ID} \in \{0,1\}^*$. It outputs the secret key sk_{ID} capable of decrypting ciphertexts intended for the holder of the given identifier ID.

$c \leftarrow \mathsf{Enc}(\mathsf{PK}, \mathsf{ID}, m)$:
The encryption algorithm Enc algorithm takes as input the public key PK, an identifier $\mathsf{ID} \in \{0,1\}^*$, and a message $m \in \mathcal{MSP}$. It then outputs a ciphertext $c \in \mathcal{CSP}$.

$m/\perp := \mathsf{Dec}(\mathsf{PK}, sk_{\mathsf{ID}}, c)$:
Given the public key PK, a secret key sk_{ID}, and a ciphertext $c \in \mathcal{CSP}$, the decryption algorithm Dec either outputs a message $m \in \mathcal{MSP}$ or the failure symbol $\perp$. We assume that Dec is deterministic.

CORRECTNESS. For every $\mathsf{ID} \in \{0,1\}^*$ and every $m \in \mathcal{MSP}$, if sk_{ID} is the secret key output by $\mathsf{Ext}(\mathsf{PK}, \mathsf{MSK}, \mathsf{ID})$, then $\mathsf{Dec}(\mathsf{PK}, sk_{\mathsf{ID}}, \mathsf{Enc}(\mathsf{PK}, \mathsf{ID}, m)) = m$.

WEAKLY ROBUST AIBE. The *Robust Encryption*, formalized by Abdalla *et al.* [29], requires that it is hard to produce a ciphertext that is valid for two different users. In [29], the authors define two types of robustness, strong and weak. Informally, an AIBE scheme is called *weakly robust*, if any adversary has negligible advantage in producing two identities $\mathsf{ID}_0, \mathsf{ID}_1$ and a message m such that the encryption of m under ID_0 can be decrypted with the private key associated with ID_1 leading to a non-$\perp$ result. In [29], the authors also provide a transformation algorithm which makes possible to obtain a weakly robust AIBE scheme from a regular AIBE one.

III. OUTSIDER-ANONYMOUS BROADCAST ENCRYPTION (oABE)

A. The Setting

Definition 2: An outsider-anonymous broadcast encryption (oABE) scheme, associated with a universe of users $U = \{1, \ldots, N\}$, a message space $\mathcal{MSP}$, and a ciphertext space $\mathcal{CSP}$, is a tuple of probabilistic polynomial algorithms $(\mathsf{Setup}, \mathsf{KeyGen}, \mathsf{Encrypt}, \mathsf{Decrypt})$ such that:

$(\mathsf{PK}, \mathsf{MSK}) \leftarrow \mathsf{Setup}(1^\lambda, N)$:

The Setup algorithm takes as input the security parameter 1^λ and the number of users in the system N. It outputs the public key PK and the master secret key MSK of the system.

$sk_i \leftarrow \mathsf{KeyGen}(\mathsf{PK}, \mathsf{MSK}, i)$:

The key generation algorithm KeyGen takes as input the public key PK, the master secret key MSK, and a user $i \in U$. It outputs the secret key sk_i of user i.

$c \leftarrow \mathsf{Encrypt}(\mathsf{PK}, S, m)$:

The Encrypt algorithm takes as input the public key PK, the set of receivers $S \subseteq U$, and a message $m \in \mathcal{MSP}$. It then outputs a ciphertext $c \in \mathcal{CSP}$.

$m/\perp := \mathsf{Decrypt}(\mathsf{PK}, sk_i, c)$:

Given the public key PK, a secret key sk_i, and a ciphertext $c \in \mathcal{CSP}$, the Decrypt algorithm either outputs a message $m \in \mathcal{MSP}$ or the failure symbol $\perp$. We assume that Decrypt is deterministic.

CORRECTNESS. For every $S \subseteq U$, every $i \in S$, and every $m \in \mathcal{MSP}$, if sk_i is the secret key output by $\mathsf{KeyGen}(\mathsf{PK}, \mathsf{MSK}, i)$ then $\mathsf{Decrypt}(\mathsf{PK}, sk_i, \mathsf{Encrypt}(\mathsf{PK}, S, m)) = m$.

Notice that the decryption algorithm in the above definition does not require the set of recipients S as an input. We stress that this is crucial for providing any level of anonymity in a broadcast encryption scheme.

B. The Security Model

DEGREES OF ANONYMITY. The degree of recipient-set anonymity captured in our security model, which we call *outsider-anonymity*, lies between the complete lack of protection that characterizes traditional broadcast encryption schemes as introduced in [2], [14], and the full anonymity provided in schemes such as [15], [16]. In an oABE scheme, when the adversary receives a ciphertext of which she is not a legal recipient, she will be unable to learn anything about the identities of the legal recipients (let alone the contents of the ciphertext). Still, for those ciphertexts for which the adversary is in the authorized set of recipients, she might also learn the identities of some the other legal recipients. This seems a natural relaxation, since often the *contents* of the communication already reveals something about the recipient set. At the same time, our new intermediate definition of security might allow the construction of more efficient anonymous broadcast encryption schemes; for example, in Sect. IV we describe the first broadcast encryption scheme with sublinear ciphertexts that attains some meaningful recipient-set anonymity guarantees.

