Journal of Computational and Applied Mathematics 259 (2014) 586–591

Contents lists available at ScienceDirect

Journal of Computational and Applied Mathematics

journal homepage: www.elsevier.com/locate/cam

Anonymous trace and revoke[☆]

Murat Ak[a,*], Serdar Pehlivanoğlu[b], Ali Aydın Selçuk[c]

[a] Department of Computer Engineering, Akdeniz University, Antalya, Turkey

[b] Department of Computer Engineering, Zirve University, Gaziantep, Turkey

[c] Department of Computer Engineering, TOBB University of Economics and Technology, Ankara, Turkey

ARTICLE INFO

Article history:
Received 15 February 2013
Received in revised form 11 September 2013

Keywords:
Digital content distribution
Digital rights management
Broadcast encryption
Trace and revoke
Anonymity
Privacy

ABSTRACT

A broadcast encryption (BE) scheme is a method for encrypting messages in a way that only a set of privileged users can decrypt it. Anonymity in a BE system is to hide any information on the privileged set. This problem has very recently had some attention and some constructions are proposed to achieve anonymity. However, anonymity in a trace and revoke (TR) scheme has not been studied yet, and to the best of our knowledge there is no construction for an anonymous TR system. In this paper, we present a generic transformation from an anonymous BE scheme into an anonymous TR scheme.

1. Introduction

Broadcast encryption (BE) is introduced in [1] and later studied in [2,3].[1] With such systems it is possible to encrypt to any chosen set of *privileged* users while the others (called *revoked users*) are precluded from the reception of the message.

A coalition of malicious users (called traitors) may use their legitimate keys to produce a pirate decryption box (called the pirate decoder) that is made available to unintended parties. Traitor tracing (TT) systems are employed [4,5] to deter users from involving in this type of piracy.

It is desired to integrate both revocation and tracing functionalities in a single system. However, combining these two features is not always easy as pointed in [6–8]. A non-trivial construction, called a trace and revoke scheme (TR), is proposed by Naor and Pinkas in [9]. Later, Naor et al. [6] proposed a TR scheme which employs a weaker tracing strategy that focuses on disabling the pirate decoder rather than identifying traitors. Yet this strategy has been shown to have its own weakness, called pirate evolution [10]. Other successful TR systems proposed so far include [7,11].

A common shortcoming of all the systems above is that the anonymity of the privileged set is not preserved in the ciphertext. Despite its importance, anonymity for BE schemes has not been considered until [12–15].

Barth et al. [12] construct an anonymous BE scheme that is secure in the random oracle model by employing a *key-private* (i.e., ciphertexts do not leak public key information) IND-CCA2 secure public key encryption scheme. Fazio and Perera [13] propose an *outsider anonymous* BE scheme where the identities are hidden only from unauthorized users (outsiders).

[☆] This work was supported in part by the Turkish Scientific and Technological Research Agency (TÜBİTAK), under grant number 111E213.

[*] Corresponding author. Tel.: +90 5366057644.
E-mail addresses: muratakcs@gmail.com, muratak@akdeniz.edu.tr (M. Ak), serdar.pehlivanoglu@zirve.edu.tr (S. Pehlivanoğlu), aliaydinselcuk@gmail.com (A.A. Selçuk).

[1] In this manuscript, we put a very dense list of references due to the lack of space.

0377-0427/$ – see front matter © 2013 Elsevier B.V. All rights reserved.
http://dx.doi.org/10.1016/j.cam.2013.10.019

A recent work by Libert et al. [15] achieves IND-CCA2 anonymity. A lower bound due to the size of the description of the privileged set is also provided. Kiayias and Samari [14] later improved this lower bound by showing that the number of encryptions in an anonymous BE scheme has to be at least linear in the length of the privileged set.

Anonymity of a TR scheme has not been studied yet, and to the best of our knowledge there is no anonymous TR system. In this paper, we present a generic transformation of an anonymous BE scheme into an anonymous TR scheme. The transformation preserves the public and private key sizes of the underlying scheme and expands the ciphertext length by a factor of two in the worst case.

When we apply our transformation to the anonymous BE schemes of Libert et al. [15] (which are in fact IND-CCA2 secure), we obtain a fully anonymous TR scheme with the same efficiency performance of [15] but our transformation inherits IND-CCA1 security rather than IND-CCA2. Another instantiation with [13] leads to a TR scheme with a weaker type of anonymity, called outsider-anonymity, while achieving a ciphertext length of $O(s)$ where s is the size of the privileged set.

2. Preliminaries

2.1. Anonymous broadcast encryption

A BE scheme consists of three algorithms: (1) $\texttt{KeyDist}(1^n)$ generates private keys $sk_i : i \in [n]$ (throughout the paper we will denote the set $\{1, 2, \dots, n\}$ by $[n]$) and a public key PK. (2) $\texttt{Encrypt}(PK, S, m)$, on input message m and a set S, prepares a ciphertext c. (3) $\texttt{Decrypt}(PK, sk_i, c)$ responds with m if and only if $i \in S$. Those users in S are called privileged users while the rest are called revoked users.

