Practical Ciphertext-Policy Attribute-Based Encryption: Traitor Tracing, Revocation, and Large Universe

Zhen Liu[1(✉)] and Duncan S. Wong[2]

[1] City University of Hong Kong, Hong Kong SAR, China
zhenliu7-c@my.cityu.edu.hk
[2] Security and Data Sciences, ASTRI, Hong Kong SAR, China
duncanwong@astri.org

Abstract. In Ciphertext-Policy Attribute-Based Encryption (CP-ABE), a user's decryption key is associated with attributes which in general are not related to the user's identity, and the same set of attributes could be shared between multiple users. From the decryption key, if the user created a decryption blackbox for sale, this malicious user could be difficult to identify from the blackbox. Hence in practice, a useful CP-ABE scheme should have some tracing mechanism to identify this 'traitor' from the blackbox. In addition, being able to revoke compromised keys is also an important step towards practicality, and for scalability, the scheme should support an exponentially large number of attributes. However, none of the existing traceable CP-ABE schemes simultaneously supports revocation and large attribute universe. In this paper, we construct the first practical CP-ABE which possesses these three important properties: (1) blackbox traceability, (2) revocation, and (3) supporting large universe. This new scheme achieves the fully collusion-resistant blackbox traceability, and when compared with the latest fully collusion-resistant blackbox traceable CP-ABE schemes, this new scheme achieves the same efficiency level, enjoying the sub-linear overhead of $O(\sqrt{N})$, where N is the number of users in the system, and attains the same security level, namely, the fully collusion-resistant traceability against policy-specific decryption blackbox, which is proven in the standard model with selective adversaries. The scheme supports large attribute universe, and attributes do not need to be pre-specified during the system setup. In addition, the scheme supports revocation while keeping the appealing capability of conventional CP-ABE, i.e. it is highly expressive and can take any monotonic access structures as ciphertext policies.

Keywords: Attribute-based encryption · Ciphertext-policy · Traitor tracing · Revocation · Large attribute universe

1 Introduction

In some emerging applications such as user-side encrypted cloud storage, users may store encrypted data on a public untrusted cloud and let other users who

© Springer International Publishing Switzerland 2015
T. Malkin et al. (Eds.): ACNS 2015, LNCS 9092, pp. 127–146, 2015.
DOI: 10.1007/978-3-319-28166-7_7

have eligible credentials decrypt and access the data. The decryption credentials could be based on the users' roles and do not have to be their identities. For example, a user Alice wants to encrypt some documents, upload to the cloud, and let all PhD students and alumni in the Department of Mathematics download and decrypt. *Attribute-Based Encryption* (ABE), introduced by Sahai and Waters [25], provides a solution to this type of applications. In a Ciphertext-Policy ABE (CP-ABE) [2,10] scheme[1], each user possesses a set of attributes and a decryption key, the encrypting party can encrypt the documents using an access policy (e.g. a Boolean formula) on attributes, and a user can decrypt if and only if the user's attributes satisfy the policy. Hence in this example, Alice can encrypt the documents under "(Mathematics AND (PhD Student OR Alumni))", which is an access policy defined over descriptive *attributes*, so that only those receivers whose attributes satisfy this policy can decrypt.

Among the recently proposed CP-ABE schemes [2,6,9,11,14,15,21,26], one of the latest works is due to Lewko and Waters [15]. Their scheme achieves high expressivity (i.e. can take any monotonic access structures as ciphertext policies), and is provably secure against adaptive adversaries in the standard model. The scheme is also efficient and removes the one-use restriction that other comparable schemes have [14,21]. As of the current Public Key Infrastructure which mandates the capabilities of key generation, revocation, and certified binding between identities and public keys, before the CP-ABE being able to deploy in practice, we should provision a practical CP-ABE scheme with three important features: (1) traceability, (2) revocation, and (3) large universe. Very recently, a handful of research works have been done on each one of these while the fundamental open problem remains is the existence of an efficient scheme which supports these three features at once.

Traceability / Traitor Tracing. Access policies in CP-ABE do not have to contain any receivers' identities, and more commonly, a CP-ABE policy is role-based and attributes are *shared* between multiple users. In practice, a malicious user, with attributes shared with multiple other users, might leak a decryption blackbox/device, which is made of the user's decryption key, for the purpose of financial gain or some other forms of incentives, as the malicious user has little risk of being identified out of all the users who can build a decryption blackbox with identical decryption capability. Being able to identify this malicious user is crucial towards the practicality of a CP-ABE system.