CCA SECURITY. We now present the security requirements for a broadcast encryption scheme to be *outsider anonymous* against chosen-ciphertext attacks (CCA). First we define the CCA of an oABE scheme as a game, which we term oABE-IND-CCA, played between a probabilistic polynomial time (PPT) adversary $\mathcal{A}$ and a challenger $\mathcal{C}$. The security requirement is that $\mathcal{A}$'s advantage of winning the oABE-IND-CCA game is negligible. The high-level idea of this game is for any two sets of recipients $S_0, S_1 \in U$, $\mathcal{A}$ cannot distinguish between a ciphertext intended for the recipient set S_0 and a ciphertext intended for the recipient set S_1 given the fact that the $\mathcal{A}$ does not possess the secret key of any user in $S_0 \cup S_1$. We require the two sets S_0, S_1 be the same size in order to avoid trivial attacks. The formal definitions follow.

Definition 3: The oABE-IND-CCA game defined for an oABE scheme $\Pi = (\mathsf{Setup}, \mathsf{KeyGen}, \mathsf{Encrypt}, \mathsf{Decrypt})$, a PPT adversary $\mathcal{A}$, and a challenger $\mathcal{C}$ is as follows:

Setup:

$\mathcal{C}$ runs $(\mathsf{PK}, \mathsf{MSK}) \leftarrow \mathsf{Setup}(1^\lambda, N)$ and gives $\mathcal{A}$ the resulting public key PK, keeping the master secret key MSK to itself. $\mathcal{C}$ also initializes the set of revoked users Rev to be empty.

Phase 1:

$\mathcal{A}$ adaptively issues queries $q_1, \ldots, q_m$ where each q_i is one of the following:

- Secret key query i:
  $\mathcal{A}$ requests the secret key of the user $i \in U$. $\mathcal{C}$ runs $sk_i \leftarrow \mathsf{KeyGen}(\mathsf{PK}, \mathsf{MSK}, i)$ to generate the secret key sk_i of the user i, adds i to Rev, and sends sk_i to $\mathcal{A}$.
- Decryption query (i, c):
  $\mathcal{A}$ issues a decryption query where $i \in U$ and $c \in \mathcal{CSP}$. First, $\mathcal{C}$ runs $sk_i \leftarrow \mathsf{KeyGen}(\mathsf{PK}, \mathsf{MSK}, i)$ to generate the secret key sk_i of the user i. Then, it runs $\mathsf{Decrypt}(\mathsf{PK}, sk_i, c)$ and gives the output to $\mathcal{A}$.

Challenge:

$\mathcal{A}$ gives $\mathcal{C}$ two equal length messages $m_0, m_1 \in \mathcal{MSP}$ and two equal length sets of user identities $S_0, S_1 \subseteq U$ with the restriction that $\mathsf{Rev} \cap (S_0 \cup S_1) = \emptyset$. $\mathcal{C}$ picks a random bit $b \in \{0, 1\}$, runs $c^* \leftarrow \mathsf{Encrypt}(\mathsf{PK}, S_b, m_b)$, and sends c^* to $\mathcal{A}$.

Phase 2:

$\mathcal{A}$ adaptively issues additional queries $q_{m+1}, \ldots, q_n$ where each q_i is one of the following:

- Secret key query i such that $i \notin S_0 \cup S_1$.
- Decryption query (i, c) such that, if $i \in S_0 \cup S_1$, then $c \neq c^*$.

In both cases, $\mathcal{C}$ responds as in *Phase 1*.

Guess:

$\mathcal{A}$ outputs a guess $b' \in \{0, 1\}$ and wins if $b' = b$.

We refer to such an adversary $\mathcal{A}$ as an oABE-IND-CCA adversary. The advantage of $\mathcal{A}$ winning the above game is defined as,

$$\mathsf{Adv}_{\mathcal{A},\Pi}^{\text{oABE-IND-CCA}} = \left|\Pr\left[b' = b\right] - \tfrac{1}{2}\right|$$

The probability is over the random bits used by the adversary $\mathcal{A}$ and the challenger $\mathcal{C}$.

Definition 4: An oABE scheme $\Pi = (\mathsf{Setup}, \mathsf{KeyGen}, \mathsf{Encrypt}, \mathsf{Decrypt})$ is $(t,\, q_{sk},\, q_d,\, \epsilon)$-secure if for any t-time oABE-IND-CCA adversary $\mathcal{A}$ making at most q_{sk} chosen secret key queries and at most q_d chosen decryption queries, we have that $\mathsf{Adv}_{\mathcal{A},\Pi}^{\text{oABE-IND-CCA}} \leq \epsilon$. As a shorthand, we say that Π is $(t, q_{sk}, q_d, \epsilon)$-oABE-IND-CCA secure.

CPA SECURITY. The chosen plaintext attack (CPA) of an oABE scheme is defined similar to the oABE-IND-CCA game with the restriction that the adversary is not allowed to issue any decryption queries during *Phase 1* and *Phase 2*. The adversary is still allowed to issue secret key queries. The CPA security game is termed oABE-IND-CPA.