In such schemes, even if all revoked users in $[n]\backslash S$ collude to decrypt a ciphertext intended for the set S, they should not be able to get any useful information about the message. This confidentiality feature is formalized below in Game 1 (see [2,3]).

Game 1 IND-CCA1 confidentiality game for BE schemes.

```
1: Initialize. The challenger $\mathcal{C}$ runs $(\{sk_i\}_{i\in[n]}, PK) \leftarrow \texttt{KeyDist}(1^n)$ and sends $PK$ to a probabilistic polynomial time (PPT) adversary $\mathcal{A}$.
2: Query Phase. $\mathcal{A}$ can corrupt polynomially many users, denoted by set $R$, and capture the keys $\{sk_i\}_{i\in R}$. A CCA1 adversary can also make polynomially many decryption queries $(i, c)$ to which the challenger $\mathcal{C}$ responds with $\texttt{Decrypt}(PK, sk_i, c)$.
3: Challenge. The adversary provides a set $S^*$ which satisfies $S^* \cap R = \emptyset$, and two messages $m_0, m_1$. $\mathcal{C}$ chooses $b \in_R \{0, 1\}$ and prepares $c^* \leftarrow \texttt{Encrypt}(PK, S^*, m_b)$ and sends $c^*$ to $\mathcal{A}$.
4: Guess. $\mathcal{A}$ guesses $b'$ for $b$.
```

Throughout the paper, we say an adversary playing a security game (e.g. Game 1 above) wins if it guesses correctly, i.e. if $b' = b$ holds. In general, we define the advantage of an adversary in a security game as $\mathrm{Adv}_{\mathcal{A}} = |\Pr[\mathcal{A} \text{ wins}] - 1/2|$.

Definition 1. A BE scheme B is IND-CCA1 secure if $\mathrm{Adv}_{\mathcal{A}}$ is negligible for any PPT adversary $\mathcal{A}$ playing Game 1.

In an anonymous broadcast encryption scheme, the adversary should be unable to distinguish between any two equal-sized sets of privileged users as long as the corrupted users do not cover the symmetric difference of the two sets. Following the terminology of [14], we define ANO-CCA1 anonymity via the following Game 2.

Game 2 ANO-CCA1 anonymity game for BE schemes.

```
1: Initialize. The challenger $\mathcal{C}$ runs $(\{sk_i\}_{i\in[n]}, PK) \leftarrow \texttt{KeyDist}(1^n)$ and sends $PK$ to a PPT adversary $\mathcal{A}$.
2: Query Phase. $\mathcal{A}$ can corrupt polynomially many users, denoted by set $R$, and captures the keys $\{sk_i\}_{i\in R}$. A CCA1 adversary can also make polynomially many decryption queries $(i, c)$ to which the challenger $\mathcal{C}$ responds with $\texttt{Decrypt}(PK, sk_i, c)$.
3: Challenge. $\mathcal{A}$ provides a message $m$ and two equal-sized sets $S_0, S_1$ which satisfy $(S_0 \Delta S_1) \cap R = \emptyset$, where $S_0 \Delta S_1 = (S_0 \setminus S_1) \cup (S_1 \setminus S_0)$. $\mathcal{C}$ chooses $b \in_R \{0, 1\}$ and prepares $c^* \leftarrow \texttt{Encrypt}(PK, S_b, m)$ and sends $c^*$ to $\mathcal{A}$.
4: Guess. $\mathcal{A}$ guesses $b'$ for $b$.
```

Definition 2. A BE scheme B is *priv-eq* ANO-CCA1 secure if $\mathrm{Adv}_{\mathcal{A}}$ is negligible for any PPT adversary $\mathcal{A}$ playing Game 2.

More generally, we define a BE scheme B to be *priv-full* ANO-CCA1 secure if the challenge sets can be any two subsets of $[n]$.

2.2. Trace and revoke scheme

A trace and revoke (TR) scheme is a multiuser encryption system that supports both revocation (as in BE) and piracy detection (as in traitor tracing schemes). In this paper, we consider an adversarial setting where the adversary corrupts a number of user keys (that we call traitor keys) and produces a pirate decoder which succeeds in decrypting ciphertexts intended for at least one subset. If the set S (or distribution over sets) for which the decoder works with non-negligible probability is infeasible to sample then it is straightforward that no tracing can take place: given the infeasibility of sampling, the tracer will fail to ever witness the decoder to work. Due to this impossibility, in the below we will assume (as is also assumed by previous works of [7,11]) without loss of generality that a set S, for which a non-negligible probability of successful decryption holds, is known to the tracer. Note that this necessary assumption is quite reasonable in practice as well: typically the set S is selected by the distribution center and it is entirely outside the control of the pirate. We denote a pirate decoder that is capable of correctly decrypting broadcasts to set S (i.e. ciphertexts generated by $\texttt{Encrypt}(PK, S, m)$) by $\mathcal{D}_S$ which can be modeled as a probabilistic polynomial time (PPT) algorithm.

