Anonymous Broadcast Encryption: Adaptive Security and Efficient Constructions in the Standard Model

Benoît Libert[1], Kenneth G. Paterson[2], and Elizabeth A. Quaglia[2]

[1] Université catholique de Louvain, ICTEAM Institute, Belgium
[2] Information Security Group, Royal Holloway, University of London, U.K.

Abstract. In this paper we consider *anonymity* in the context of Broadcast Encryption (BE). This issue has received very little attention so far and *all but one* of the currently available BE schemes fail to provide anonymity. Yet, we argue that it is intrinsically desirable to provide anonymity in standard applications of BE and that it can be achieved at a moderate cost. We provide a security definition for Anonymous Broadcast Encryption (ANOBE) and show that it is achievable assuming only the existence of IND-CCA secure public key encryption (PKE). Focusing on reducing the size of ciphertexts, we then give two generic constructions for ANOBE. The first is from any anonymous (key-private) IND-CCA secure PKE scheme, and the second is from any IBE scheme that satisfies a weak security notion in the multi-TA setting. Furthermore, we show how randomness re-use techniques can be deployed in the ANOBE context to reduce computational and communication costs, and how a new cryptographic primitive – anonymous hint systems – can be used to speed up the decryption process in our ANOBE constructions. All of our results are in the standard model, achieving fully collusion-resistant ANOBE schemes secure against *adaptive* IND-CCA adversaries.

Keywords: Broadcast Encryption, Anonymity

1 Introduction

Anonymity. In a world that is increasingly relying on digital technologies, addressing the issue of protecting users' privacy is of crucial importance. This is reflected by the great attention given to *anonymity* in all the main fields of modern cryptography. In the area of Public-Key Encryption (PKE), anonymity is often referred to as key-privacy [6]. This notion captures the property that an eavesdropper is not able to tell under which one of several public keys a ciphertext was created. The analogous concept in the ID-based setting was studied in [1]. The benefit of preserving receivers' privacy is relevant in more elaborate systems involving for example Hierarchical IBE [12], Attribute-Based Encryption (ABE) or Predicate Encryption [26], where achieving anonymity guarantees becomes increasingly challenging. Furthermore, in the context of digital signatures, a number of primitives effectively *rely* on anonymity: group signatures [16] and anonymous credentials [15] are well-known examples of this.

Broadcast Encryption. Broadcast Encryption (BE) addresses the issue of confidentially broadcasting a message to an arbitrary subset drawn from a universe of users. We will call the universe of n users U and the target, or privileged, set S, where $S \subseteq U$. Since its introduction in 1993 by Fiat and Naor [22], various flavours of BE have been introduced: the scheme can be in a symmetric or asymmetric setting; the set of receivers could be static or dynamic; revocation and traitor-tracing algorithms could be integrated into the system, users' keys might or might not be updated and then forward secrecy may be achieved. We refer to some of the relevant work in the area and the references therein [22,32,19,39,9,18,17,24,36]. One of the fundamental properties of a BE scheme is *collusion resistance* in the sense that no coalition of users in $U \setminus S$ should be able to recover the message. In the literature we can find several schemes that resist collusion attacks mounted by coalitions of at most $t < n$ users; only some schemes are *fully* collusion-resistant, *i.e.* they can tolerate attacks by coalitions of any size. For the purpose of this paper, we will consider systems that are *public-key*, allow *stateless receivers* (users that are not required to update their private keys) and are *fully* collusion-resistant. These are by now standard objectives for a BE scheme in the public-key setting.

Several additional practical aspects need to be taken into consideration, especially in view of the real-life applications of BE: strength of security notions, public and private storage requirements, ciphertext length, and computational costs. The specific nature of the primitive has led researchers to focus in particular on solutions having ciphertexts that are as short as possible. In this respect, the results of [9] and [24] are nearly optimal. However, designing BE schemes for real-life applications to broadcasting should not only involve efficiency and confidentiality issues. In particular, the privacy of users should be protected as much as possible. We believe that, to date, this aspect has not been adequately dealt with. Our study of the literature reveals that anonymity in BE has only been considered in a single paper [5], in the context of encrypted file systems[3]. Surprisingly, almost all subsequent work on BE has ignored the issue of anonymity. Moreover, as we shall explain below, state-of-the-art BE schemes are inherently incapable of providing any kind of anonymity.

Anonymity in Broadcast Encryption. According to commonly accepted definitions [24,10,17], a BE scheme consists of four algorithms: `Setup`, `KeyGen`, `Enc` and `Dec`. Each user in the system can obtain his private key from the `KeyGen` algorithm, and the sender can choose an *arbitrary* target set of users S to which he wishes to broadcast a message. To decrypt, a legitimate user, i.e. a user in S, has to run the decryption algorithm on input the ciphertext, his private key *and* a description of the target set S. This set S is required specifically as an input to `Dec` in the existing definitions of BE. Hence the user needs to somehow know to which set S the message was broadcast, otherwise he cannot decrypt. Unfortunately, solving this problem is not just a matter of removing this requirement

[3] We observe that [25] addresses the issue of hiding the identity of the *sender* in a broadcast protocol, which is *not* what we intend by anonymous broadcast encryption.

from the model, as current schemes explicitly *rely* on S as an input to `Dec` for decryption to work. Thus these schemes cannot provide any anonymity.

This limitation in the existing BE model and schemes clearly causes serious privacy issues: imagine we deploy a BE scheme, as defined above, for television broadcasting. Suppose the privileged set is the set of all users who have paid a subscription to a certain channel. Each customer should have access to that channel using his private key. The problem is that, to decrypt, he will have to know who *else* has paid for the specific subscription! Not only is this requirement very inconvenient for the practical deployment of BE schemes, it is also a severe violation of the individual subscriber's privacy. Ideally, a BE scheme should protect users' privacy by guaranteeing that ciphertexts do not leak any information about the privileged set S.

Current BE schemes such as those in [24,10,17] do not account for the cost of broadcasting a description of S when calculating the size of ciphertexts. In the most general usage scenario intended for BE, where S is dynamic and may be unpredictable from message to message, the ciphertexts in such schemes must effectively include a description of S as part of the ciphertexts themselves. This means that the true ciphertext size in these schemes is linear in n rather than constant-size, as a cursory examination of the schemes might suggest[4]. However, achieving linear-sized ciphertext is already an impressive achievement, since there is a simple counting argument showing that, for a universe of n users in which every possible subset S should be reachable by secure broadcast, ciphertexts must contain at least n bits.

Further Details on Related Work. As mentioned above, the only prior work addressing the issue of anonymity in BE appears to be that of Barth *et al.* [5] (there, it is called *privacy*). In [5], several BE systems used in practice were examined with respect to anonymity. In addition, a generic construction for a BE scheme using a key-private, IND-CCA secure PKE scheme was given, with the scheme achieving anonymity and IND-CCA security against static adversaries. The construction encrypts the message for each intended receiver using the PKE scheme, and then ties together the resulting ciphertexts using a strongly secure one-time signature. Barth *et al.* [5] also provided a technique which can be used to speed-up decryption, but this technique was only analysed in the Random Oracle Model.

In very recent work [21] that builds on [5] and this paper, the authors have given constructions for anonymous broadcast encryption schemes with compact ciphertexts, but using a much weaker notion of anonymity that does not seem to relate very closely to real-world requirements.

In [11] the authors provide a private linear broadcast encryption (PLBE) scheme to realise a fully collusion-resistant traitor-tracing scheme. A PLBE, however, is a BE system with limited capabilities (i.e. it cannot address arbitrary sets of users) and hence this work does not provide a solution to the

[4] This does not rule the use of compact encodings of S being transmitted with ciphertexts in more restrictive usage scenarios, for example, only sending the difference in S when the set S changes only slowly from message to message.

problem considered so far.