Given a well-formed decryption key, if the *tracing algorithm* of a CP-ABE scheme can identify the malicious user who created the key, the scheme is called Whitebox Traceable CP-ABE [17]. Given a decryption blackbox, while the decryption key and even the decryption algorithm could be hidden inside the blackbox, if the *tracing algorithm* can still find out the traitor whose key has been used in constructing the blackbox, the scheme is called Blackbox Traceable CP-ABE [16]. In this stronger notion, there are two types of blackboxes: key-like decryption blackbox and policy-specific decryption blackbox. A key-like

[1] Due to page limitation, here we focus on CP-ABE, while skipping discussions about Key-Policy ABE.

decryption blackbox has an attribute set associated and can decrypt encrypted messages with policies being satisfied by the attribute set. A policy-specific decryption blackbox has a policy associated and can decrypt encrypted messages with the same policy. Liu et al. [18] formally proved that if a CP-ABE scheme is traceable against policy-specific decryption blackbox then it is also traceable against key-like decryption blackbox, and proved that the CP-ABE scheme in [16] is *fully collusion-resistant traceable* against policy-specific decryption blackbox in the standard model with selective adversaries. The scheme in [16] is highly expressive, and as a fully collusion-resistant blackbox traceable CP-ABE scheme, it achieves the most efficient level to date, i.e. the overhead for the fully collusion-resistant traceability is in $O(\sqrt{N})$, where N is the number of users in the system. Note that fully collusion-resistant traceability means that the number of colluding users in constructing a decryption blackbox is not limited and can be arbitrary. Another recent blackbox traceable CP-ABE scheme is due to Deng et al. [7], but the scheme is only *t-collusion-resistant* traceable, where the number of colluding users is limited, i.e., less than a parameter t, and the scheme's security is proven in the random oracle model.

Revocation. For any encryption systems that involve many users, private keys might get compromised, users might leave or be removed from the systems. When any of these happens, the corresponding user keys should be revoked. In the literature, several revocation mechanisms have been proposed in the context of CP-ABE. In [24][2], Sahai et al. proposed an *indirect* revocation mechanism, which requires an authority to periodically broadcast a key update information so that only the non-revoked users can update their keys and continue to decrypt messages. In [1], Attrapadung and Imai proposed a *direct* revocation mechanism, which allows a revocation list to be specified directly during encryption so that the resulting ciphertext cannot be decrypted by any decryption key which is in the revocation list even though the associated attribute set of the key satisfies the ciphertext policy. The direct revocation mechanism does not need any periodic key updates that an indirect revocation mechanism requires. It does not affect any non-revoked users either. In direct revocation, a system-wide revocation list could be made public and revocation could be taken into effect promptly as the revocation list could be updated immediately once a key is revoked. In this paper, we focus on achieving direct revocation in CP-ABE.

Large Attribute Universe. In most CP-ABE schemes, the size of the attribute universe is polynomially bounded in the security parameter, and the attributes have to be fixed during the system setup. In a large universe CP-ABE, the attribute universe can be exponentially large, any string can be used as an attribute, and attributes do not need to be pre-specified during setup. Although "somewhat" large universe CP-ABE schemes have been proposed or discussed previously [1,14,22,26], as explained by Rouselakis and Waters [23], limitations

[2] Note that in this paper we focus on the conventional revocation, which is to prevent a compromised or revoked user from decrypting newly encrypted messages. In [24], revoking access on previously encrypted data is also considered.

exist. The first "truly" large universe CP-ABE construction, in which there is no restriction on ciphertext policies or attributes associated with the decryption keys, was proposed in [23].

1.1 Our Results

We propose the first practical CP-ABE scheme that simultaneously supports (1) traceability against policy-specific decryption blackbox, (2) (direct) revocation and (3) "truly" large attribute universe. The scheme's traceability is fully collusion-resistant, that is, the number of colluding users in constructing a decryption blackbox is not limited and can be arbitrary. Furthermore, the traceability is public, that is, anyone can run the tracing algorithm. The scheme is also highly expressive that allows any monotonic access structures to be the ciphertext policies.

The scheme is proven selectively secure and traceable in the standard model. This is comparable to the policy-specific blackbox traceability of the fully collusion-resistant traceable CP-ABE [18] and also to the security of the "truly" large universe CP-ABE [23]. The selective security is indeed a weakness when compared with the full security of [15,16], but as discussed in [23], selective security is still a meaningful notion and can be a reasonable trade off for performance in some circumstances. Furthermore, in light of the proof method of [15] that achieves full security through selective techniques, we can see that developing selectively secure schemes could be an important stepping stone towards building fully secure ones.

Table 1 compares this new scheme with the representative results in conventional CP-ABE [15], blackbox traceable CP-ABE [16], and large universe CP-ABE [23], all of which are provably secure in the standard model and highly expressive. The scheme's overhead is in $O(\sqrt{N})$, where N is the number of users in a system, and for fully collusion-resistant blackbox traceable CP-ABE, this is the most efficient one to date. Furthermore, when compared with the existing fully collusion-resistant blackbox traceable CP-ABE scheme in [16], at the cost of $\sqrt{N}$ additional elements in private key, our construction achieves revocation and "truly" large universe. For achieving better performance, this new scheme is constructed on prime order groups, rather than composite order groups, as it has been showed (e.g. in [8,13]) that constructions on composite order groups will result in significant loss of efficiency.