Throughout the paper, we consider $\mathcal{D}_S$ to be perfect, i.e., it decrypts the ciphertexts generated by $\texttt{Encrypt}(PK, S, m)$ with probability 1. An imperfect decoder can be dealt with employing δ-robust Boneh–Shaw fingerprinting codes of [4], which will not be elaborated further in this paper.

In addition to the three algorithms of a BE scheme, it is equipped with a tracing algorithm `Trace` which gets as input a pirate decoder $\mathcal{D}_S$ that is produced by a PPT adversary corrupting a set of users T with $|T| \leq t$. `Trace` outputs a user u who has contributed in the construction of the pirate decoder $\mathcal{D}_S$. We say the algorithm succeeds if $u \in T$ holds, i.e. $\texttt{Trace}(S, \mathcal{D}_S, PK) \in T$ where $\mathcal{D}_S$ is produced by using the traitor keys $\{sk_i\}_{i \in T}$.

In this paper, we consider black-box tracing where decoders cannot be reverse engineered and the keys inside cannot be revealed directly. In this setting, the `Trace` algorithm has black-box access to the pirate decoder, i.e., it is only allowed to query with ciphertexts and observe the way the decoder responds. The pirate decoders are also assumed to be resettable (does not maintain state during tracing) and available (remains available as long as the tracing process continues). In the literature, almost all of the positive results in designing traitor tracing schemes (including the schemes that we compare to our constructions) are successful against such decoders.

Since a TR scheme is a BE scheme with a tracing procedure and tracing is not relevant to IND-CCA1 security and anonymity, we do not need to define IND-CCA1 security and anonymity for TR schemes separately. They are inherited from their BE counterparts. We need one new security definition for TR schemes regarding tracing capability, which we state as follows.

Definition 3. We say $\texttt{T}$ is a TR scheme against a coalition $\mathcal{T}$ of t traitors, if, given a PPT pirate decoder $\mathcal{D}_S$ forged by $\mathcal{T}$, the PPT `Trace` algorithm of $\texttt{T}$ succeeds to find a traitor index with non-negligible probability.

2.3. Boneh–Shaw fingerprinting code

The Boneh–Shaw fingerprinting code, a pair of $\langle\texttt{CodeGen}, \texttt{Identify}\rangle$ algorithms, plays a crucial role in our generic construction. Due to the lack of space, only a brief description will be presented here which would be sufficient to follow and check the correctness and traceability of our construction. For a complete discussion on the Boneh–Shaw code, we refer the reader to [16].

The `CodeGen` algorithm of the Boneh–Shaw fingerprinting code works as follows: a binary matrix of n by nd is formed[2] where n is the number of codewords, $d = O(n^2 \log(1/\epsilon))$ holds and ϵ is a security parameter. The columns are grouped in n blocks each of length d. The first i blocks of the i-th row consist of all 0s and the remaining $n - i$ blocks consist of all 1's. Below, an example matrix for $n = 4$, $d = 3$ is given on the left:

$$\begin{pmatrix} 0&0&0&1&1&1&1&1&1&1&1&1\\ 0&0&0&0&0&0&1&1&1&1&1&1\\ 0&0&0&0&0&0&0&0&0&1&1&1\\ 0&0&0&0&0&0&0&0&0&0&0&0 \end{pmatrix} \xrightarrow{\pi} \begin{pmatrix} 0&1&1&1&0&1&1&1&0&1&1&1\\ 0&1&1&0&0&1&1&0&0&0&1&1\\ 0&0&1&0&0&1&0&0&0&0&1&0\\ 0&0&0&0&0&0&0&0&0&0&0&0 \end{pmatrix}.$$

Finally, a permutation π over the columns of the matrix is applied as shown above. The rows of the resulting matrix constitute the codewords of the (ℓ, n)-code with length $\ell = nd$. We say a piracy has occurred if a pirate codeword $p \in \{0, 1\}^\ell$ is produced from some codewords (we call them traitor codewords). The pirate codeword is produced under a marking assumption, which essentially says if all traitor codewords have 0 (or 1) in a particular column, the corresponding column of the pirate codeword must be 0 (resp. 1) as well.

The `Identify` algorithm of Boneh–Shaw fingerprinting code is given a pirate codeword p and it returns a row-index which is supposed to be the index of a traitor codeword. The secrecy of the permutation π is crucial for the correctness of the code. We say the code is (ϵ, t)-identifier if `Identify` algorithm identifies a traitor regardless of the traitor strategy upon observation of a pirate codeword produced by a coalition of t traitors.