There is much work, both cryptographic and non-cryptographic, on pseudonymous systems. In principle, pseudonyms could be used to enhance the anonymity of BE schemes: now users would not be identifiable directly, since a certificate would link a public key to a pseudonym rather than a real name. However, ciphertexts would still be linkable, in the sense that it would be possible to detect if two ciphertexts were intended for the same set of recipients or not. The approach we take here offers much stronger levels of privacy, removing ciphertext linkability in particular.

Our Contributions. Despite its importance, anonymous broadcast encryption has not received much attention since the initial work of Barth *et al.* [5]. This paper aims to raise the profile of this neglected primitive.

We start by giving a unified security definition for Anonymous Broadcast Encryption (ANOBE). Instead of separating anonymity and confidentiality as in [5], we use a combined security notion for ANOBE which helps to streamline our presentation and proofs. In addition, we strengthen the model to allow the adversary to make *adaptive* corruptions, *with all of our constructions achieving security in this setting.* In contrast, the definition of [5] is static, requiring the adversary to choose whom to corrupt before seeing the public keys in the system. As a first step, we show that our enhanced security definition is satisfiable: adaptively secure ANOBE can be built based only on the existence of IND-CCA secure PKE (without requiring the base PKE scheme to have anonymity properties itself). This construction results in a very efficient (constant-time) decryption procedure but has ciphertexts whose size is linear in n, the number of users in the universe U.

Our second contribution is to show that the generic construction for ANOBE suggested by Barth *et al.* [5] actually possesses adaptive security, and not merely static security as was established in [5]. This construction starts from any weakly robust (in the sense of [2]), key-private PKE scheme with chosen-ciphertext security. In comparison with our first generic construction, this result imposes stronger requirements on the underlying encryption scheme. However, it achieves shorter ciphertexts, with the size being linear in the size of the target set S. We also provide a variant of this construction that replaces the IND-CCA secure PKE component with an identity-based encryption (IBE) scheme having suitable security properties. This alternative further increases the set of components that can be used to obtain ANOBE.

One major drawback of the latter constructions is that decryption takes linear time in the size of the set S. Our third result is a technique allowing for constant decryption cost and which we prove secure in the standard model (*i.e.*, without random oracles) using our enhanced security definition. So far, the only known technique – put forth by Barth *et al.* [5] – enabling constant-time decryption requires the random oracle heuristic in the security analysis. To eliminate the random oracle, we introduce a new primitive, which we call an *anonymous hint system.* In essence, this primitive provides a way for an encrypter to securely tell receivers which ciphertext component is intended for them, allowing them

to ignore all but one ciphertext component and so decrypt more efficiently. The hint primitive, for which we provide an implementation based on the Decision-Diffie-Hellman (DDH) assumption, is defined and realized in such a way that its integration with our generic ANOBE constructions maintains compatibility with our proofs of adaptive security.

Our fourth contribution is to show how randomness re-use techniques originally developed for PKE in [28,8,7] can be modified for secure deployment in the ANOBE setting. In particular, we identify a slightly stronger notion of reproducibility that we call *key-less reproducibility*. We show that if our base PKE scheme has this property (in addition to the other properties needed in our generic construction) then it can be used with the same randomness across all ciphertext components in our main ANOBE construction. This not only allows the size of ciphertexts to be reduced further (by eliminating repeated ciphertext elements) but also reduces the sender's computational overhead.

In the full version of the paper [30], we establish that the Kurosawa-Desmedt (KD) [29] hybrid encryption scheme can be tweaked to have all the properties that are needed of the base PKE scheme in our constructions. The KD scheme is an ideal starting point since it is one of most efficient PKE schemes with IND-CCA security in the standard model.

Tying everything together and using KD* as the base scheme, we obtain the *currently most efficient instantiation* of an ANOBE scheme, for which ciphertexts contain only 2 group elements and $|S|$ symmetric ciphertexts (plus a signature and a verification key). Decryption can be achieved in constant time by combining this scheme with our DDH-based hint system, with an additional $2|S| + 1$ group elements in the ciphertext.

As can be seen from the details of our constructions, achieving anonymity does not add *any* cost to the encryption process compared to non-anonymous schemes (for example, [9,24]): in our ANOBE schemes, encryption requires a number of group operations that is linear in $|S|$. As for decryption, our speed-up technique allows the legitimate user to recover the message in constant time. Our ciphertext size is linear in $|S|$ (and thus linear in n and of the same order of magnitude as the *true* ciphertext size in existing BE schemes). Thus one interpretation of our results is that anonymity does not "cost" anything in an asymptotic sense. Naturally, the constants matter in practice, and reducing the constant in the ciphertext size for ANOBE to something closer to what can be achieved in the non-anonymous setting is a major open problem. However, we reiterate that reducing the *true* size of ciphertexts below linear in n in either the anonymous or non-anonymous setting is impossible.

2 Anonymous Broadcast Encryption

We define a model of public-key Broadcast Encryption, where algorithms are specified to allow for anonymity (similarly to [5]) and they are general enough to include the identity-based variant of BE introduced in [17].

Definition 1. *Let $U = \{1, ..., n\}$ be the universe of users. A* broadcast encryption (BE) scheme *is defined by four algorithms and has associated message space $\mathcal{MSP}$ and ciphertext space $\mathcal{CSP}$.*

`BE.Setup`(λ, n)**:** *This algorithm takes as input the security parameter λ and the number of users in the system n. It outputs a master public key BE-MPK and a master secret key BE-MSK.*

`BE.Key-Gen`$(\text{BE-MPK}, \text{BE-MSK}, i)$**:** *This algorithm takes as input BE-MPK, BE-MSK and an index $i \in U$ and outputs the private key sk_i for user i.*

`BE.Enc`$(\text{BE-MPK}, m, S)$**:** *This algorithm takes as input BE-MPK, a message $m \in \mathcal{MSP}$ and a subset $S \subseteq U$, the broadcast target set. It outputs a ciphertext $c \in \mathcal{CSP}$.*

`BE.Dec`$(\text{BE-MPK}, sk_i, c)$**:** *This algorithm takes as input BE-MPK, a private key sk_i and a ciphertext $c \in \mathcal{CSP}$. It outputs either a message $m \in \mathcal{MSP}$ or a failure symbol $\perp$.*

For all $S \subseteq U$ and $i \in U$, if $c =$ `BE.Enc`$(\text{BE-MPK}, m, S)$ and sk_i is the private key for $i \in S$, then `BE.Dec`$(\text{BE-MPK}, sk_i, c) = m$ with overwhelming probability.

We observe that this definition no longer requires the set S as an input to the decryption algorithm. This is crucial in developing the notion of anonymous broadcast encryption (ANOBE), for which we next provide an appropriate security model for the case of *adaptive* adversaries.

Definition 2. *We define the ANO-IND-CCA security game for BE as follows.*
Setup. *The challenger $\mathcal{C}$ runs `BE.Setup`(λ, n) to generate the master key pair (BE-MPK, BE-MSK) and gives BE-MPK to the adversary $\mathcal{A}$.*
Phase 1. *$\mathcal{A}$ can issue queries to a private key extraction oracle for any index $i \in U$. The oracle will respond by returning $sk_i =$ `BE.Key-Gen`$(\text{BE-MPK}, \text{BE-MSK}, i)$. $\mathcal{A}$ can also issue decryption queries of the form (c, i), where $i \in U$, and the oracle will return the decryption `BE.Dec`$(\text{BE-MPK}, sk_i, c)$.*
Challenge. *$\mathcal{A}$ selects two equal-length messages m_0, $m_1 \in \mathcal{MSP}$ and two distinct sets $S_0, S_1 \subseteq U$ of users. We require that S_0 and S_1 be of equal size and also impose the restriction that $\mathcal{A}$ has not issued key queries for any $i \in S_0 \triangle S_1 = (S_0 \setminus S_1) \cup (S_1 \setminus S_0)$. Further, if there exists an $i \in S_0 \cap S_1$ for which $\mathcal{A}$ has queried the key, then we require that $m_0 = m_1$. The adversary $\mathcal{A}$ passes m_0, m_1 and S_0, S_1 to $\mathcal{C}$. The latter picks a random bit $b \in \{0, 1\}$ and computes $c^* =$ `BE.Enc`$(\text{BE-MPK}, m_b, S_b)$ which is returned to $\mathcal{A}$.*
Phase 2. *$\mathcal{A}$ continues to make queries to the private key extraction oracle with the restrictions that $i \notin S_0 \triangle S_1$ and that, if $i \in S_0 \cap S_1$, then $m_0 = m_1$. $\mathcal{A}$ may continue issuing decryption queries (c, i) with the restriction that if $c = c^*$ then either $i \notin S_0 \triangle S_1$ or $i \in S_0 \cap S_1$ and $m_0 = m_1$.*
Guess. *The adversary outputs its guess b' for b.*