Paper Outline. In Sect. 2, we propose a definition for CP-ABE supporting policy-specific blackbox traceability, direct revocation and large attribute universe. As of [16], the definition is 'functional', namely each decryption key is uniquely indexed by $k \in \{1, \ldots, N\}$ (N is the number of users in the system) and given a policy-specific decryption blackbox, the tracing algorithm Trace can return the index k of a decryption key which has been used for building the decryption blackbox. On direct revocation, in our definition, the $\mathsf{Encrypt}$ algorithm takes a revocation list $R \subseteq \{1, \ldots, N\}$ as an additional input so that a message encrypted under the (revocation list, access policy) pair $(R, \mathbb{A})$ would only allow

Table 1. Features and efficiency comparison

[a] [b]	Blackbox traceability	Revocation	Large universe	Public key size	Ciphertext size	Private key size	Pairings in decryption
[15][c]	×	×	×	$14 + 6\|\mathcal{U}\|$	$7 + 6l$	$6 + 6\|S\|$	$9 + 6\|I\|$
[23]	×	×	✓	6	$2 + 3l$	$2 + 2\|S\|$	$1 + 3\|I\|$
[16, 18][d]	✓	×	×	$3 + 4\sqrt{N} + \|\mathcal{U}\|$	$17\sqrt{N} + 2l$	$4 + \|S\|$	$10 + 2\|I\|$
this work	✓	✓	✓	$5 + 5\sqrt{N}$	$16\sqrt{N} + 3l$	$2 + 2\|S\| + \sqrt{N}$	$9 + 3\|I\|$

[a] All the four schemes are provably secure in the standard model and highly expressive.
[b] Let N be the number of users in the system, $|\mathcal{U}|$ the size of the attribute universe, l the number of rows of the LSSS matrix for an access policy, $|S|$ the size of the attribute set of a decryption key, and $|I|$ the number of attributes for a decryption key to satisfy a ciphertext policy.
[c] The efficiency evaluation here is based on the prime order construction in the full version.
[d] The construction in [16,18] is on composite order groups where the group order is the product of three large primes, and the efficiency evaluation is based on the composite order groups.

users whose (index, attribute set) pair (k, S) satisfies $(k \notin R) \wedge (S\ satisfies\ \mathbb{A})$ to decrypt.

On the construction, we refer to the 'functional' CP-ABE in Sect. 2 as Revocable CP-ABE (R-CP-ABE), then extend the R-CP-ABE to a primitive called Augmented R-CP-ABE (AugR-CP-ABE), which will lastly be transformed to a policy-specific blackbox *traceable* R-CP-ABE. More specifically, in Sect. 3, we define the encryption algorithm of AugR-CP-ABE as $\mathsf{Encrypt}_{\mathsf{A}}(\mathsf{PP}, M, R, \mathbb{A}, \bar{k})$ which takes one more parameter $\bar{k} \in \{1, \ldots, N+1\}$ than the original one in R-CP-ABE. This also changes the decryption criteria in AugR-CP-ABE in such a way that an encrypted message can be recovered using a decryption key $\mathsf{SK}_{k,S}$, which is identified by index $k \in \{1, \ldots, N\}$ and associated with an attribute set S, only if $(k \notin R) \wedge (S\ satisfies\ \mathbb{A}) \wedge (k \geq \bar{k})$. On the security, we formalize and show that a message-hiding and index-hiding AugR-CP-ABE can be transformed to a secure R-CP-ABE with policy-specific blackbox traceability.

In Sect. 4, we propose a *large universe* AugR-CP-ABE construction, and prove its message-hiding and index-hiding properties in the standard model. Combining it with the results in Sect. 3, we obtain a large universe R-CP-ABE construction, which is efficient (with overhead size in $O(\sqrt{N})$), highly expressive, and provably secure and traceable in the standard model.

To construct the AugR-CP-ABE, we borrow ideas from the CP-ABE constructions in [16,23] and Trace & Revoke scheme in [8]. However, the combination is not trivial and may result in inefficient or insecure systems. In particular, besides achieving the important features for practicality, such as traitor tracing, revocation, large universe, high expressivity and efficiency, we achieve provable security and traceability in the standard model. As we will discuss later in Sect. 4, proving the blackbox traceability while supporting the large attribute universe is one of the most challenging tasks in this work. As we can see, the proof techniques for blackbox traceability in [16] are no longer applicable for large universe, while that for large universe in [23] are only for confidentiality rather than for blackbox traceability.