3. Generic transformation

In this section, we show how to transform an anonymous BE scheme into an anonymous TR scheme. A message is broadcasted for a privileged set S as follows: first S is partitioned into $P = \{S_0, S_1\}$ and the message is encrypted for each subset S_0, S_1 separately using the BE scheme. This transformation preserves the revocation capability of the underlying BE scheme.

[2] We deviate slightly from the original description where the matrix is n by $(n - 1)d$. This change does not affect the correctness of the code.

Let $\mathtt{B}$ be an anonymous BE scheme consisting of three algorithms $\mathtt{BKeyDist}$, $\mathtt{BEncrypt}$, and $\mathtt{BDecrypt}$. We design the algorithms of our generic scheme $\mathtt{T}$ as follows.

$\mathtt{TKeyDist}(1^n)$ runs $\mathtt{BKeyDist}(1^n)$ which produces a public key PK_{B} and set of private keys sk_i, $i \in [n]$. The public key PK_{T} is set to be PK_{B}.

$\mathtt{TEncrypt}(PK_{\mathrm{T}}, S, m)$ parses PK_{B} and chooses a random split (S_0, S_1) of the set S where S_0 and S_1 are disjoint sets and $S = S_0 \cup S_1$ holds. The algorithm outputs $c = \langle c_0, c_1 \rangle$ where $c_b \leftarrow \mathtt{BEncrypt}(PK_{\mathrm{B}}, S_b, m)$ for $b \in \{0, 1\}$.

$\mathtt{TDecrypt}(PK_{\mathrm{T}}, sk_i, c)$ parses PK_{B} and both c_0 and c_1 from the input. The algorithm tries to decrypt both c_0 and c_1 by sk_i. If one of them, say c_b, leads to a valid decryption, then return $BDecrypt(PK_{\mathrm{B}}, sk_i, c_b)$.[3]

Our tracing algorithm follows a similar tracing strategy of the works by [4,8,17]. Fixing a particular partition $P = \{S_0, S_1\}$ of set $S = S_0 \cup S_1$, two types of tracing ciphertext are constructed: in type $(P, 0)$ we prepare $c = \langle c_0, c_1 \rangle$ where $c_0 \leftarrow \mathtt{BEncrypt}(PK_{\mathrm{B}}, S_0, m)$ and $c_1 \leftarrow \mathtt{BEncrypt}(PK_{\mathrm{B}}, S_1, m)$; in type $(P, 1)c_0$ is changed to be an encryption of a random message m_r.

Note that the message m can be retrieved by all privileged users in the tracing transmission of type $(P, 0)$. If a tracing transmission of type $(P, 1)$ is decryptable by the pirate decoder $\mathcal{D}_S$ (that is constructed by a traitor coalition T) with probability greater than 1/2, then we conclude the existence of a traitor in set S_1. Otherwise, we observe a substantial gap (greater than 1/2 success probability drop) in the behavior of the decoder for tracing transmissions of type $(P, 0)$ and $(P, 1)$ which suggests that a traitor exists in set S_0.

We extend the above basic strategy into a full-fledged tracing algorithm for our generic construction. Toward this goal, we will simulate the random partitions of the $\mathtt{TEncrypt}$ algorithm by employing a Boneh–Shaw fingerprinting code $\mathcal{C} = \{c_1, \ldots, c_n\}$ of length ℓ. First, the user indices are shuffled by a random permutation π and some *column index* $j \in [\ell]$ value is chosen. A partition of type (π, j), denoted by $P_{\pi,j} = \{S_{j,0}, S_{j,1}\}$, is constructed by setting $S_{j,0} = S \cap \{v : c_{\pi(v)}[j] = 0\}$ and $S_{j,1} = S \cap \{v : c_{\pi(v)}[j] = 1\}$. Due to the random choice of the permutation π and the choice of j, a partition of type (π, j) is indistinguishable from a random split of the set S.

Fixing the permutation π, we query the pirate decoder with tracing ciphertexts of types $(P_{\pi,j}, 0)$, $(P_{\pi,j}, 1)$ for each $1 \le j \le \ell$. Based on the success probability of the pirate decoder in decrypting those ciphertexts we conclude the existence of a traitor in either set $S_{j,0}$ or set $S_{j,1}$ and further produce a pirate codeword w (we denote the jth bit by w_j): we set $w_j = 1$ if a traitor exists in set $S_{j,1}$ and set $w_j = 0$ otherwise. We argue that the pirate codeword would be in the descendant set of the codewords available to the user-set $T_\pi = \{\pi(u) : u \in T\}$, i.e. if $w_j = b$ then $S_{j,b} \cap T_\pi \neq \emptyset$ holds. The correctness of this argument can be trivially checked and will be elaborated later in Section 4. Finally, the $\mathtt{Identify}$ algorithm of the Boneh–Shaw code, on input w, will identify a user index t from set T_π, i.e. $\pi^{-1}(t)$ would be a traitor-index involved in the production of the pirate decoder.