Definition 3. *We say that a BE scheme is* anonymous and semantically secure against chosen-ciphertext attacks *(ANO-IND-CCA) if all polynomial-time adaptive adversaries $\mathcal{A}$ have at most negligible advantage in the above game, where $\mathcal{A}$'s advantage is defined as $Adv_{\mathcal{A},BE}^{ANO\text{-}IND\text{-}CCA}(\lambda) = \left|\Pr[b' = b] - \frac{1}{2}\right|$.*

Like the definition of [5], Definition 2 does not require the ANOBE ciphertext to hide the number of receivers. However, specific schemes (such as the one in Section 3.1) can also conceal the cardinality of S.

We will next show that this notion is indeed *feasible* by presenting a generic construction that relies solely on the existence of IND-CCA secure PKE schemes. We will then improve its performance by giving alternative generic constructions whose underlying primitives require additional security properties.

3 Generic Constructions for ANOBE from PKE

3.1 ANOBE from Minimal Assumptions

Since our aim is to provide a formal treatment of anonymous broadcast encryption, we begin by showing that ANOBE *can be achieved*. Indeed, by simply assuming the existence of an IND-CCA secure PKE scheme we can construct an ANOBE scheme as follows.

Let $\pi^{\text{pke}} = (\mathsf{Gen}, \mathsf{KeyGen}, \mathsf{Encrypt}, \mathsf{Decrypt})$ be a PKE scheme with message space $\mathcal{M} = \{0,1\}^m$. Here, algorithm Gen takes as input a security parameter and outputs public parameters par, used by KeyGen to generate a key pair (pk, sk). Let $\Sigma = (\mathcal{G}, \mathcal{S}, \mathcal{V})$ be a one-time signature scheme consisting of a key generation algorithm $\mathcal{G}$, a signing algorithm $\mathcal{S}$ and a verification algorithm $\mathcal{V}$. We assume that the key space of Σ is $\mathcal{K} = \{0,1\}^v$, for some $v \in \mathsf{poly}(\lambda)$. We use π^{pke} and Σ to generically instantiate a BE scheme, with message space $\{0,1\}^{m-v}$. In the description hereafter, we include the symbol ε as a valid but distinguished message in $\{0,1\}^{m-v}$: in other words, all the messages that receivers accept as legal plaintexts are different from ε.

`BE.Setup`(λ, n)**:** Generate $par \leftarrow \mathsf{Gen}(\lambda)$ and, for $i = 1$ to n, generate $(sk_i, pk_i) \leftarrow \mathsf{Keygen}(par)$. The master private key is BE-MSK $= \{sk_i\}_{i=1}^n$ and the master public key consists of BE-MPK $= \big(par,\ \Sigma,\ \{pk_i\}_{i=1}^n\big)$.

`BE.Key-Gen`(BE-MPK, BE-MSK, i)**:** parse the master secret key BE-MSK as $\{sk_i\}_{i=1}^n$ and output sk_i.

`BE.Enc`(BE-MPK, M, S)**:** to encrypt M for a receiver set $S \subseteq \{1, \ldots, n\}$, generate a one-time key pair $(\mathsf{SK}, \mathsf{VK}) \leftarrow \mathcal{G}(\lambda)$. For each $j = 1$ to n, compute $C_j = \mathsf{Encrypt}(par, pk_j, M||\mathsf{VK})$ if $j \in S$ and $C_j = \mathsf{Encrypt}(par, pk_j, \varepsilon||\mathsf{VK})$ if $j \notin S$. Finally, output $C = \big(C_1, \ldots, C_n, \sigma\big)$, where $\sigma = \mathcal{S}\big(\mathsf{SK}, (C_1, \ldots, C_n)\big)$.

`BE.Dec`(BE-MPK, sk_i, C)**:** given the ANOBE ciphertext $C = \big(C_1, \ldots, C_n, \sigma\big)$, compute $M' = \mathsf{Decrypt}(sk_i, C_i)$. If $M' \neq \perp$, parse M' as $M' = M||\mathsf{VK}$ for some bitstrings $M \in \{0,1\}^{m-v}$ and $\mathsf{VK} \in \{0,1\}^v$. Then, if it holds that $\mathcal{V}\big(\mathsf{VK}, (C_1, \ldots, C_n), \sigma\big) = 1$ and $M \neq \varepsilon$ return M. Otherwise, output $\perp$.

The correctness follows directly from the correctness of π^{pke} and Σ. This construction is reminiscent of generic constructions of chosen-ciphertext-secure multiple encryption [20] and it is easily seen to yield a secure ANOBE. A proof of the following theorem is available in the full version of the paper [30].

Theorem 1. *Let π^{pke} be an IND-CCA secure PKE scheme and let Σ be a strongly unforgeable one-time signature scheme. The BE scheme constructed above is ANO-IND-CCA secure against adaptive adversaries.*

We have described an ANOBE scheme from minimal assumptions. We note that encryption time is linear in n but decryption is performed in *constant* time, since a user simply selects the ciphertext component to decrypt according to its index. However, the ciphertext size is *linear* in n, as we encrypt to each user in the universe. It is desirable to improve on this and achieve a realization of ANOBE with more compact ciphertexts.

We will next see how to modify this first generic construction, obtaining an ANOBE scheme whose ciphertext size is linear in the size of the *target set* S.

3.2 Adaptively Secure ANOBE from Robust, Anonymous PKE

A simple solution to the broadcast problem is to encrypt the message under the public key of each user in the privileged set. This naive approach, so often discarded in most BE literature due to efficiency reasons, turns out to provide another generic construction for ANOBE, which differs from the previous one as now we deploy a public-key encryption scheme only to encrypt the *message* to the users *in the target set*.

For this approach, the underlying PKE scheme has to be key-private (or IK secure [6]), in that the ciphertext does not leak under which public key it was created. We also require the PKE scheme to be weakly robust, in the sense of [2], not only for correctness but also for consistency in the CCA security proof simulation. This property can be generically achieved [2] for any PKE scheme using a simple redundancy-based transformation.

This is essentially the construction that was already suggested by Barth, Boneh and Waters [5]. We now prove that it is actually *adaptively* secure, rather than just statically secure, as was established in [5].

Let $\pi^{\mathrm{pke}} = (\mathsf{Gen}, \mathsf{Keygen}, \mathsf{Encrypt}, \mathsf{Decrypt})$ be a PKE scheme and $\Sigma = (\mathcal{G}, \mathcal{S}, \mathcal{V})$ be a one-time signature. Our ANOBE scheme, $\mathrm{ANOBE}^{\pi^{\mathrm{pke}}, \Sigma}$, is as follows.

`BE.Setup`(λ, n)**:** Run $\mathsf{Gen}(\lambda, n)$ to obtain public parameters par. For $i = 1$ to n, run $\mathsf{Keygen}(par)$ to generate (sk_i, pk_i). The master private key is BE-MSK $= \{sk_i\}_{i=1}^n$ and the master public key is BE-MPK $= \big(par,\ \Sigma,\ \{pk_i\}_{i=1}^n\big)$.

`BE.Key-Gen`(BE-MPK, BE-MSK, i)**:** given BE-MSK $= \{sk_i\}_{i=1}^n$, output sk_i.