2 Revocable CP-ABE and Blackbox Traceability

In this section, we define Revocable CP-ABE (or R-CP-ABE for short) and its security, which are based on conventional (non-traceable, non-revocable) CP-ABE (e.g. [15,23]). Similar to the traceable CP-ABE in [16], in our 'functional' definition, we explicitly assign and identify users using unique indices. Then we formalize traceability against policy-specific decryption blackbox on R-CP-ABE.

2.1 Revocable CP-ABE

Given a positive integer n, let $[n]$ be the set $\{1, 2, \ldots, n\}$. A Revocable CP-ABE (R-CP-ABE) scheme consists of four algorithms:

$\mathsf{Setup}(\lambda, N) \rightarrow (\mathsf{PP}, \mathsf{MSK})$. The algorithm takes as input a security parameter λ and the number of users in the system N, runs in polynomial time in λ, and outputs a public parameter PP and a master secret key MSK. We assume that PP contains the description of the attribute universe $\mathcal{U}$[3].

$\mathsf{KeyGen}(\mathsf{PP}, \mathsf{MSK}, S) \rightarrow \mathsf{SK}_{k,S}$. The algorithm takes as input PP, MSK, and an attribute set S, and outputs a secret key $\mathsf{SK}_{k,S}$ corresponding to S. The secret key is assigned and identified by a unique index $k \in [N]$.

$\mathsf{Encrypt}(\mathsf{PP}, M, R, \mathbb{A}) \rightarrow CT_{R,\mathbb{A}}$. The algorithm takes as input PP, a message M, a revocation list $R \subseteq [N]$, and an access policy $\mathbb{A}$ over $\mathcal{U}$, and outputs a ciphertext $CT_{R,\mathbb{A}}$. $(R, \mathbb{A})$ is included in $CT_{R,\mathbb{A}}$.

$\mathsf{Decrypt}(\mathsf{PP}, CT_{R,\mathbb{A}}, \mathsf{SK}_{k,S}) \rightarrow M$ or $\perp$. The algorithm takes as input PP, a ciphertext $CT_{R,\mathbb{A}}$, and a secret key $\mathsf{SK}_{k,S}$. If $(k \in [N] \setminus R)$ AND $(S\ satisfies\ \mathbb{A})$, the algorithm outputs a message M, otherwise it outputs $\perp$ indicating the failure of decryption.

Correctness. For any attribute set $S \subseteq \mathcal{U}$, index $k \in [N]$, revocation list $R \subseteq [N]$, access policy $\mathbb{A}$, and message M, suppose $(\mathsf{PP}, \mathsf{MSK}) \leftarrow \mathsf{Setup}(\lambda, N)$, $\mathsf{SK}_{k,S} \leftarrow \mathsf{KeyGen}(\mathsf{PP}, \mathsf{MSK}, S)$, $CT_{R,\mathbb{A}} \leftarrow \mathsf{Encrypt}(\mathsf{PP}, M, R, \mathbb{A})$. If $(k \in [N] \setminus R) \wedge (S$ satisfies $\mathbb{A})$ then $\mathsf{Decrypt}(\mathsf{PP}, CT_{R,\mathbb{A}}, \mathsf{SK}_{k,S}) = M$.

Security. The security of the R-CP-ABE is defined using the following message-hiding game, which is a typical semantic security game and is similar to that for conventional CP-ABE [15,23] security.

$\mathsf{Game}_{\mathsf{MH}}$. The message-hiding game is defined between a challenger and an adversary $\mathcal{A}$ as follows:

Setup. The challenger runs $\mathsf{Setup}(\lambda, N)$ and gives the public parameter PP to $\mathcal{A}$.

[3] For large universe and also in our work, the attribute universe depends only on the size of the underlying group $\mathbb{G}$, which depends on λ and the group generation algorithm.

Phase 1. For $i = 1$ to Q_1, $\mathcal{A}$ adaptively submits (index, attribute set) pair (k_i, S_{k_i}) to ask for secret key for attribute set S_{k_i}. For each (k_i, S_{k_i}) pair, the challenger responds with a secret key $\mathsf{SK}_{k_i,S_{k_i}}$, which corresponds to attribute set S_{k_i} and has index k_i.

Challenge. $\mathcal{A}$ submits two equal-length messages M_0, M_1 and a (revocation list, access policy) pair $(R^*, \mathbb{A}^*)$. The challenger flips a random coin $b \in \{0,1\}$, and sends $CT_{R^*,\mathbb{A}^*} \leftarrow \mathsf{Encrypt}(\mathsf{PP}, M_b, R^*, \mathbb{A}^*)$ to $\mathcal{A}$.

Phase 2. For $i = Q_1 + 1$ to Q, $\mathcal{A}$ adaptively submits (index, attribute set) pair (k_i, S_{k_i}) to ask for secret key for attribute set S_{k_i}. For each (k_i, S_{k_i}) pair, the challenger responds with a secret key $\mathsf{SK}_{k_i,S_{k_i}}$, which corresponds to attribute set S_{k_i} and has index k_i.