- $\mathtt{Trace}(S, \mathcal{D}_S, PK_{\mathrm{T}})$ first parses PK_{B} from PK_{T}. It produces a Boneh–Shaw fingerprinting code $\mathcal{C} = \{c_1, \ldots, c_n\} \leftarrow \mathtt{CodeGen}(1^n)$, initializes a pirate codeword $w \leftarrow 0^\ell$ and chooses a permutation π. Denoting the length of the code $\mathcal{C}$ by ℓ, it creates partitions $P_{\pi,j} = \{S_{j,0}, S_{j,1}\}$ for each $j = 1, \ldots, \ell$ such that $S_{j,0} = S \cap \{v : c_{\pi(v)}[j] = 0\}$ and $S_{j,1} = S \cap \{v : c_{\pi(v)} [j] = 1\}$.
  A tracing ciphertext $c = \langle c_0, c_1 \rangle$ of type $(P_{\pi,j}, 1)$ is prepared where $c_0 \leftarrow \mathtt{BEncrypt}(PK_{\mathrm{B}}, S_{j,0}, m_r)$ for a randomly chosen m_r and $c_1 \leftarrow \mathtt{BEncrypt}(PK_{\mathrm{B}}, S_{j,1}, m)$. We say that the pirate decoder $\mathcal{D}_S$ succeeds in decrypting tracing ciphertext of type $(P_{\pi,j}, 1)$ if it returns m. The success probability of $\mathcal{D}_S$ denoted by $p_{j,1}$ is approximated. We set $w_j = 1$ if $p_{j,1} \ge 1/2$ and set $w_j = 0$ otherwise.
  By repeating the above for $j \in [\ell]$, we produce a pirate codeword w and compute $t \leftarrow \mathtt{Identify}(w)$ for which $\pi^{-1}(t)$ would be a traitor-index involved in the production of the pirate decoder.

4. Security properties of the construction

Lemma 1 (*Confidentiality*). *Let B be an* IND-CCA1 *secure BE scheme employed in our generic construction of the Section 3. The resulting trace and revoke scheme T is also* IND-CCA1 *secure.*

Proof. Let us denote the original confidentiality Game 1 of Section 2 by $\mathbf{G_0}$. For a particular choice of PPT adversary A_T, we define its advantage in $\mathbf{G_0}$ by $\mathrm{Adv}_{A_T}[0] = |Pr[\mathcal{A}_T \text{ wins in } \mathbf{G_0}] - 1/2|$. In this game the challenger is given a triple (S^*, m_0, m_1) which satisfies $S^* \cap R = \emptyset$ where R is the set of corrupted users whose keys are possessed by the adversary. The challenger prepares a ciphertext $c^* \leftarrow \mathtt{TEncrypt}(PK, S^*, m_b)$ for a randomly chosen $b \in \{0, 1\}$. Due to the nature of our transformation, the set S^* is partitioned into two subsets S_0^* and S_1^* and the challenge is prepared as $c^* = \langle c_0^*, c_1^* \rangle$ where $c_d^* \leftarrow \mathtt{BEncrypt}(PK, S_d^*, m_b)$ holds for $d \in \{0, 1\}$.

We deviate slightly from the above game by preparing a random encryption for the subset S_0^*. Denoting this new game by $\mathbf{G_1}$, c_0^* is set to be $\mathtt{BEncrypt}(PK, S_0^*, m_r)$ for a randomly chosen message m_r. The rest of the game is identical to the Game $\mathbf{G_0}$. We define the advantage of the adversary A_T in Game $\mathbf{G_1}$ by $\mathrm{Adv}_{A_T}[1] = |Pr[\mathcal{A}_T \text{ wins in } \mathbf{G_1}] - 1/2|$.

[3] Deciding a valid decryption, i.e. applying a wrong key to a ciphertext results to a special fail message to be returned requires some sort of strong correctness property which is assumed for our construction in this paper. Motivating further exploration in this quest, a possible solution can be achieved e.g. by appending a value $H(m)$ (where H is a hash function) to the plaintext m; we omit further details.

The difference between $\mathbf{G_0}$ and $\mathbf{G_1}$ is replacing $\texttt{BEncrypt}(PK, S_0^*, m_b)$ with $\texttt{BEncrypt}(PK, S_0^*, m_r)$. Following the standard proof arguments, the adversary A_T can be used to distinguish between the games $\mathbf{G_0}$ and $\mathbf{G_1}$ with an advantage of $|\mathrm{Adv}_{A_T}[1] - \mathrm{Adv}_{A_T}[0]|/2$. This can further be used to mount an attack to distinguish $\texttt{BEncrypt}(PK, S_0^*, m_b)$ from $\texttt{BEncrypt}$ (PK, S_0^*, m_r) using the standard game hopping arguments (we omit the details due to the lack of space). Hence, it holds that $|\mathrm{Adv}_{A_T}[1] - \mathrm{Adv}_{A_T}[0]| \leq 2\mathrm{Adv}_{A_B}$ since any PPT adversary can win the confidentiality game of BE with advantage at most Adv_{A_B}.