`BE.Enc`(BE-MPK, M, S)**:** to encrypt M for a receiver set $S = \{i_1, \ldots, i_\ell\} \subseteq \{1, \ldots, n\}$ of size $\ell = |S|$, generate a signature key pair $(\mathsf{SK}, \mathsf{VK}) \leftarrow \mathcal{G}(\lambda)$. For $j = 1$ to ℓ, compute $C_j = \mathsf{Encrypt}(par, pk_{i_j}, M||\mathsf{VK})$. The ANOBE ciphertext is $C = \big(\mathsf{VK}, C_{\tau(1)}, \ldots, C_{\tau(\ell)}, \sigma\big)$, where $\sigma = \mathcal{S}\big(\mathsf{SK}, C_{\tau(1)}, \ldots, C_{\tau(\ell)}\big)$ and $\tau : \{1, \ldots, \ell\} \to \{1, \ldots, \ell\}$ is a random permutation.

`BE.Dec`(BE-MPK, sk_i, C)**:** parse C as a tuple $\big(\mathsf{VK}, C_1, \ldots, C_\ell, \sigma\big)$. Return $\perp$ if $\mathcal{V}\big(\mathsf{VK}, C_1, \ldots, C_\ell, \sigma\big) = 0$. Otherwise, repeat these steps for $j = 1$ to ℓ.

1. Compute $M' = \mathsf{Decrypt}(sk_i, C_j)$. If $M' \neq \perp$ and can moreover be parsed as $M' = M||\mathsf{VK}$ for some M of appropriate length, return M.

2. If $j = \ell$ output $\perp$.

The correctness of ANOBE$^{\pi^{\text{pke}},\Sigma}$ follows directly from the correctness and weak robustness of π^{pke}.

Theorem 2. *ANOBE$^{\pi^{\text{pke}},\Sigma}$ is adaptively ANO-IND-CCA secure assuming that: (i) π^{pke} is key-private and IND-CCA (AI-CCA) secure and weakly robust under chosen-ciphertext attacks (as defined in [2]); (ii) Σ is a strongly unforgeable one-time signature scheme.*

In our proof (given in the full version of the paper) we make use of a sequence of hybrid arguments where ciphertext components are gradually modified at each step and each hybrid argument requires the reduction to guess upfront the identity of an uncorrupted user.

In terms of efficiency, from this construction we will obtain secure ANOBE schemes with typically very small (constant) private key storage requirements and ciphertexts which are $|S|$ times the size of the ciphertext of the underlying PKE scheme. Encryption and decryption have both cost linear in the size of S.

If we look at recent efficient instantiations of BE, for example that of Gentry-Waters [24], we have private keys whose size is linear in the number of users, and ciphertexts which consist of n bits plus 3 group elements (if we include the cost of transmitting a description of S as part of the ciphertext). It is clear that in general the solution of [24] is more efficient in terms of ciphertext size. The key point though is that it is not anonymous.

4 Generic Construction for ANOBE from IBE

An IBE scheme I typically consists of four algorithms $(\mathsf{Setup}, \mathsf{KeyExt}, \mathsf{Enc}, \mathsf{Dec})$, where Setup and KeyExt are run by a trusted authority (TA). Our construction uses a multi-TA IBE scheme $I' = (\mathsf{CommonSetup}, \mathsf{TASetup}, \mathsf{KeyDer}, \mathsf{Enc}', \mathsf{Dec}')$ as formalized in [34]. We recall from [34] that $\mathsf{CommonSetup}$, on input the security parameter, outputs the system's parameters par and a set of labels of the TAs in the system, and that $\mathsf{TASetup}$, on input par, outputs a master public key mpk and a master secret key msk. This algorithm is randomized and executed independently for each TA in the system. The remaining algorithms are as per a normal IBE scheme. For this primitive we consider the notion of TA anonymity, as defined in [34], which formally models the inability of the adversary to distinguish two ciphertexts corresponding to the same message and identity, but created using different TA master public keys. An example of a TA-anonymous IBE scheme is the multi-TA version of Gentry's IBE [23] developed in [35].

Now, let $I' = (\mathsf{CommonSetup}, \mathsf{TASetup}, \mathsf{KeyDer}, \mathsf{Enc}', \mathsf{Dec}')$ be a weakly robust (in the sense of a definition of robustness deferred to the full version of the paper), multi-TA IBE scheme and let $\Sigma = (\mathcal{G}, \mathcal{S}, \mathcal{V})$ be a signature scheme. We will use I' and Σ to generically instantiate a BE scheme in the following way.

`BE.Setup`(λ, n)**:** Run $\mathsf{CommonSetup}$ on input of $\lambda \in \mathbb{N}$ to obtain the system's parameters par. Run $\mathsf{TASetup}(par)$ n times to obtain n distinct master key pairs $\{mpk_i, msk_i\}_{i \in U}$. Return the par, Σ and n public keys $\{mpk_i\}_{i \in U}$.

$\texttt{BE.Key-Gen}(par, \lambda, i)$**:** Return msk_i, the secret key corresponding to the public key mpk_i of user i.

$\texttt{BE.Enc}(par, M, S)$**:** Run $\mathcal{G}$ to obtain a one-time signature key pair $(\mathsf{SK}, \mathsf{VK})$. For each $i \in S$ run $\mathsf{Enc}'(mpk_i, M, \mathsf{VK})$ to obtain ciphertext C_i. The ANOBE ciphertext is $C = \big(\mathsf{VK}, C_{\tau(1)}, \ldots, C_{\tau(\ell)}, \sigma\big)$, where $\sigma = \mathcal{S}\big(\mathsf{SK}, C_{\tau(1)}, \ldots, C_{\tau(\ell)}\big)$ and $\tau : \{1, \ldots, \ell\} \to \{1, \ldots, \ell\}$ is a random permutation.

$\texttt{BE.Dec}(par, msk_i, C)$**:** Parse C as $\big(\mathsf{VK}, C_1, \ldots, C_\ell, \sigma\big)$. If $\mathcal{V}\big(\mathsf{VK}, C_1, \ldots, C_\ell, \sigma\big) = 0$, return $\perp$. Otherwise, compute $sk_{i_{\mathsf{VK}}} = \mathsf{KeyDer}(mpk_i, msk_i, \mathsf{VK})$ and repeat the following steps for $j = 1$ to ℓ.

1. Compute $M' = \mathsf{Dec}'(mpk_i, sk_{i_{\mathsf{VK}}}, C_j)$. If $M' \neq \perp$ and can moreover be parsed as $M' = M||\mathsf{VK}$ for some M of appropriate length, return M.
2. If $j = \ell$ output $\perp$.

The correctness of the BE scheme follows directly from the correctness and the weak robustness of the IBE scheme I' used to construct it.

If instantiated with the multi-TA version of Gentry's IBE scheme [23,35] (which can be made weakly robust simply by applying the transform in [2]), this construction yields very short constant size private keys (just one element in $\mathbb{Z}_p^*$) and ciphertexts consisting of roughly $3 \cdot |S|$ group elements ($|S|$ in $\mathbb{G}$ and $2 \cdot |S|$ in $\mathbb{G}_T$) plus a signature and a verification key. Encryption and decryption have both cost linear in the size of S.

Theorem 3. *Let I' be a TA-anonymous, sID-IND-CPA secure IBE scheme and let Σ be a strongly unforgeable one-time signature. Then, the above BE scheme is adaptively ANO-IND-CCA secure.*

We give some intuition for the proof. We observe that, in [35], the authors apply a modified version of the Canetti-Halevi-Katz (CHK) transform [13] using the same primitives as our generic construction to obtain a key-private IND-CCA PKE scheme. We introduce further modifications to build a BE scheme achieving ANO-IND-CCA security. The idea is that, within this transform, we encrypt m for the *same* identity VK under the $|S|$ different public keys. We then sign all ciphertexts and append the verification key VK (note that this signature binds all these ciphertexts together). Upon decryption, a user verifies the signature against VK and, if valid, proceeds to derive the decryption key for identity VK by running the IBE key-extraction algorithm on input his private key. By similar arguments to those in [13] and [35], and by applying techniques analogous to those proving adaptive security in Theorem 2, we can show that adaptive ANO-IND-CCA security is achieved.