Guess. $\mathcal{A}$ outputs a guess $b' \in \{0,1\}$ for b.

$\mathcal{A}$ wins the game if $b' = b$ under the **restriction** that none of the queried $\{(k_i, S_{k_i})\}_{i=1}^{Q}$ can satisfy $(k_i \in [N] \setminus R^*)$ AND $(S_{k_i}\ satisfies\ \mathbb{A}^*)$. The advantage of $\mathcal{A}$ is defined as $\mathsf{MHAdv}_{\mathcal{A}} = |\Pr[b' = b] - \frac{1}{2}|$.

Definition 1. *An N-user R-CP-ABE scheme is secure if for all probabilistic polynomial time (PPT) adversaries $\mathcal{A}$, $\mathsf{MHAdv}_{\mathcal{A}}$ is negligible in λ.*

We say that an N-user R-CP-ABE scheme is *selectively* secure if we add an **Init** stage before **Setup** where the adversary commits to the access policy $\mathbb{A}^*$.

Remark: (1) Although the KeyGen algorithm is responsible for determining/assigning the index of each user's secret key, to capture the security that an adversary can adaptively choose secret keys to corrupt, the above model allows $\mathcal{A}$ to specify the index when querying for a key, i.e., for $i = 1$ to Q, $\mathcal{A}$ submits pairs of (k_i, S_{k_i}) for secret keys with attribute sets corresponding to S_{k_i}, and the challenger will assign k_i to be the index of the corresponding secret key, where $Q \leq N$, $k_i \in [N]$, and $k_i \neq k_j\ \forall 1 \leq i \neq j \leq Q$ (this is to guarantee that each user/key can be *uniquely* identified by an index). (2) For $k_i \neq k_j$ we do not require $S_{k_i} \neq S_{k_j}$, i.e., different users/keys may have the same attribute set.

Remark: (1) The R-CP-ABE defined above extends the conventional definition for non-revocable CP-ABE [15,16,23], where the revocation list R is always empty. (2) When the revocation list R needs an update due to, for example, some secret keys being compromised or some users leaving the system, the updated R needs to be disseminated to encrypting parties. In practice, this can be done in a similar way to the certificate revocation list distribution in the existing Public Key Infrastructure, namely an authority may update R, and publish it together with the authority's signature generated on it. (3) From the view of the public, R is just a set of numbers (in $[N]$). These numbers (or indices) do not have to provide any information on the corresponding users, in fact, besides the authority who runs KeyGen, each user only knows his/her own index. Also, encrypting parties do not need to know the indices of any users in order to encrypt but only the access policies. Although associating a revocation list with

a ciphertext might make the resulting CP-ABE look less purely attribute-based, it does not undermine the capability of CP-ABE, that is, enabling fine-grained access control on encrypted messages.

2.2 Blackbox Traceability

A policy-specific decryption blackbox $\mathcal{D}$ is described by a (revocation list, access policy) pair $(R_{\mathcal{D}}, \mathbb{A}_{\mathcal{D}})$ and a non-negligible probability value ϵ (i.e. $\epsilon = 1/f(\lambda)$ for some polynomial f), and this blackbox $\mathcal{D}$ can decrypt ciphertexts generated under $(R_{\mathcal{D}}, \mathbb{A}_{\mathcal{D}})$ with probability at least ϵ. Such a blackbox can reflect most practical scenarios, which include the key-like decryption blackbox for sale and decryption blackbox "found in the wild", which are discussed in [16,18]. In particular, once a blackbox is found being able to decrypt ciphertexts (regardless of how this is found, for example, an explicit description of the blackbox's decryption ability is given, or the law enforcement agency finds some clue), we can regard it as a policy-specific decryption blackbox with the corresponding (revocation list, access policy) pair (which is associated to the ciphertext).

We now define the tracing algorithm and traceability against policy-specific decryption blackbox.

$\mathsf{Trace}^{\mathcal{D}}(\mathsf{PP}, R_{\mathcal{D}}, \mathbb{A}_{\mathcal{D}}, \epsilon) \rightarrow \mathbb{K}_T \subseteq [N]$. Trace *is an oracle algorithm that interacts with a policy-specific decryption blackbox* $\mathcal{D}$*. By given the public parameter* PP*, a revocation list* $R_{\mathcal{D}}$*, an access policy* $\mathbb{A}_{\mathcal{D}}$*, and a probability value* ϵ*, the algorithm runs in time polynomial in* λ *and* $1/\epsilon$*, and outputs an index set* $\mathbb{K}_T \subseteq [N]$ *which identifies the set of malicious users. Note that* ϵ *has to be polynomially related to* λ*, i.e.* $\epsilon = 1/f(\lambda)$ *for some polynomial* f*.*

Traceability. The following tracing game captures the notion of **fully collusion-resistant traceability** against policy-specific decryption blackbox. In the game, the adversary targets to build a decryption blackbox $\mathcal{D}$ that can decrypt ciphertexts under some (revocation list, access policy) pair $(R_{\mathcal{D}}, \mathbb{A}_{\mathcal{D}})$.