We take one step ahead and create a new game, denoted by $\mathbf{G_2}$, where we set c_1^* to be $\texttt{BEncrypt}(PK, S_1^*, m_r)$, i.e. c_1^* also encrypts to the random message in this game. The rest of the game is identical to the Game $\mathbf{G_1}$. We define the advantage of the adversary A_T in Game $\mathbf{G_2}$ by $\mathrm{Adv}_{A_T}[2] = |Pr[\mathcal{A}_T \text{ wins in } \mathbf{G_2}] - 1/2|$.

The difference between $\mathbf{G_1}$ and $\mathbf{G_2}$ is replacing $\texttt{BEncrypt}(PK, S_1^*, m_b)$ with $\texttt{BEncrypt}(PK, S_1^*, m_r)$. Following the similar arguments above, we obtain $|\mathrm{Adv}_{A_T}[2] - \mathrm{Adv}_{A_T}[1]| \leq 2\mathrm{Adv}_{A_B}$. Note that, in Game $\mathbf{G_2}$, the adversary A_T is challenged with encryption of a random message that is neither m_0 nor m_1. Hence, the advantage of the adversary winning in Game $\mathbf{G_2}$ is negligible. Applying the triangular inequality we obtain $\mathrm{Adv}_{A_T}[0] \leq 4\mathrm{Adv}_{A_B} + negl$ which completes the proof of the lemma. □

Lemma 2 (*Anonymity*). *Let B be a* priv-eq ANO-CCA1 *secure BE scheme employed in our generic construction of the Section* 3. *The resulting trace and revoke scheme T is also* priv-eq ANO-CCA1 *secure. Our transformation preserves anonymity for the* priv-full *notion as well.*

Proof. Throughout the proof, we will focus on the *priv-eq* ANO-CCA1 security. By relaxing the conditions (which will be clear later in the proof) on the cardinality of the partitions, our proof works for the *priv-full* notion as well.

Let us recall the anonymity game 2, denoted by $\mathbf{G}_T$, defined for the scheme T. A PPT adversary A_T of the game provides the triple (m, S_0, S_1) as a challenge query. The challenger picks $b \in \{0, 1\}$ randomly and prepares the challenge ciphertext $c_b^* = \langle c_{b,0}^*, c_{b,1}^* \rangle$ where $c_{b,0}^* \leftarrow \texttt{BEncrypt}(PK_{\mathrm{B}}, S_{b,0}, m)$ and $c_{b,1}^* \leftarrow \texttt{BEncrypt}(PK_{\mathrm{B}}, S_{b,1}, m)$ hold for a random split $(S_{b,0}, S_{b,1})$ of the set S.

Through a sequence of games, we will show that $b = 0$ case in game $\mathbf{G}_T$ (denoted by $\mathbf{G}_{T_0}$) is indistinguishable from $b = 1$ case (denoted by $\mathbf{G}_{T_1}$). Let $\mathbf{G_0}$ be a new game almost identical to the game $\mathbf{G}_{T_0}$ where, in addition, a random partition $S_1 = S_{1,0} \cup S_{1,1}$ is also computed for the subset S_1. This partition will be a conditionally random split strictly bounded by the following equalities: $|S_{0,0}| = |S_{1,0}|$, $|S_{0,1}| = |S_{1,1}|$, $S_{0,0} \cap S_{1,1} = \emptyset$ and $S_{0,1} \cap S_{1,0} = \emptyset$. Since the partition of the set S_1 is not used in the challenge, a direct reduction from game $\mathbf{G}_0$ to the game $\mathbf{G}_{T_0}$ is trivially available.

We define a new game $\mathbf{G_1}$ which is almost identical to $\mathbf{G_0}$: instead of preparing the challenge ciphertext for the subsets $S_{0,0}$ and $S_{0,1}$, a pair of encryption $(c_{0,0}^*, c_{1,1}^*)$ is produced for the subsets $S_{0,0}$ and $S_{1,1}$. The difference in success probabilities of games $\mathbf{G_0}$ and $\mathbf{G_1}$ is bounded by distinguishing if a ciphertext is prepared for $S_{0,1}$ or $S_{1,1}$. Since both of these subsets do not intersect with subset $S_{0,0}$, an adversary distinguishing these two games can be turned into an adversary (through standard ways which are omitted due to the lack of space) which is capable of breaking the *priv-eq* ANO-CCA1 security of the BE scheme.