5 Efficient Decryption in the Standard Model

The generic constructions for ANOBE presented in Section 3.2 and 4 both suffer from linear time decryption. This arises from the fact that users do not know which ciphertext component is intended for them, and hence will have to perform an average of $|S|/2$ decryptions before recovering the message. Clearly this

procedure is quite cumbersome. We now present a technique which achieves *constant* time decryption in the standard model. We make use of a new primitive, called tag-based *anonymous hint systems*, for which we provide a definition, the relevant security models and a concrete instantiation.

5.1 Tag-Based Anonymous Hint Systems

A tag-based anonymous *hint* system is a tag-based encryption scheme [27] allowing to generate weak forms of encryption under a tag t and a public key pk. The result of the process consists of a *value* U and a *hint* H. The pair (U, H) should be pseudo-random (in particular, hints generated under two distinct public keys should be indistinguishable) when only the public key pk is available. Also, the private key sk makes it possible to check whether a given hint H is valid w.r.t. a tag t. A value-hint pair can be seen as an extractable commitment to a public key. Formally, such a system is defined in terms of the following algorithms.

$\texttt{Keygen}(\texttt{cp})$: takes as input a set of common public parameters $\texttt{cp}$ and outputs a key pair (sk, pk). We assume that $\texttt{cp}$ specifies a randomness space $\mathcal{R}^h$ and a space $\mathcal{T}^h$ of acceptable tags for the scheme.

$\texttt{Hint}(\texttt{cp}, t, pk, r)$**:** is a deterministic algorithm taking as input common public parameters $\texttt{cp}$, a public key pk, a tag t and random coins $r \in_R \mathcal{R}^h$. It outputs pair (U, H) consisting of a value U and a hint H. It is required that U only depends on the random coins r and not on pk.

$\texttt{Invert}(\texttt{cp}, sk, t, U)$**:** is a deterministic "inversion" algorithm taking as input a value U, a tag t and a private key sk. It outputs either a hint H or $\perp$ if U is not in the appropriate domain.

Correctness requires that, for any pair $(sk, pk) \leftarrow \texttt{Keygen}(\lambda)$ and any possible random coins r, if $(U, H) \leftarrow \texttt{Hint}(t, pk, r)$, then $\texttt{Invert}(\texttt{cp}, sk, t, U) = H$.

Although hint systems bear similarities with tag-KEMs, as formalized by Abe *et al.* [3], the two primitives are different and incomparable. In the tag-KEM syntax, the symmetric "session key" is chosen first and it does not depend on the tag. In hint schemes, the syntax requires to choose a pair (U, H), where U does not depend on pk but the session key H can depend on both pk and the tag (this is what happens in the construction we give). The security definitions are also different since, in Definition 4 hereafter, there is no inversion oracle (that would return H given U and t) but only a verification oracle that determines if (U, H, t) form a valid triple with respect to public keys pk_0 and pk_1.

In certain aspects, hint schemes are reminiscent of extractable hash proof systems [38] but there are several differences. In [38], in addition to the value that we call U, the random coins allowing to compute U are used to compute a witness S such that (U, S) satisfies some relation. From U, the element S is also computable using the private key and the value that we call H (which is termed "hash value" in [38]). At the same time, S should be infeasible to compute without the private key or the random coins used to sample U. Hint schemes are different in that they rather require the hardness of computing H

from U without the private key. In addition, tag-based hints require that it be hard to decide if a pair (U, H) is valid for a certain tag $t^\star$ (*i.e.*, to decide if $H = \mathtt{Invert}(\mathsf{cp}, sk, t^\star, U)$) even with access to a decision oracle for tags $t \neq t^\star$.

Definition 4. *A tag-based hint system* $(\mathtt{Keygen}, \mathtt{Hint}, \mathtt{Invert})$ *is* anonymous *if no PPT adversary has non-negligible advantage in the following game:*

1. *On input of common public parameters* cp*, the adversary* $\mathcal{A}$ *chooses a tag* $t^\star$ *and sends it to the challenger.*
2. *The challenger generates two key pairs* $(sk_0, pk_0) \leftarrow \mathtt{Keygen}(\lambda)$*,* $(sk_1, pk_1) \leftarrow \mathtt{Keygen}(\lambda)$ *and gives* pk_0, pk_1 *to* $\mathcal{A}$*.*
3. *On polynomially-many occasions,* $\mathcal{A}$ *adaptively invokes a verification oracle on value-hint-tag triples* (U, H, t) *such that* $t \neq t^\star$*. The challenger replies by returning bits* $(d_0, d_1) \in \{0,1\}^2$ *where* $d_0 = 1$ *if and only if* $H = \mathtt{Invert}(\mathsf{cp}, sk_0, t, U)$ *and* $d_1 = 1$ *if and only if* $H = \mathtt{Invert}(\mathsf{cp}, sk_1, t, U)$*.*
4. *When* $\mathcal{A}$ *decides to enter the challenge phase, the challenger flips a binary coin* $b \xleftarrow{\$} \{0,1\}$ *and chooses other random coins* $r^\star \xleftarrow{\$} \mathcal{R}^h$*. It outputs* $(U^\star, H^\star) = \mathtt{Hint}(\mathsf{cp}, t^\star, pk_b, r^\star)$*.*
5. $\mathcal{A}$ *makes further queries but is not allowed to make queries involving the target tag* $t^\star$*.*
6. $\mathcal{A}$ *outputs a bit* $b' \in \{0,1\}$ *and wins if* $b' = b$*.*

As usual, $\mathcal{A}$*'s advantage is the distance* $\mathbf{Adv}^{\text{anon-hint}}(\mathcal{A}) = |\Pr[b' = b] - 1/2|$*.*

Definition 5. *A tag-based hint system* $(\mathtt{Keygen}, \mathtt{Hint}, \mathtt{Invert})$ *is strongly* robust *if no PPT adversary* $\mathcal{A}$ *has non-negligible advantage in the following game, where* $\mathcal{A}$*'s advantage is its probability of success.*

1. *The challenger chooses public parameters* cp *and generates pairs* $(sk_0, pk_0) \leftarrow \mathtt{Keygen}(\lambda)$*,* $(sk_1, pk_1) \leftarrow \mathtt{Keygen}(\lambda)$*. It gives* cp *and* pk_0, pk_1 *to* $\mathcal{A}$*.*
2. $\mathcal{A}$ *invokes a verification oracle on arbitrary value-hint-tag triples* (U, H, t)*. The challenger replies by returning bits* $(d_0, d_1) \in \{0,1\}^2$ *where* $d_0 = 1$ *if and only if* $H = \mathtt{Invert}(\mathsf{cp}, sk_0, t, U)$ *and* $d_1 = 1$ *if and only if* $H = \mathtt{Invert}(\mathsf{cp}, sk_1, t, U)$*.*
3. $\mathcal{A}$ *outputs a triple* $(U^\star, H^\star, t^\star)$ *and wins if* $H^\star = \mathtt{Invert}(\mathsf{cp}, sk_0, t^\star, U^\star) = 1$ *and* $H^\star = \mathtt{Invert}(\mathsf{cp}, sk_1, t^\star, U^\star) = 1$*.*

Analogously to the PKE case [2], weak robustness is defined for tag-based hint schemes by letting the adversary simply make a challenge request in step 3. The challenger then chooses a tag $t^\star$ as well as random coins $r^\star$, generates a value-hint pair $(U^\star, H^\star) = \mathtt{Hint}(\mathsf{cp}, t^\star, pk_0, r^\star)$ and $\mathcal{A}$ wins if $H^\star = \mathtt{Invert}(\mathsf{cp}, sk_1, t^\star, U^\star) = 1$. Weak robustness will be sufficient for our purposes but the scheme hereafter is also strongly robust assuming that the discrete logarithm assumption holds in $\mathbb{G}$.