Definition 5: An oABE scheme $\Pi = (\mathsf{Setup}, \mathsf{KeyGen}, \mathsf{Encrypt}, \mathsf{Decrypt})$ is (t, q_{sk}, ϵ)-oABE-IND-CPA secure if Π is $(t, q_{sk}, 0, \epsilon)$-oABE-IND-CCA secure.

Remark 6: Our definition of security of an outsider-anonymous broadcast encryption scheme can be easily transformed to a definition of security of a fully anonymous broadcast encryption scheme by changing the restriction in the challenge phase, which is currently $\mathsf{Rev} \cap (S_0 \cup S_1) = \emptyset$, to $\mathsf{Rev} \cap (S_0 \triangle S_1) = \emptyset$.[1]

IV. Our Constructions

We now present our constructions of outsider-anonymous broadcast encryption (oABE) schemes. In a nutshell, the key point of our constructions is to combine an anonymized version of the public-key extension by Dodis and Fazio [8] of the CS method by Naor *et al.* [6] with a fully secure weakly robust AIBE scheme such as [22]. Notice that our approach can be seen as a *framework* for achieving an oABE scheme by using any weakly robust AIBE scheme as an underlying primitive.

The ciphertext length in all constructions is $\mathcal{O}\left(r \log\left(\frac{N}{r}\right)\right)$ times the ciphertext length of the underlying AIBE scheme, and the user secret key length is $\mathcal{O}\left(\log N\right)$ times the user secret key length of the underlying AIBE scheme, where r is the number of revoked users and N is the total number of users in the system.

We provide two generic public-key constructions: a CPA secure construction in Sect. IV-A and a CCA secure construction in Sect. IV-B. The limitation with both of these constructions is that on average, the Decrypt algorithm attempts $\mathcal{O}\left(r \log\left(\frac{N}{r}\right) \log N\right)$ decryption operations of the underlying AIBE scheme. In Sect. IV-C, we present an enhanced CCA secure construction in which for a given oABE ciphertext, the Decrypt algorithm executes a single AIBE decryption operation. The proofs of all the constructions are provided in the full version of this paper [30], [31].

For the simplicity of exposition, our constructions encrypt the actual message m. The ciphertext length could be further reduced by using a hybrid encryption where m is encrypted using a symmetric key encryption algorithm with a symmetric key k, and k is then encrypted using the oABE scheme.

In all constructions, $\mathcal{T}$ denotes the binary tree of N users in the system with respect to the CS method. For simplicity, we assume that $N = 2^n$.

A. A Generic CPA Secure Public-Key Construction

Given a weakly robust AIBE scheme $\Pi' = (\mathsf{Init}, \mathsf{Ext}, \mathsf{Enc}, \mathsf{Dec})$, we construct an oABE-IND-CPA secure scheme $\Pi = (\mathsf{Setup}, \mathsf{KeyGen}, \mathsf{Encrypt}, \mathsf{Decrypt})$ in Construction 1.

The correctness of Construction 1 follows from the correctness of the underlying AIBE scheme. In Theorem 7, we establish the security of this construction based on the security of the underlying AIBE scheme.

Theorem 7: If $\Pi' = (\mathsf{Init}, \mathsf{Ext}, \mathsf{Enc}, \mathsf{Dec})$ is (t, q_{sk}, ϵ)-AIBE-IND-CPA secure, then Construction 1 is $(t, q_{sk}, 2\,\epsilon\, r \log\left(\frac{N}{r}\right))$-oABE-IND-CPA secure.

PARAMETERS. When Construction 1 is instantiated with Gentry's Fully Secure IBE scheme in the CPA setting [22], we obtain the following parameter lengths. MSK is just one element in $\mathbb{Z}_p$ and the integer N. PK is only 3 group elements in $\mathbb{G}$. The user secret key consists of $(\log N + 1)$ elements in $\mathbb{Z}_p$ and $(\log N + 1)$ elements in $\mathbb{G}$. The ciphertext consists of $\left\lfloor r \log\left(\frac{N}{r}\right)\right\rfloor$ elements in $\mathbb{G}$ and $2\left\lfloor r \log\left(\frac{N}{r}\right)\right\rfloor$ elements in $\mathbb{G}_\mathrm{T}$. Also notice that the Enc algorithm in Gentry's AIBE scheme does not require any pairing computations since they can be pre-computed.

B. A Generic CCA Secure Public-Key Construction

Given a weakly robust AIBE scheme $\Pi' = (\mathsf{Init}, \mathsf{Ext}, \mathsf{Enc}, \mathsf{Dec})$ and a strongly existentially unforgeable one-time signature scheme $\Sigma = (\mathsf{Gen}, \mathsf{Sign}, \mathsf{Vrfy})$, we construct an oABE-IND-CCA secure scheme $\Pi = (\mathsf{Setup}, \mathsf{KeyGen}, \mathsf{Encrypt}, \mathsf{Decrypt})$ in Construction 2.

The correctness of Construction 2 follows from the correctness of the underlying Σ and AIBE schemes. Next, in Theorem 8, we establish the security of this construction based on the security of the underlying Σ and AIBE schemes.