We take a another step and define a new game $\mathbf{G_2}$ which is almost identical to $\mathbf{G_1}$: instead of preparing the challenge ciphertext for the subsets $S_{0,0}$ and $S_{1,1}$, a pair of encryption $(c_{1,0}^*, c_{1,1}^*)$ is produced for the subsets $S_{1,0}$ and $S_{1,1}$. Similarly, these two games are indistinguishable with respect to our *priv-eq* ANO-CCA1 security assumption of the BE scheme.

In all these three games $\mathbf{G_0}$, $\mathbf{G_1}$ and $\mathbf{G_2}$, a conditionally random split (those conditions are listed above) for S_1 is computed after a random split for S_0 is chosen. In our final game $\mathbf{G_3}$, we change this order: we first compute a random split for S_1 and later a random split for S_0 satisfying the equations/conditions stated above. Since the challenge ciphertexts of game $\mathbf{G_2}$ and $\mathbf{G_3}$ are independent from the partition of S_0, a conditionally random split of S_1 in game $\mathbf{G_2}$ would be indistinguishable from a random split of S_1 in game $\mathbf{G_3}$. Since the partition of the set S_0 is not used in the challenge, a direct reduction from game $\mathbf{G}_{T_1}$ to the game $\mathbf{G}_3$ is trivially available. Hence, the success probability of an adversary winning in game $\mathbf{G}_T$ is bounded with the advantage of breaking the anonymity of the BE scheme. □

Lemma 3 (*Traceability*). *Consider* IND-CCA1 *secure BE scheme B and (ϵ_f, t)-identifier Boneh–Shaw code F. The generic construction of Section* 3*, employing the schemes B and F, produces a trace and revoke scheme against a coalition of t traitors.*

Proof. In order to prove this lemma, we will show that the tracing algorithm of T succeeds (i.e. detects a traitor) against any resettable and available pirate decoder $\mathcal{D}_S$ constructed by a coalition of size at most t, with a probability $1 - \epsilon_f - \ell\epsilon$ by using an (ϵ_f, t)-identifier Boneh–Shaw code with codeword length ℓ, where ϵ is a negligible probability that can arbitrarily be decreased by increasing the number of experiments in the tracing procedure.

The tracing process can be considered as three stages: approximating the success probability of the decoder in decrypting tracing ciphertexts, producing the pirate codeword w, and finally identifying a traitor index.

We would like to first argue that the tracing partitions and the regular partitions are chosen from the same distribution space. A tracing partition of type (π, j), denoted by $P_{\pi,j} = \{S_{j,0}, S_{j,1}\}$, is constructed by setting $S_{j,0} = S \cap \{v : c_{\pi(v)}[j] = 0\}$ and $S_{j,1} = S \cap \{v : c_{\pi(v)}[j] = 1\}$ where $\mathcal{C} = \{c_1, \ldots, c_n\}$ of length ℓ is a fingerprinting code that is produced by running $\texttt{CodeGen}(1^n)$ algorithm of Boneh–Shaw code. The random choice of the permutation π and the choice of j makes the tracing partition of type (π, j) indistinguishable (a full proof is omitted due to the lack of space) from a random split of the set S.

Before proceeding with the analysis of the tracing stages, let us introduce some notation: we call the number of approximation queries by λ, the expected number of times the decoder succeeds in decrypting tracing ciphertext of type

$(P_{\pi,j}, b)$ is denoted by $\mu_{j,b} = \lambda \cdot \sigma_{j,b}$ and the actual number of successes during the approximation process is denoted by $\rho_{j,b} = \lambda \cdot p_{j,b}$ where $\sigma_{j,b}$ and $p_{j,b}$ are defined as the success probabilities respectively.

Approximation phase: due to the allowed resettability of the decoder after each tracing query, we can keep the same permutation π fixed throughout tracing process. Applying a two-tailed form of the Chernoff bound (since the consecutive tracing experiments are independent due to resettability), for the choice of $\lambda = 48\ln(2/\epsilon)$, we obtain $|\rho_{j,1} - \mu_{j,1}| \geq \lambda/4$ with probability at most ϵ.

$$\Pr[|\rho_{j,1} - \mu_{j,1}| \geq \alpha] \leq 2e^{-\frac{\alpha^2}{3\mu_{j,1}}} \leq 2e^{-\frac{\alpha^2}{3\lambda}}.$$

Substituting $\alpha = \lambda/4$ and $\lambda = 48\ln(2/\epsilon)$ we obtain

$$2e^{-\alpha^2/3\lambda} = 2e^{-\frac{\lambda^2}{3\cdot 16\cdot\lambda}} = 2e^{-\lambda/48} = 2e^{-\ln(2/\epsilon)} = \epsilon.$$

Conclusion: $|\rho_{j,1} - \mu_{j,1}| \leq \lambda/4$, which is equivalent of saying $|p_{j,1} - \sigma_{j,1}| \leq 1/4$, holds with probability at least $1 - \epsilon$.