To show that this newly defined primitive is indeed feasible, we give an example of an anonymous hint system based on the DDH assumption and the CCA-secure public key encryption scheme described in [14].

Let the common public parameters $\mathsf{cp} = \{\mathbb{G}, p, g\}$ consist of a group $\mathbb{G}$ of prime order $p > 2^\lambda$ with a generator $g \in_R \mathbb{G}$. We assume that tags are elements of $\mathcal{T}^h = \mathbb{Z}_p^*$ and that the randomness space is $\mathcal{R}^h = \mathbb{Z}_p^*$.

$\texttt{Keygen}(\mathsf{cp})$**:** chooses random $x_1, x_2, y_1, y_2 \xleftarrow{\$} \mathbb{Z}_p^*$ and computes $X_i = g^{x_i}$ and $Y_i = g^{y_i}$ for each $i \in \{1, 2\}$. The public key is $pk = \big(X_1, X_2, Y_1, Y_2\big)$ and the private key is $sk = (x_1, x_2, y_1, y_2)$.

$\texttt{Hint}(\mathsf{cp}, t, pk, r)$**:** given $pk = (\mathbb{G}, p, g, X_1, X_2, Y_1, Y_2)$, return $\perp$ if $r \notin \mathcal{R}^h = \mathbb{Z}_p^*$. Otherwise, compute (U, H) as

$$U = g^r, \qquad\qquad H = (V, W) = \big((X_1^t X_2)^r,\ (Y_1^t Y_2)^r\big).$$

$\texttt{Invert}(\mathsf{cp}, sk, t, U)$**:** return $\perp$ if $U \notin \mathbb{G}$. Otherwise, parse the private key sk as $(x_1, x_2, y_1, y_2) \in (\mathbb{Z}_p^*)^4$ and output $H = (V, W) = (U^{t \cdot x_1 + x_2}, U^{t \cdot y_1 + y_2})$

In the full version of the paper, we prove that the scheme provides anonymity in the sense of Definition 4 under the DDH assumption and strong robustness (in the sense of Definition 5) under the discrete logarithm assumption.

5.2 ANOBE with Efficient Decryption

Let $\pi^{\text{hint}} = (\texttt{Keygen}, \texttt{Hint}, \texttt{Invert})$ be an anonymous hint system with its set of common public parameters cp. Let $\pi^{\text{pke}} = (\mathsf{Gen}, \mathsf{Keygen}, \mathsf{Encrypt}, \mathsf{Decrypt})$ be a PKE scheme and $\Sigma = (\mathcal{G}, \mathcal{S}, \mathcal{V})$ be a signature scheme.

$\texttt{BE.Setup}(\lambda, n)$**:** Obtain $(par) \leftarrow \mathsf{Gen}(\lambda)$ and, for each $i \in \{1, \ldots, n\}$, and generate encryption key pairs $(\tilde{sk}_i, \tilde{pk}_i) \leftarrow \pi^{\text{pke}}.\mathsf{Keygen}(par)$ as well as hint key pairs $(sk_i^h, pk_i^h) \leftarrow \pi^{\text{hint}}.\texttt{Keygen}(\mathsf{cp})$. The master private key consists of BE-MSK $= \{\tilde{sk}_i, sk_i^h\}_{i=1}^n$ and the master public key is

$$\text{BE-MPK} = \Big(\mathsf{cp}, par,\ \{\big(\tilde{pk}_i, pk_i^h\big)\}_{i=1}^n,\ \Sigma\Big).$$

$\texttt{BE.Key-Gen}(\text{BE-MPK}, \text{BE-MSK}, i)$**:** parse BE-MSK as $\{\tilde{sk}_i, sk_i^h\}_{i=1}^n$ and output $sk_i = (\tilde{sk}_i, sk_i^h)$.

$\texttt{BE.Enc}(\text{BE-MPK}, M, S)$**:** given a receiver set $S = \{i_1, \ldots, i_\ell\} \subseteq \{1, \ldots, n\}$ of size $\ell = |S|$ and a message M, generate a one-time signature key pair $(\mathsf{SK}, \mathsf{VK}) \leftarrow \mathcal{G}(\lambda)$. Then, choose random coins $r \xleftarrow{\$} \mathcal{R}^h$ for the hint scheme and compute $(U, H_j) = \pi^{\text{hint}}.\texttt{Hint}(\mathsf{cp}, \mathsf{VK}, pk_{i_j}^h, r)$ for $j = 1$ to ℓ (recall that the first output U of $\texttt{Hint}$ does not depend on the public key). For $j = 1$ to ℓ, compute $C_j = \pi^{\text{pke}}.\mathsf{Encrypt}(par, \tilde{pk}_{i_j}, M||\mathsf{VK})$. Choose a random permutation $\tau : \{1, \ldots, \ell\} \to \{1, \ldots, \ell\}$ and set the final ciphertext as

$$C = \big(\mathsf{VK}, U, (H_{\tau(1)}, C_{\tau(1)}), \ldots, (H_{\tau(\ell)}, C_{\tau(\ell)}), \sigma\big),$$

where $\sigma = \mathcal{S}\big(\mathsf{SK}, U, (H_{\tau(1)}, C_{\tau(1)}), \ldots, (H_{\tau(\ell)}, C_{\tau(\ell)})\big)$.

$\texttt{BE.Dec}(\text{BE-MPK}, sk_i, C)$**:** on input of $C = \big(\mathsf{VK}, U, (H_1, C_1), \ldots, (H_\ell, C_\ell), \sigma\big)$ and $sk_i = (\tilde{sk}_i, sk_i^h)$, return $\perp$ if $\mathcal{V}\big(\mathsf{VK}, U, (H_1, C_1), \ldots, (H_\ell, C_\ell), \sigma\big) = 0$ or if U is not in the appropriate space defined by π^{hint}. Otherwise, compute $H = \pi^{\text{hint}}.\texttt{Invert}(\mathsf{cp}, sk_i^h, \mathsf{VK}, U)$. If $H \neq H_j$ for all $j \in \{1, \ldots, \ell\}$, return $\perp$. Otherwise, let j be the smallest index such that $H = H_j$ and compute $M' = \pi^{\text{pke}}.\mathsf{Decrypt}(\tilde{sk}_i, C_j)$. If M' can be parsed as $M' = M||\mathsf{VK}$ for some M of appropriate length, return M. Otherwise, output $\perp$.

The correctness of this scheme follows directly from the correctness and weak robustness of its component schemes π^{hint} and π^{pke}.

The following security result is proved in the full version of the paper.

Theorem 4. *The above construction is adaptively ANO-IND-CCA secure if (i) π^{hint} is anonymous; (ii) π^{pke} is AI-CCA secure and weakly robust under chosen-ciphertext attacks; (iii) Σ is a strongly unforgeable one-time signature.*

In [5] a technique to speed up decryption was presented. The scheme of [5] can be seen as using a hint scheme where tags are empty strings and pairs (U, H_j) consist of $U = g^r$ and $H_j = H(X_{i_j}^r)$, where H is a random oracle and $X_{i_j} \in \mathbb{G}$ is the public key of the hint scheme. In the present context, it is tempting to believe that simple hints of the form $X_{i_j}^r$ suffice to achieve efficient decryption in the standard model. Indeed, one step of the proof consists of a DDH-based transition from one hybrid game to another and, during that specific transition, the simulator $\mathcal{B}$ could simply handle all decryption queries using the private keys $\{\tilde{sk}_i\}_{i=1}^n$ in the underlying encryption scheme since it knows them all. For reasons that will become apparent in the proof of a key lemma for Theorem 4 below, this does not suffice. The reason is that, the adversary can issue decryption queries where $(g, U = g^r, X_{i_j}, H_{i_j} = X_{i_j}^{r'})$ does *not* form a Diffie-Hellman tuple. In this case, the answer of the simulator would differ from that of the real decryption procedure in the chosen-ciphertext scenario: more precisely, the simulation could accept a ciphertext that would be rejected by a real decryption.