Theorem 8: If $\Sigma = (\mathsf{Gen}, \mathsf{Sign}, \mathsf{Vrfy})$ is (t, ϵ_1)-strongly existentially unforgeable and $\Pi' = (\mathsf{Init}, \mathsf{Ext}, \mathsf{Enc}, \mathsf{Dec})$ is $(t, q_{sk}, q_d, \epsilon_2)$-AIBE-IND-CCA secure, then Construction 2 is $(t, q_{sk}, q_d, 2(\epsilon_1 + \epsilon_2)\, r \log\left(\frac{N}{r}\right))$-oABE-IND-CCA secure.

[1] For any two sets S_0, S_1, their symmetric difference is denoted by $S_0 \triangle S_1$.

Construction 1 Our Generic CPA Secure PK Construction

```
algorithm Setup($1^\lambda, N$)
    $(\mathsf{PK}', \mathsf{MSK}') \leftarrow \mathsf{Init}(1^\lambda)$
    $\mathsf{PK} = (\mathsf{PK}', N)$
    $\mathsf{MSK} = \mathsf{MSK}'$
    return $\mathsf{PK}, \mathsf{MSK}$

algorithm KeyGen($\mathsf{PK}, \mathsf{MSK}, i$)
    ▷ Let $\mathsf{HID}_i$ be the hierarchical identifier of $i$ in $\mathcal{T}$
    for $j = 1$ to $n + 1$ do
        $sk_{i,j} \leftarrow \mathsf{Ext}(\mathsf{PK}', \mathsf{MSK}', \mathsf{HID}_{i|j})$
    $sk_i = (sk_{i,1}, \ldots, sk_{i,n+1})$
    return $sk_i$

algorithm Encrypt($\mathsf{PK}, S, m$)
    $r = N - |S|$
    $L = \left\lfloor r \log\left(\frac{N}{r}\right)\right\rfloor$
    ▷ Let $\mathsf{Cov}$ be the set of subtrees covering $S$ in $\mathcal{T}$
    for $j = 1$ to $|\mathsf{Cov}|$ do
        ▷ Let $T$ be the subtree in $\mathsf{Cov}$ indexed by $j$
        ▷ Let $\mathsf{HID}$ be the hierarchical identifier of $T$
        $c_j \leftarrow \mathsf{Enc}(\mathsf{PK}', \mathsf{HID}, m)$
    $\widetilde{m} \xleftarrow{\$} \{0,1\}^{|m|}$
    for $j = |\mathsf{Cov}| + 1$ to $L$ do
        ▷ dummy is a special ID used for padding ciphertexts
        $c_j \leftarrow \mathsf{Enc}(\mathsf{PK}', \mathtt{dummy}, \widetilde{m})$
    ▷ Let $\pi$ be a random permutation
    $c = \left(c_{\pi(1)}, \ldots, c_{\pi(L)}\right)$
    return $c$

algorithm Decrypt($\mathsf{PK}, sk_i, c$)
    Parse $sk_i$ as $(sk_{i,1}, \ldots, sk_{i,n+1})$
    Parse $c$ as $(c_1, \ldots, c_L)$
    for $k = 1$ to $n + 1$ do
        for $j = 1$ to $L$ do
            $m \leftarrow \mathsf{Dec}(\mathsf{PK}', sk_{i,k}, c_j)$
            if $m \neq \perp$ then
                return $m$
    return $\perp$
```

Construction 2 Our Generic CCA Secure PK Construction

```
algorithm Setup($1^\lambda, N$)
    $(\mathsf{PK}', \mathsf{MSK}') \leftarrow \mathsf{Init}(1^\lambda)$
    $\mathsf{PK} = (\mathsf{PK}', N)$
    $\mathsf{MSK} = \mathsf{MSK}'$
    return $\mathsf{PK}, \mathsf{MSK}$

algorithm KeyGen($\mathsf{PK}, \mathsf{MSK}, i$)
    ▷ Let $\mathsf{HID}_i$ be the hierarchical identifier of $i$ in $\mathcal{T}$
    for $j = 1$ to $n + 1$ do
        $sk_{i,j} \leftarrow \mathsf{Ext}(\mathsf{PK}', \mathsf{MSK}', \mathsf{HID}_{i|j})$
    $sk_i = (sk_{i,1}, \ldots, sk_{i,n+1})$
    return $sk_i$

algorithm Encrypt($\mathsf{PK}, S, m$)
    $(\mathsf{VK}, \mathsf{SK}) \leftarrow \mathsf{Gen}(1^\lambda)$
    $r = N - |S|$
    $L = \left\lfloor r \log\left(\frac{N}{r}\right)\right\rfloor$
    ▷ Let $\mathsf{Cov}$ be the set of subtrees covering $S$ in $\mathcal{T}$
    for $j = 1$ to $|\mathsf{Cov}|$ do
        ▷ Let $T$ be the subtree in $\mathsf{Cov}$ indexed by $j$
        ▷ Let $\mathsf{HID}$ be the hierarchical identifier of $T$
        $c_j \leftarrow \mathsf{Enc}(\mathsf{PK}', \mathsf{HID}, \mathsf{VK}||m)$
    $\widetilde{m} \xleftarrow{\$} \{0,1\}^{|\mathsf{VK}||m|}$
    for $j = |\mathsf{Cov}| + 1$ to $L$ do
        ▷ dummy is a special ID used for padding ciphertexts
        $c_j \leftarrow \mathsf{Enc}(\mathsf{PK}', \mathtt{dummy}, \widetilde{m})$
    ▷ Let $\pi$ be a random permutation
    $c = \left(c_{\pi(1)}, \ldots, c_{\pi(L)}\right)$
    $\sigma \leftarrow \mathsf{Sign}(\mathsf{SK}, \mathsf{VK}||c)$
    $C = \sigma||c$
    return $C$

algorithm Decrypt($\mathsf{PK}, sk_i, c$)
    Parse $sk_i$ as $(sk_{i,1}, \ldots, sk_{i,n+1})$
    Parse $C$ as $\sigma||c$
    Parse $c$ as $(c_1, \ldots, c_L)$
    for $k = 1$ to $n + 1$ do
        for $j = 1$ to $L$ do
            $m' \leftarrow \mathsf{Dec}(\mathsf{PK}', sk_{i,k}, c_j)$
            if $m' = \mathsf{VK}||m \wedge \mathsf{Vrfy}(\mathsf{VK}, \sigma, \mathsf{VK}||c)$ then
                return $m$
    return $\perp$
```