$\mathsf{Game}_{\mathsf{TR}}$. The tracing game is defined between a challenger and an adversary $\mathcal{A}$ as follows:

Setup. The challenger runs $\mathsf{Setup}(\lambda, N)$ and gives the public parameter PP to $\mathcal{A}$.

Key Query. For $i = 1$ to Q, $\mathcal{A}$ adaptively submits (index, attribute set) pair (k_i, S_{k_i}) to ask for secret key for attribute set S_{k_i}. For each (k_i, S_{k_i}) pair, the challenger responds with a secret key $\mathsf{SK}_{k_i, S_{k_i}}$, which corresponds to attribute set S_{k_i} and has index k_i.

Decryption Blackbox Generation. $\mathcal{A}$ outputs a decryption blackbox $\mathcal{D}$ associated with a (revocation list, access policy) pair $(R_{\mathcal{D}}, \mathbb{A}_{\mathcal{D}})$ and a non-negligible probability value ϵ.

Tracing. The challenger runs $\mathsf{Trace}^{\mathcal{D}}(\mathsf{PP}, R_{\mathcal{D}}, \mathbb{A}_{\mathcal{D}}, \epsilon)$ to obtain an index set $\mathbb{K}_T \subseteq [N]$.

Let $\mathbb{K}_{\mathcal{D}} = \{k_i | 1 \leq i \leq Q\}$ be the index set of secret keys corrupted by the adversary. We say that $\mathcal{A}$ wins the game if the following two conditions hold:

1. $\Pr[\mathcal{D}(\mathsf{Encrypt}(\mathsf{PP}, M, R_{\mathcal{D}}, \mathbb{A}_{\mathcal{D}})) = M] \geq \epsilon$, where the probability is taken over the random choices of message M and the random coins of $\mathcal{D}$. A decryption blackbox satisfying this condition is said to be a *useful policy-specific decryption blackbox*.
2. $\mathbb{K}_T = \emptyset$, or $\mathbb{K}_T \not\subseteq \mathbb{K}_{\mathcal{D}}$, or $((k_t \in R_{\mathcal{D}})\ \mathsf{OR}\ (S_{k_t}\ does\ not\ satisfy\ \mathbb{A}_{\mathcal{D}})\ \forall k_t \in \mathbb{K}_T)$.

We denote by $\mathsf{TRAdv}_{\mathcal{A}}$ the probability that $\mathcal{A}$ wins.

Remark: For a useful policy-specific decryption blackbox $\mathcal{D}$, the traced $\mathbb{K}_T$ must satisfy $(\mathbb{K}_T \neq \emptyset) \wedge (\mathbb{K}_T \subseteq \mathbb{K}_{\mathcal{D}}) \wedge (\exists k_t \in \mathbb{K}_T\ s.t.\ (k_t \in [N] \backslash R_{\mathcal{D}})\ \mathsf{AND}\ (S_{k_t}\ satisfies\ \mathbb{A}_{\mathcal{D}}))$ for traceability. (1) $(\mathbb{K}_T \neq \emptyset) \wedge (\mathbb{K}_T \subseteq \mathbb{K}_{\mathcal{D}})$ captures the preliminary traceability that the tracing algorithm can extract at least one malicious user and the coalition of malicious users cannot frame any innocent user. (2) $(\exists k_t \in \mathbb{K}_T\ s.t.\ (k_t \in [N] \setminus R_{\mathcal{D}})\ \mathsf{AND}\ (S_{k_t}\ satisfies\ \mathbb{A}_{\mathcal{D}}))$ captures the *strong traceability* that the tracing algorithm can extract at least one malicious user whose secret key enables $\mathcal{D}$ to have the decryption ability corresponding to $(R_{\mathcal{D}}, \mathbb{A}_{\mathcal{D}})$, i.e. whose index is not in $R_{\mathcal{D}}$ and whose attribute set satisfies $\mathbb{A}_{\mathcal{D}}$. We refer to [12,16] on why strong traceability is desirable.

Note that, as of [4,5,8,12,16], we are modeling a stateless (resettable) decryption blackbox – such a blackbox is just an oracle and maintains no state between activations. Also note that we are modeling public traceability, namely, the Trace algorithm does not need any secrets and anyone can perform the tracing.

Definition 2. *An N-user R-CP-ABE scheme is traceable against policy-specific decryption blackbox if for all PPT adversaries $\mathcal{A}$, $\mathsf{TRAdv}_{\mathcal{A}}$ is negligible in λ.*

We say that an N-user R-CP-ABE is *selectively* traceable against policy-specific decryption blackbox if we add an **Init** stage before **Setup** where the adversary commits to the access policy $\mathbb{A}_{\mathcal{D}}$.