In [5], Barth, Boneh and Waters addressed this problem using a random oracle and the Gap Diffie-Hellman assumption [33]: each hint was of the form $H_j = H(X_{i_j}^r)$, where H is the random oracle. By invoking the DDH-oracle at each random oracle query, the simulator was able to figure out which ciphertext components had to be decrypted so as to perfectly emulate the real decryption algorithm. Here, we address this issue in the standard model using the tag-based anonymous hint primitive.

It is convenient to instantiate the above construction by combining our DDH-based hint scheme with an encryption scheme based on the same assumption such as the Cramer-Shoup cryptosystem. Interestingly both schemes can be instantiated using the same DDH-hard cyclic group. Considering efficiency, it is moreover possible to recycle the group element g^r of the hint system and simultaneously use it as part of a Cramer-Shoup ciphertext. In the security proof, everything goes through with these optimizations.

6 Shortening Ciphertexts with Randomness Re-Use

This section considers *randomness re-use* [7,4], which is a powerful tool providing computational and bandwidth savings, as a technique to optimize ANOBE schemes. In [7], Bellare *et al.* introduce a property, called *reproducibility*, providing a condition under which randomness re-use is secure. We define the notion of *key-less reproducibility*, which is better suited for the anonymity setting.

Definition 6. *Let* $\pi^{\text{pke}} = (\mathsf{Gen}, \mathsf{Keygen}, \mathsf{Encrypt}, \mathsf{Decrypt})$ *be a PKE scheme. Let* $\mathcal{M}$ *and* $\mathcal{R}$ *be the message and randomness space of* π^{pke}*. Let* R *be an algorithm that takes as input the public parameters, a ciphertext, another random message and a key pair* (sk, pk)*, and outputs a ciphertext. Consider the experiment:*

$\mathbf{Exp}^{\mathsf{KLR}}_{\pi^{\text{pke}},R}(\lambda)$
$(par) \xleftarrow{\$} \mathsf{Gen}(\lambda)$
$(pk, sk) \xleftarrow{\$} \mathsf{Keygen}(par)$
$m \xleftarrow{\$} \mathcal{M}$; $r \xleftarrow{\$} \mathcal{R}$
$c = \mathsf{Encrypt}(pk, m; r)$
$(pk', sk') \xleftarrow{\$} \mathsf{Keygen}(par)$
$m' \xleftarrow{\$} \mathcal{M}$
return 1 *if* $\mathsf{Encrypt}(par, pk', m'; r) = R(par, c, m', pk', sk')$ *and* 0 *otherwise.*

π^{pke} *is* key-less reproducible *if, for any* λ*, there is a PPT algorithm* R *such that the above experiment outputs 1 with probability 1.*

We note that this definition differs from the one in [7] since the algorithm R does not take pk (the public key under which c was created) as an input. Indeed, this is a crucial difference which allows extending the notion of reproducibility to the context where anonymity is required. We now reconsider the generic construction for ANOBE presented in Section 3.2.

Let $\pi^{\text{pke}} = (\mathsf{Gen}, \mathsf{Keygen}, \mathsf{Encrypt}, \mathsf{Decrypt})$ be a key-less reproducible PKE scheme and let $\Sigma = (\mathcal{G}, \mathcal{S}, \mathcal{V})$ be a one-time signature. We call $\text{ANOBE}_{rr}^{\pi^{\text{pke}},\Sigma}$ the scheme constructed from Σ and π^{pke} as follows.

`BE.Setup, BE.Key-Gen, BE.Dec` are as in Section 3.2.

`BE.Enc`(BE-MPK, M, S)**:** to encrypt M for a receiver set $S = \{i_1, \ldots, i_\ell\} \subseteq \{1, \ldots, n\}$ of size $\ell = |S|$, generate a signature key pair $(\mathsf{SK}, \mathsf{VK}) \leftarrow \mathcal{G}(\lambda)$. Choose $r \xleftarrow{\$} \mathcal{R}$, where $\mathcal{R}$ is the randomness space of π^{pke}_{par}. Then, for each $j = 1$ to ℓ, compute $C_j = \mathsf{Encrypt}(par, pk_{i_j}, M||\mathsf{VK}; r)$. The final BE ciphertext consists of $C = \big(\mathsf{VK}, C_{\tau(1)}, \ldots, C_{\tau(\ell)}, \sigma\big)$, where $\sigma = \mathcal{S}\big(\mathsf{SK}, C_{\tau(1)}, \ldots, C_{\tau(\ell)}\big)$ and $\tau : \{1, \ldots, \ell\} \to \{1, \ldots, \ell\}$ is a random permutation.

Theorem 5. *Let* $\pi^{\text{pke}} = (\mathsf{Gen}, \mathsf{Keygen}, \mathsf{Encrypt}, \mathsf{Decrypt})$ *be an AI-CCA secure, weakly robust and key-less reproducible PKE scheme. Let* Σ *be a strongly unforgeable one-time signature scheme. Then,* $\text{ANOBE}_{rr}^{\pi^{\text{pke}},\Sigma}$ *is adaptively ANO-IND-CCA secure.*

The proof for Theorem 5 (which is given in the full paper) is analogous to that of Theorem 2, the only difference being the use of algorithm R in the simulation.

We have shown that the key-less reproducibility of a PKE scheme guarantees that randomness can be re-used securely. We can exploit this property to compress the ANOBE ciphertexts and, depending on the concrete instantiation, significantly increase the efficiency of the scheme. More precisely, given

an $\text{ANOBE}_{rr}^{\pi^{\text{pke}},\Sigma}$ ciphertext $C = (\mathsf{VK}, C_{\tau(1)}, \ldots, C_{\tau(\ell)}, \sigma)$, let ccc denote the *common ciphertext components* that may arise in $C_{\tau(1)}, \ldots, C_{\tau(\ell)}$ from sharing randomness across PKE components, *i.e.*,

$$C_{\tau(1)} = (\mathsf{ccc}, \tilde{c}_{\tau(1)}), \ldots, C_{\tau(\ell)} = (\mathsf{ccc}, \tilde{c}_{\tau(\ell)}).$$

The compressed ANOBE ciphertext will be $\tilde{C} = (\mathsf{VK}, \mathsf{ccc}, \tilde{c}_{\tau(1)}, \ldots, \tilde{c}_{\tau(\ell)}, \sigma)$. Upon receipt, the user simply reconstitutes the original ciphertext C and runs `BE.Dec` as usual. We explore instantiations of this idea in the full version.

7 Conclusions and Open Problems

In the context of broadcast encryption the main focus of research to date has been on reducing ciphertext size. Achieving this has entailed sacrificing *all* anonymity properties. Yet we have argued that anonymity is a *fundamental property* to strive for in broadcast encryption. With the aim of highlighting the importance of this overlooked feature, we have formally defined the notion of anonymous broadcast encryption (ANOBE) and given several constructions for this primitive. We have also shown how these constructions can be improved via anonymous hint systems (to optimize decryption performance) and randomness re-use (to reduce the ciphertext size and the computational costs of encryption).

Much work still needs to be done in this area, from improving the efficiency of ANOBE schemes to considering all the additional properties that can be found in standard BE, such as traitor tracing, revocation, dynamism of users joining the system, and realising them in the anonymous setting. There is still a gap between the sizes of ciphertexts in state-of-the-art BE schemes and our ANOBE schemes. This gap is hidden in the constants in an asymptotic evaluation of ciphertext size (when the true size of ciphertexts is measured) but is nevertheless significant in practice. A major challenge, then, is to further reduce the size of ciphertexts in ANOBE, whilst maintaining its full anonymity properties.