PARAMETERS. The parameter lengths of Construction 2 when instantiated with Gentry's Fully Secure IBE scheme in the CCA setting [22] are as follows. MSK is one element in $\mathbb{Z}_p$ and the integer N. PK consists of 5 group elements in $\mathbb{G}$ and the definition of a hash function H from a family of universal one-way hash functions. The user secret key consists of $3(\log N + 1)$ elements in $\mathbb{Z}_p$ and $3(\log N + 1)$ elements in $\mathbb{G}$. The ciphertext consists of $\left\lfloor r \log\left(\frac{N}{r}\right)\right\rfloor$ elements in $\mathbb{G}$ and $3\left\lfloor r \log\left(\frac{N}{r}\right)\right\rfloor$ elements in $\mathbb{G}_\mathrm{T}$. Similar to Gentry's CPA secure AIBE construction, the Enc algorithm in the CCA secure construction does not require any pairing computations since they can be pre-computed.

C. *An Enhanced* CCA *Secure Public-Key Construction*

The main limitation of our generic public-key constructions is the running time of the decryption algorithm. As described in the opening paragraphs of Sect. IV, decryption amounts to performing $\mathcal{O}\left(r \log\left(\frac{N}{r}\right) \log N\right)$ AIBE decryption attempts on average. The root cause behind this limitation is the decryption process's inability to identify the correct AIBE ciphertext component efficiently. In this section, we describe an enhancement of our generic public-key construction under the Diffie-Hellman assumption, in the random oracle model. The main idea of this enhancement is to adapt the techniques of [15] to the structure of our ciphertexts and attach a unique tag to each AIBE ciphertext component of a given oABE ciphertext. With this optimization, the Decrypt algorithm is able to identify the correct AIBE ciphertext component via a linear search through the whole oABE ciphertext components, at which point a single AIBE decryption operation suffices to recover the original plaintext. This yields an asymptotic decryption time of $\mathcal{O}\left(r \log\left(\frac{N}{r}\right) \log N\right)$, but in fact this is in a sense an overestimate, since the cost of searching for the correct ciphertext component is much less than carrying out multiple decryption attempts.

Given a weakly robust AIBE scheme $\Pi' = (\mathsf{Init}, \mathsf{Ext}, \mathsf{Enc}, \mathsf{Dec})$ and a strongly existentially unforgeable one-time signature scheme $\Sigma = (\mathsf{Gen}, \mathsf{Sign}, \mathsf{Vrfy})$, we construct an enhanced oABE-IND-CCA secure scheme $\Pi = (\mathsf{Setup}, \mathsf{KeyGen}, \mathsf{Encrypt}, \mathsf{Decrypt})$ in Construction 3. In this construction, $\mathbb{G} = \langle g \rangle$ denotes a group with prime order $q > 2^\lambda$ in which CDH is hard and DDH is easy and g is a group generator. $H' : \mathbb{G} \to \{0,1\}^\lambda$ is a cryptographic hash function that will be modeled as a random oracle in the security analysis.

Theorem 9: If $\Sigma = (\mathsf{Gen}, \mathsf{Sign}, \mathsf{Vrfy})$ is (t, ϵ_1)-strongly existentially unforgeable, $\Pi' = (\mathsf{Init}, \mathsf{Ext}, \mathsf{Enc}, \mathsf{Dec})$ is $(t, q_{sk}, q_d, \epsilon_2)$-AIBE-IND-CCA secure, and CDH is (t, ϵ_3)-hard in $\mathbb{G}$ and DDH is efficiently computable in $\mathbb{G}$, then Construction 3 is $(t, q_{sk}, q_d, 2(\epsilon_1+\epsilon_2+\epsilon_3) r \log\left(\frac{N}{r}\right))$-oABE-IND-CCA secure, in the random oracle model.