Pirate codeword generation: we now prove that w is in the descendent set of the codewords available to the user-set $T_\pi = \{\pi(u) : u \in T\}$ with high probability: i.e., for all $1 \leq j \leq \ell$, if $w_j = b$ then $S_{j,b} \cap T_\pi \neq \emptyset$ for $b \in \{0, 1\}$.

Consider the case $w_j = 0$: due to the tracing strategy this can happen only when $p_{j,1} < 1/2$ holds. Assume the contradiction of our claim that is $S_{j,0} \cap T_\pi = \emptyset$, i.e. there exists no traitor in $S_{j,0}$. Since we consider only perfect decoders it holds that $\sigma_{j,0} = p_{j,0} = 1$. It further holds that $|\sigma_{j,1} - p_{j,1}| \leq 1/4$ with probability at least $1 - \epsilon$ due to approximation.

Replacing all these into the following triangular inequality $|\sigma_{j,0} - \sigma_{j,1}| + |\sigma_{j,1} - p_{j,1}| \geq |\sigma_{j,0} - p_{j,1}|$, we obtain $|\sigma_{j,0} - \sigma_{j,1}| \geq 1/4$. This suggests that the pirate decoder is capable of distinguishing the tracing ciphertexts of type $(P_{\pi,j}, 0)$ from $(P_{\pi,j}, 1)$ with probability at least 1/4, hence distinguishing whether a random message or the actual message is broadcasted to the set $S_{j,0}$. Since our contradiction assumption states that $S_{j,0} \cap T_\pi = \emptyset$, we can use such capability of the pirate decoder to break the IND-CCA1 security claim of the scheme B. Thus we end up with a contradiction.

Consider now the case $w_j = 1$, so does $p_{j,1} \geq 1/2$: following the similar analysis we obtain $\sigma_{j,1} \geq 1/4$ with probability at least $1 - \epsilon$. Under the security assumptions of the underlying BE scheme, we conclude that there should exist a traitor in set $S_{j,1}$.

The above is repeated for all j values. Applying the union probability over the choices of j, we conclude that the constructed pirate codeword is in the descendent set of traitors with probability at least $1 - \ell\epsilon$.

Traitor identification. Finally, `Identify`(w) algorithm of Boneh–Shaw returns a traitor index with probability at least $1 - \epsilon_f$. This completes the proof of the traceability. The overall failure probability of accusing an innocent user is bounded by $\epsilon_f + \ell\epsilon$ (for the failures in identification and approximations, respectively) for the given parameters. □

Acknowledgments

We thank Aggelos Kiayias for his many contributions and his valuable comments. We also thank the anonymous referees for their valuable and helpful comments on the manuscript.

References

[1] A. Fiat, M. Naor, Broadcast encryption, in: CRYPTO'93.
[2] D. Boneh, C. Gentry, B. Waters, Collusion resistent broadcast encryption with shorter ciphertexts and private keys, in: CRYPTO'05.
[3] C. Gentry, B. Waters, Adaptive security in broadcast encryption systems (with short ciphertexts), in: EUROCRYPT'09.
[4] D. Boneh, M. Naor, Traitor tracing with constant size ciphertext, in: 15th ACM Conference, CCS'08, 2008, pp. 501–510.
[5] B. Chor, A. Fiat, M. Naor, Tracing traitors, in: CRYPTO'94.
[6] D. Naor, M. Naor, J. Lotspiech, Revocation and tracing schemes for stateless receivers, in: CRYPTO'01.
[7] D. Boneh, B. Waters, A fully collusion resistent broadcast, trace, and revoke system, in: CCS'06.
[8] A. Kiayias, S. Pehlivanoğlu, Tracing and revoking pirate rebroadcasts, in: ACNS'09.
[9] M. Naor, B. Pinkas, Efficient trace and revoke schemes, in: FC'00.
[10] A. Kiayias, S. Pehlivanoğlu, Pirate evolution: how to make the most of your traitor keys, in: CRYPTO'07.
[11] J. Furukawa, N. Attrapadung, Fully collusion resistant black-box traitor revocable broadcast encryption with short private keys, in: ICALP'07.
[12] A. Barth, D. Boneh, B. Waters, Privacy in encrypted content distribution using private broadcast encryption, in: FC'06.
[13] N. Fazio, I.M. Perera, Outsider-anonymous broadcast encryption with sublinear ciphertexts, in: PKC'12.
[14] A. Kiayias, K. Samari, Lower bounds for private broadcast encryption, in: 14th Information Hiding Conf., Berkeley, California, 2012.
[15] B. Libert, K.G. Paterson, E.A. Quaglia, Anonymous broadcast encryption: adaptive security and efficient constructions in the standard model, in: PKC'12.
[16] D. Boneh, J. Shaw, Collusion-secure fingerprinting for digital data, IEEE Trans. Inform. Theory 44 (5) (1998) 1897–1905.
[17] A. Kiayias, S. Pehlivanoglu, Improving the round complexity of traitor tracing schemes, in: ACNS'10.