In the traceable CP-ABE of [16], given a decryption blackbox, it is guaranteed that at least one secret key in the blackbox will be traced. But in the traceable R-CP-ABE above, it is possible to trace *all the active secret keys* in the blackbox. In particular, given a decryption blackbox $\mathcal{D}$ described by $(R_{\mathcal{D}}, \mathbb{A}_{\mathcal{D}})$ and non-negligible probability ϵ, we can run Trace to obtain an index set $\mathbb{K}_T$ so that $(\mathbb{K}_T \neq \emptyset) \wedge (\mathbb{K}_T \subseteq \mathbb{K}_{\mathcal{D}}) \wedge (\exists k_t \in \mathbb{K}_T\ s.t.\ (k_t \in [N] \backslash R_{\mathcal{D}})\ \mathsf{AND}\ (S_{k_t}\ satisfies\ \mathbb{A}_{\mathcal{D}}))$. Then, we can set a new revocation list $R'_{\mathcal{D}} = R_{\mathcal{D}} \cup \{k_t \in \mathbb{K}_T \mid (k_t \in [N] \setminus R_{\mathcal{D}})\ \mathsf{AND}\ (S_{k_t}\ satisfies\ \mathbb{A}_{\mathcal{D}})\}$ and test whether $\mathcal{D}$ can decrypt ciphertexts under $(R'_{\mathcal{D}}, \mathbb{A}_{\mathcal{D}})$. If $\mathcal{D}$ can still decrypt the ciphertexts with non-negligible probability ϵ', we can run Trace on $(R'_{\mathcal{D}}, \mathbb{A}_{\mathcal{D}}, \epsilon')$ and obtain a new index set $\mathbb{K}'_T$, where $(\mathbb{K}'_T \neq \emptyset) \wedge (\mathbb{K}'_T \subseteq \mathbb{K}_{\mathcal{D}}) \wedge (\exists k_t \in \mathbb{K}'_T\ s.t.\ (k_t \in [N] \backslash R'_{\mathcal{D}})\ \mathsf{AND}\ (S_{k_t}\ satisfies\ \mathbb{A}_{\mathcal{D}}))$. By repeating this process, iteratively expanding the revocation list, until $\mathcal{D}$ can no longer decrypt the corresponding ciphertexts, we have finished finding out *all the active* malicious users of $\mathcal{D}$.

3 Augmented R-CP-ABE

As outlined in Sect. 1.1, we now define Augmented R-CP-ABE (AugR-CP-ABE) from the R-CP-ABE above, and formalize its security notions, then show that a

secure AugR-CP-ABE can be transformed to a secure R-CP-ABE with blackbox traceability. In Sect. 4, we propose a concrete construction of AugR-CP-ABE.

3.1 Definitions

An AugR-CP-ABE scheme has four algorithms: $\mathsf{Setup_A}$, $\mathsf{KeyGen_A}$, $\mathsf{Encrypt_A}$, and $\mathsf{Decrypt_A}$. The setup and key generation algorithms are the same as that of R-CP-ABE. For the encryption algorithm, it takes one more parameter $\bar{k} \in [N+1]$ as input, and is defined as follows.

$\mathsf{Encrypt_A}(\mathsf{PP}, M, R, \mathbb{A}, \bar{k}) \rightarrow CT_{R,\mathbb{A}}$. The algorithm takes as input PP, a message M, a revocation list $R \subseteq [N]$, an access policy $\mathbb{A}$, and an index $\bar{k} \in [N+1]$, and outputs a ciphertext $CT_{R,\mathbb{A}}$. $(R, \mathbb{A})$ **is included in** $CT_{R,\mathbb{A}}$**, but the value of** $\bar{k}$ **is not.**

The decryption algorithm is also defined in the same way as that of R-CP-ABE. However, the correctness definition is changed to the following.

Correctness. For any attribute set $S \subseteq \mathcal{U}$, index $k \in [N]$, revocation list $R \subseteq [N]$, access policy $\mathbb{A}$ over $\mathcal{U}$, encryption index $\bar{k} \in [N+1]$, and message M, suppose $(\mathsf{PP}, \mathsf{MSK}) \leftarrow \mathsf{Setup_A}(\lambda, N)$, $\mathsf{SK}_{k,S} \leftarrow \mathsf{KeyGen_A}(\mathsf{PP}, \mathsf{MSK}, S)$, $CT_{R,\mathbb{A}} \leftarrow \mathsf{Encrypt_A}(\mathsf{PP}, M, R, \mathbb{A}, \bar{k})$. If $(k \in [N] \setminus R) \wedge (S \text{ satisfies } \mathbb{A}) \wedge (k \geq \bar{k})$ then $\mathsf{Decrypt_A}(\mathsf{PP}, CT_{R,\mathbb{A}}, \mathsf{SK}_{k,S}) = M$.