Remark 10: Notice that the check in Line 3.48 can be performed in expected time $\mathcal{O}(n+L) = \mathcal{O}(L)$, *e.g.*, using a hash table $\mathrm{H_T}$ to compute the intersection between $\{tag_k\}_{k \in [1,n+1]}$ and $\{\overline{c}_j\}_{j \in [1,L]}$ as follows:

```
Initialize $\mathrm{H_T}$ to be empty
for $k = 1$ to $n+1$ do
    Insert $(tag_k, k)$ in $\mathrm{H_T}$
for $j = 1$ to $L$ do
    Look up an entry of the form $(\overline{c}_j, k)$ in $\mathrm{H_T}$
    if found then
        return $k$
```

Remark 11: Using the twin Diffie-Hellman methodology [32] via techniques similar to [16], it is possible to modify Construction 3 to get an outsider-anonymous broadcast encryption scheme that is adaptive CCA secure, in the standard model, under the decisional Diffie-Hellman assumption.

D. An Enhanced CCA *Secure Symmetric-Key Construction*

The enhanced CCA secure public key construction achieves a major performance gain in the Decrypt algorithm compared to the generic CCA secure construction, but it also changes the length of the public key from $\mathcal{O}(1)$ to $\mathcal{O}(N)$. This increase in public key length may not be a concern for many practical constructions, since the public key can be stored as a static data file on a server on the Internet and also in users' computers. Still, for the symmetric-key setting it is possible to accommodate storage-sensitive systems and attain constant key storage at the Center, while maintaining efficient decryption and logarithmic storage at the receivers.

In particular, recall from Sect. II-A that in the symmetric-key setting, only the Center can broadcast messages to the receivers. Thus, the $\mathcal{O}(N)$ information from which the tags for efficient decryption are created does not need to be published. Therefore, this information can be compressed into $\mathcal{O}(1)$ key storage using a standard trick based on any length-tripling pseudo-random number generator G (*cf. e.g.*, the SD method of Naor *et al.* [6]). In other words, the random exponents associated with the subtrees of $\mathcal{T}$ (*cf.* Sect. IV-C) are now pseudo-randomly generated from a single seed, by repeated invocations of G on the left or right third of the result of the previous iteration, based on the path to the root of the subtree at hand. Finally, upon reaching the subtree root, the

Construction 3 Our Enhanced CCA Secure PK Construction

```
 1: algorithm $\mathsf{Setup}(1^\lambda, N)$
 2:     $(\mathsf{PK}', \mathsf{MSK}') \leftarrow \mathsf{Init}(1^\lambda)$
 3:     ▷ Let $\mathsf{Fam}$ be the set of all the subtrees in $\mathcal{T}$
 4:     for $j = 1$ to $|\mathsf{Fam}|$ do
 5:         ▷ Let $T$ be the subtree in $\mathsf{Fam}$ indexed by $j$
 6:         ▷ Let $\mathsf{HID}$ be the hierarchical identifier of $T$
 7:         $a_{\mathsf{HID}} \xleftarrow{\$} \mathbb{Z}_q^*$
 8:         $A_{\mathsf{HID}} = g^{a_{\mathsf{HID}}}$
 9:     $\mathsf{PK} = \left(\mathsf{PK}', N, \mathbb{G}, g, \{A_{\mathsf{HID}}\}_{\mathsf{HID} \in \mathcal{T}}\right)$
10:     $\mathsf{MSK} = \left(\mathsf{MSK}', \{a_{\mathsf{HID}}\}_{\mathsf{HID} \in \mathcal{T}}\right)$
11:     return $\mathsf{PK}, \mathsf{MSK}$

12: algorithm $\mathsf{KeyGen}(\mathsf{PK}, \mathsf{MSK}, i)$
13:     ▷ Let $\mathsf{HID}_i$ be the hierarchical identifier of $i$ in $\mathcal{T}$
14:     for $j = 1$ to $n+1$ do
15:         $\overline{sk}_{i,j} = a_{\mathsf{HID}_{i|j}}$
16:         $sk_{i,j} \leftarrow \mathsf{Ext}(\mathsf{PK}', \mathsf{MSK}', \mathsf{HID}_{i|j})$
17:     $sk_i = \left(\left(\overline{sk}_{i,1}, sk_{i,1}\right), \ldots, \left(\overline{sk}_{i,n+1}, sk_{i,n+1}\right)\right)$
18:     return $sk_i$

19: algorithm $\mathsf{Encrypt}(\mathsf{PK}, S, m)$
20:     $(\mathsf{VK}, \mathsf{SK}) \leftarrow \mathsf{Gen}(1^\lambda)$
21:     $r = N - |S|$
22:     $L = \left\lfloor r \log\left(\frac{N}{r}\right)\right\rfloor$
23:     $s \xleftarrow{\$} \mathbb{Z}_q^*$
24:     $\overline{c}_0 = g^s$
25:     ▷ Let $\mathsf{Cov}$ be the set of subtrees covering $S$ in $\mathcal{T}$
26:     for $j = 1$ to $|\mathsf{Cov}|$ do
27:         ▷ Let $T$ be the subtree in $\mathsf{Cov}$ indexed by $j$
28:         ▷ Let $\mathsf{HID}$ be the hierarchical identifier of $T$
29:         $\overline{c}_j = H'(A_{\mathsf{HID}}^s)$
30:         $c_j \leftarrow \mathsf{Enc}(\mathsf{PK}', \mathsf{HID}, \mathsf{VK}||A_{\mathsf{HID}}^s||m)$
31:     $\widetilde{m} \xleftarrow{\$} \{0,1\}^{|\mathsf{VK}||\overline{c}_0||m|}$
32:     for $j = |\mathsf{Cov}| + 1$ to $L$ do
33:         $s_j \xleftarrow{\$} \mathbb{Z}_q^*$
34:         $\overline{c}_j = H'(g^{s_j})$
35:         ▷ $\mathtt{dummy}$ is a special ID used for padding ciphertexts
36:         $c_j \leftarrow \mathsf{Enc}(\mathsf{PK}', \mathtt{dummy}, \widetilde{m})$
37:     ▷ Let $\pi$ be a random permutation
38:     $c = \left(\overline{c}_0, \left(\overline{c}_{\pi(1)}, c_{\pi(1)}\right), \ldots, \left(\overline{c}_{\pi(L)}, c_{\pi(L)}\right)\right)$
39:     $\sigma \leftarrow \mathsf{Sign}(\mathsf{SK}, \mathsf{VK}||c)$
40:     $C = \sigma||c$
41:     return $C$

42: algorithm $\mathsf{Decrypt}(\mathsf{PK}, sk_i, c)$
43:     Parse $sk_i$ as $\left(\left(\overline{sk}_{i,1}, sk_{i,1}\right), \ldots, \left(\overline{sk}_{i,n+1}, sk_{i,n+1}\right)\right)$
44:     Parse $C$ as $\sigma||c$
45:     Parse $c$ as $\left(\overline{c}_0, \left(\overline{c}_{\pi(1)}, c_{\pi(1)}\right), \ldots, \left(\overline{c}_{\pi(L)}, c_{\pi(L)}\right)\right)$
46:     for $k = 1$ to $n+1$ do
47:         $tag_k = H'(\overline{c}_0^{\,\overline{sk}_{i,k}})$
48:     if $\exists\ k \in [1, n+1],\ \exists\ j \in [1, L]\ tag_k = \overline{c}_j$ then
49:         $m' \leftarrow \mathsf{Dec}(\mathsf{PK}', sk_{i,k}, c_j)$
50:         $\overline{m} = \overline{c}_0^{\,\overline{sk}_{i,k}}$
51:         if $m' = \mathsf{VK}||\overline{m}||m \ \wedge\ \mathsf{Vrfy}(\mathsf{VK}, \sigma, \mathsf{VK}||c)$ then
52:             return $m$
53:     return $\perp$
```