Acknowledgements

The work in this paper was supported in part by the European Commission through the ICT programme under contract ICT-2007-216676 ECRYPT II. The work in this paper was sponsored in part by the US Army Research Laboratory and the UK Ministry of Defense 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 US Army Research Laboratory, the US Government, the UK Ministry of Defense, or the UK Government. The US and UK Governments are authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation hereon. The first author acknowledges the Belgian Fund for Scientific Research (F.R.S.- F.N.R.S.) for his "Collaborateur scientifique" fellowship. The second author was supported by an EPSRC Leadership Fellowship, EP/H005455/1.

References

1. 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. *J. Cryptology*, 21(3):350–391, 2008.
2. M. Abdalla, M. Bellare, and G. Neven. Robust encryption. In *TCC'10*, *LNCS 5978*, pp. 480–497. Springer, 2010.
3. M. Abe, R. Gennaro, K. Kurosawa, and V. Shoup. Tag-KEM/DEM: A new framework for hybrid encryption and a new analysis of kurosawa-desmedt kem. In *Eurocrypt 2005*, *LNCS* 3494, pp. 128–146. Springer, 2005.
4. M. Barbosa and P. Farshim. Randomness reuse: Extensions and improvements. In *IMA Int. Conf.*, *LNCS* 4887, pp. 257–276. Springer, 2007.
5. A. Barth, D. Boneh, and B. Waters. Privacy in encrypted content distribution using private broadcast encryption. In *Financial Cryptography 2006*, *LNCS* 4107, pp. 52–64. Springer, 2006.
6. M. Bellare, A. Boldyreva, A. Desai, and D. Pointcheval. Key-privacy in public-key encryption. In *Asiacrypt'01*, *LNCS* 2248, pp. 566–582. Springer, 2001.
7. M. Bellare, A. Boldyreva, K. Kurosawa, and J. Staddon. Multirecipient encryption schemes: How to save on bandwidth and computation without sacrificing security. *IEEE Trans. on Information Theory*, 53(11):3927–3943, 2007.
8. M. Bellare, A. Boldyreva, and J. Staddon. Randomness re-use in multi-recipient encryption schemes. In *PKC'03*, *LNCS* 2567, pp. 85–99. Springer, 2003.
9. D. Boneh, C. Gentry, and B. Waters. Collusion resistant broadcast encryption with short ciphertexts and private keys. In *Crypto'05*, *LNCS* 3621, pp. 258–275. Springer, 2005.
10. D. Boneh and M. Hamburg. Generalized identity based and broadcast encryption schemes. In *Asiacrypt'08*, *LNCS* 5350, pp. 455–470. Springer, 2008.
11. D. Boneh, A. Sahai, and B. Waters. Fully collusion resistant traitor tracing with short ciphertexts and private keys. In *Eurocrypt'06*, *LNCS* 4004, pp. 573–592. Springer, 20006.
12. X. Boyen and B. Waters. Anonymous hierarchical identity-based encryption (without random oracles). In *Crypto'06*, *LNCS* 4117, pp. 290–307. Springer, 2006.
13. R. Canetti, S. Halevi, and J. Katz. Chosen-ciphertext security from identity-based encryption. In *Eurocrypt'04*, *LNCS* 3027, pp. 207–222. Springer, 2004.
14. D. Cash, E. Kiltz, and V. Shoup. The twin Diffie-Hellman problem and applications. In *Eurocrypt'08*, *LNCS* 4965, pp. 127–145. Springer, 2008.
15. D. Chaum. Security without identification: Transaction systems to make Big Brother obsolete. *Commun. ACM 1985*, 28(10):1030–1044, 1985.
16. D. Chaum and E. van Heyst. Group signatures. In *Eurocrypt'91*, *LNCS* 547, pp. 257–265. Springer, 1991.
17. C. Delerablée. Identity-based broadcast encryption with constant size ciphertexts and private keys. In *Asiacrypt'07*, *LNCS* 4833, pp. 200–215. Springer, 2007.
18. C. Delerablée, P. Paillier, and D. Pointcheval. Fully collusion secure dynamic broadcast encryption with constant-size ciphertexts or decryption keys. In *Pairing'07*, *LNCS* 4575, pp. 39–59. Springer, 2007.
19. Y. Dodis and N. Fazio. Public key broadcast encryption for stateless receivers. In *Digital Rights Management Workshop 2002 (DRM'02)*, *LNCS* 2696, pp. 61–80. Springer, 2002.

20. Y. Dodis and J. Katz. Chosen-ciphertext security of multiple encryption. In *TCC'05*, *LNCS* 3378, pp. 188–209. Springer, 2005.
21. N. Fazio, I. Perera. Outsider-Anonymous Broadcast Encryption with Sublinear Ciphertexts. In *Public Key Cryptography 2012 (PKC'12)*, *LNCS* series. Springer, 2012.
22. A. Fiat and M. Naor. Broadcast encryption. In *Crypto'93*, *LNCS* 773, pp. 480–491. Springer, 1993.
23. C. Gentry. Practical identity-based encryption without random oracles. In *Eurocrypt'06*, *LNCS* 4004, pp. 445–464. Springer, 20006.
24. C. Gentry and B. Waters. Adaptive security in broadcast encryption systems (with short ciphertexts). In *Eurocrypt'09*, *LNCS* 5479, pp. 171–188. Springer, 2009.
25. J. Groth. Efficient maximal privacy in boardroom voting and anonymous broadcast. In *Financial Cryptography 2004*, *LNCS* 3110, pp. 90–104. Springer, 2004.
26. J. Katz, A. Sahai, and B. Waters. Predicate encryption supporting disjunctions, polynomial equations, and inner products. In *Eurocrypt'08*, *LNCS* 4965, pp. 146–162. Springer, 2008.
27. E. Kiltz. Chosen-ciphertext security from tag-based encryption. In *Theory of Cryptograhy Conference 2006*, *LNCS* 3876, pp. 581–600. Springer, 2006.
28. K. Kurosawa. Multi-recipient public-key encryption with shortened ciphertext. In *PKC'02*, *LNCS* 2274, pp. 48–63. Springer, 2002.
29. K. Kurosawa and Y. Desmedt. A new paradigm of hybrid encryption scheme. In *Crypto'04*, *LNCS* 3152, pp. 426–442. Springer, 2004.
30. B. Libert, K. G. Paterson and E. A. Quaglia. Anonymous Broadcast Encryption: Adaptive Security and Efficient Constructions in the Standard Model. Cryptology ePrint Archive: Report 2011/476.
31. P. Mohassel. A closer look at anonymity and robustness in encryption schemes. In *Asiacrypt'10*, *LNCS* 6477, pp. 501–518. Springer, 2010.
32. D. Naor, M. Naor, and J. Lotspiech. Revocation and tracing schemes for stateless receivers. In *Crypto'01*, *LNCS* 2139, pp. 41–62, 2001.
33. T. Okamoto and D. Pointcheval. The gap-problems: A new class of problems for the security of cryptographic schemes. In *PKC'01*, *LNCS* 1992, pp. 104–118. Springer, 2001.
34. K. G. Paterson and S. Srinivasan. Security and anonymity of identity-based encryption with multiple trusted authorities. In *Pairing'08*, *LNCS* 5209, pp. 354–375. Springer, 2008.
35. K. G. Paterson and S. Srinivasan. Building key-private public-key encryption schemes. In *ACISP'09*, *LNCS* 5594, pp. 276–292. Springer, 2009.
36. D.-H. Phan, D. Pointcheval, M. Strefler. Security Notions for Broadcast Encryption. In *ACNS'11*, *LNCS* 6715, pp. 377–394. Springer, 2011.
37. V. Shoup. A proposal for an iso standard for public key encryption (version 2.1). Manuscript, 2001.
38. H. Wee. Efficient chosen-ciphertext security via extractable hash proofs. In *Crypto'10*, *LNCS* 6223, pp. 314–332. Springer, 2010.
39. D. Yao, N. Fazio, Y. Dodis, and A. Lysyanskaya. ID-based encryption for complex hierarchies with applications to forward security and broadcast encryption. In *ACM-CCS'04*, pp. 354–363. ACM, 2004.