middle third of the pseudorandom output is used to generate the required exponent.

V. Conclusions and Future Work

In this work, we introduced the notion of outsider-anonymity in the broadcast encryption setting and showed that it enables efficient constructions of broadcast encryption schemes with sublinear communication complexity and meaningful anonymity guarantees. It remains an interesting open problem to construct receiver-anonymous broadcast encryption schemes that at once afford full anonymity to the receivers and attain performance levels comparable to those of standard broadcast encryption systems. While the lower bounds of [17] suggest that the trade-off proposed in this paper may be unavoidable, it is still conceivable that alternative techniques may be able to escape the notion of "atomicity" of [17] by somehow making the decryption algorithm "oblivious" to the specific key used. One approach might be to devise a mechanism whereby users are only given 'packaged' keysets, and decryption succeeds only when the keyset is provided without alterations. The challenge here is to prevent users from tampering with the decryption algorithm.

Acknowledgments

This research was sponsored in part by NSF award #1117675 and by the U.S. Army Research Laboratory and the U.K. Ministry of Defence and was accomplished under Agreement Number W911NF-06-3-0001. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the U.S. Army Research Laboratory, the U.S. Government, the U.K. Ministry of Defence or the U.K. Government. The U.S. and U.K. Governments are authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation hereon. This research project was also partially supported by PSC-CUNY Awards 63356-00 41 and 64578-00 42, jointly funded by the Professional Staff Congress and CUNY.

References

[1] S. Berkovits, "How to broadcast a secret," in *Advances in Cryptology—EUROCRYPT 1991*, 1991, pp. 535–541.

[2] A. Fiat and M. Naor, "Broadcast encryption," in *Advances in Cryptology—CRYPTO 1993*, 1993, pp. 480–491.

[3] AACS, "Advanced access content system," http://www.aacsla.com/.

[4] D. Boneh, E.-J. Goh, N. Modadugu, and H. Shacham, "Sirius: Securing remote untrusted storage," in *ISOC Network and Distributed System Security Symposium—NDSS 2003*, 2003, pp. 131–145.

[5] J. A. Garay, J. Staddon, and A. Wool, "Long-lived broadcast encryption," in *Advances in Cryptology—CRYPTO 2000*, 2000, pp. 333–352.

[6] D. Naor, M. Naor, and J. Lotspiech, "Revocation and tracing schemes for stateless receivers," in *Advances in Cryptology—CRYPTO 2001*, 2001, pp. 41–62.

[7] D. Halevy and A. Shamir, "The LSD broadcast encryption scheme," in *Advances in Cryptology—CRYPTO 2002*, 2002, pp. 47–60.

[8] Y. Dodis and N. Fazio, "Public-key broadcast encryption for stateless receivers," in *Digital Rights Management—DRM 2002*, 2002, pp. 61–80.

[9] ——, "Public-key trace and revoke scheme secure against adaptive chosen ciphertext attack," in *Public Key Cryptography—PKC 2003*, 2003, pp. 100–115.

[10] Y. Dodis, N. Fazio, A. Kiayias, and M. Yung, "Scalable public-key tracing and revoking," in *ACM Symposium on Principles of Distributed Computing—PODC 2003*, 2003, pp. 190–199, invited to the Special Issue of Journal of Distributed Computing PODC 2003.

[11] Y. Dodis, N. Fazio, A. Lysyanskaya, and D. Yao, "ID-based encryption for complex hierarchies with applications to forward security and broadcast encryption," in *ACM Conference on Computer and Communications Security—CCS 2004*, 2004, pp. 354–363.

[12] D. Boneh, C. Gentry, and B. Waters, "Collusion resistant broadcast encryption with short ciphertexts and private keys," in *Advances in Cryptology—CRYPTO 2005*, 2005, pp. 258–275.

[13] D. Boneh and B. Waters, "A fully collusion resistant broadcast, trace, and revoke system," in *ACM Conference on Computer and Communications Security—CCS 2006*, 2006, pp. 211–220.

[14] C. Gentry and B. Waters, "Adaptive security in broadcast encryption systems (with short ciphertexts)," in *Advances in Cryptology—EUROCRYPT 2009*, 2009, pp. 171–188.

[15] A. Barth, D. Boneh, and B. Waters, "Privacy in encrypted content distribution using private broadcast encryption," in *Financial Cryptography and Data Security—FC 2006*, 2006, pp. 52–64.

[16] B. Libert, K. G. Paterson, and E. A. Quaglia, "Anonymous broadcast encryption," in *Public Key Cryptography—PKC 2012*, 2012, pp. 206–224.

[17] A. Kiayias and K. Samari, "Lower bounds for private broadcast encryption," in *Information Hiding—IH 2012*, 2012, pp. 176–190.

[18] L. Krzywiecki, P. Kubiak, and M. Kutylowski, "A revocation scheme preserving privacy," in *Information Security and Cryptology—Inscrypt 2006*, 2006, pp. 130–143.

[19] K. Ren, S. Yu, and W. Lou, "Attribute-based on-demand multicast group setup with receiver anonymity," in *Security and Privacy in Communication Networks—SecureComm 2008*, 2008, pp. 18:1–18:6.

[20] S. Jarecki and X. Liu, "Unlinkable secret handshakes and key-private group key management schemes," in *Applied Cryptography and Network Security—ACNS 2007*, 2007, pp. 270–287.

[21] M. Abdalla, M. Bellare, D. Catalano, E. Kiltz, T. Kohno, T. Lange, J. Malone-Lee, G. Neven, P. Paillier, and H. Shi, "Searchable encryption revisited: Consistency properties, relation to Anonymous IBE, and extensions," in *Advances in Cryptology—CRYPTO 2005*, 2005, pp. 205–222.

[22] C. Gentry, "Practical identity-based encryption without random oracles," in *Advances in Cryptology—EUROCRYPT 2006*, 2006, pp. 445–464.

[23] A. Shamir, "Identity-based cryptosystems and signature schemes," in *Advances in Cryptology—CRYPTO 1984*, 1984, pp. 47–53.

[24] D. Boneh and M. K. Franklin, "Identity-based encryption from the weil pairing," in *Advances in Cryptology—CRYPTO 2001*, 2001, pp. 213–229.

[25] D. Boneh and X. Boyen, "Secure identity based encryption without random oracles," in *Advances in Cryptology—CRYPTO 2004*, 2004, pp. 443–459.

[26] D. Boneh, C. Gentry, and M. Hamburg, "Space-efficient identity based encryption without pairings," in *IEEE Symposium on Foundations of Computer Science—FOCS 2007*, 2007, pp. 647–657.

[27] B. Waters, "Dual system encryption: Realizing fully secure IBE and HIBE under simple assumptions," in *Advances in Cryptology—CRYPTO 2009*, 2009, pp. 619–636.

[28] X. Boyen and B. Waters, "Anonymous hierarchical identity-based encryption (without random oracles)," in *Advances in Cryptology—CRYPTO 2006*, 2006, pp. 290–307.

[29] M. Abdalla, M. Bellare, and G. Neven, "Robust encryption," in *Theory of Cryptography—TCC 2010*, 2010, pp. 480–497.

[30] N. Fazio and I. M. Perera, "Outsider-anonymous broadcast encryption with sublinear ciphertexts," in *Public Key Cryptography—PKC 2012*, 2012, pp. 225–242.

[31] ——, "Outsider-anonymous broadcast encryption with sublinear ciphertexts," Cryptology ePrint Archive, Report 2012/129, 2012, full Version of [30].

[32] D. Cash, E. Kiltz, and V. Shoup, "The twin Diffie-Hellman problem and applications," in *Advances in Cryptology—EUROCRYPT 2008*, 2008, pp. 127–145.