Note that during decryption, as long as $(k \in [N] \setminus R) \wedge (S \text{ satisfies } \mathbb{A})$, the decryption algorithm outputs a message, but only when $k \geq \bar{k}$, the output message is equal to the correct message, that is, if and only if $(k \in [N] \setminus R) \wedge (S \text{ satisfies } \mathbb{A}) \wedge (k \geq \bar{k})$, can $\mathsf{SK}_{k,S}$ correctly decrypt a ciphertext under $(R, \mathbb{A}, \bar{k})$. If we always set $\bar{k} = 1$, the functions of AugR-CP-ABE are identical to that of R-CP-ABE. In fact, the idea behind transforming an AugR-CP-ABE to a traceable R-CP-ABE, that we will show shortly, is to construct an AugR-CP-ABE with index-hiding property, and then always sets $\bar{k} = 1$ in normal encryption, while using $\bar{k} \in [N+1]$ to generate ciphertexts for tracing.

Security. We define the security of AugR-CP-ABE in two games. The first game is a **message-hiding game** and says that a ciphertext created using index $N+1$ is unreadable by anyone. The second game is an **index-hiding game** and captures the intuition that a ciphertext created using index $\bar{k}$ reveals no non-trivial information about $\bar{k}$.

$\mathsf{Game}^{\mathsf{A}}_{\mathsf{MH}}$. The **message-hiding game** $\mathsf{Game}^{\mathsf{A}}_{\mathsf{MH}}$ is similar to $\mathsf{Game}_{\mathsf{MH}}$ except that the **Challenge** phase is

Challenge. $\mathcal{A}$ submits two equal-length messages M_0, M_1 and a (revocation list, access policy) pair $(R^*, \mathbb{A}^*)$. The challenger flips a random coin $b \in \{0,1\}$, and sends $CT_{R^*,\mathbb{A}^*} \leftarrow \mathsf{Encrypt_A}(\mathsf{PP}, M_b, R^*, \mathbb{A}^*, N+1)$ to $\mathcal{A}$.

$\mathcal{A}$ wins the game if $b' = b$. The advantage of $\mathcal{A}$ is defined as $\mathsf{MH}^{\mathsf{A}}\mathsf{Adv}_{\mathcal{A}} = |\Pr[b' = b] - \frac{1}{2}|$.

Definition 3. *An N-user Augmented R-CP-ABE scheme is message-hiding if for all PPT adversaries $\mathcal{A}$ the advantage $\mathsf{MH}^{\mathsf{A}}\mathsf{Adv}_{\mathcal{A}}$ is negligible in λ.*

$\mathsf{Game}^{\mathsf{A}}_{\mathsf{IH}}$. In the **index-hiding game**, we require that, for any (revocation list, access policy) pair $(R^*, \mathbb{A}^*)$, an adversary cannot distinguish between a ciphertext under $(R^*, \mathbb{A}^*, \bar{k})$ and $(R^*, \mathbb{A}^*, \bar{k}+1)$ without a secret key $\mathsf{SK}_{\bar{k},S_{\bar{k}}}$ such that $(\bar{k} \in [N] \setminus R^*) \wedge (S_{\bar{k}}\ satisfies\ \mathbb{A}^*)$. The game takes as input a parameter $\bar{k} \in [N]$ which is given to both the challenger and the adversary. The game is similar to $\mathsf{Game}_{\mathsf{MH}}$ except that the **Challenge** phase is

Challenge. $\mathcal{A}$ submits a message M and a (revocation list, access policy) pair $(R^*, \mathbb{A}^*)$. The challenger flips a random coin $b \in \{0,1\}$, and sends $CT_{R^*,\mathbb{A}^*} \leftarrow \mathsf{Encrypt}_{\mathsf{A}}(\mathsf{PP}, M, R^*, \mathbb{A}^*, \bar{k}+b)$ to $\mathcal{A}$.

$\mathcal{A}$ wins the game if $b' = b$ under the **restriction** that none of the queried pairs $\{(k_i, S_{k_i})\}_{i=1}^{Q}$ can satisfy $(k_i = \bar{k}) \wedge (k_i \in [N] \setminus R^*) \wedge (S_{k_i}\ satisfies\ \mathbb{A}^*)$. The advantage of $\mathcal{A}$ is defined as $\mathsf{IH}^{\mathsf{A}}\mathsf{Adv}_{\mathcal{A}}[\bar{k}] = |\Pr[b' = b] - \frac{1}{2}|$.

Definition 4. *An N-user Augmented R-CP-ABE scheme is index-hiding if for all PPT adversaries $\mathcal{A}$ the advantages $\mathsf{IH}^{\mathsf{A}}\mathsf{Adv}_{\mathcal{A}}[\bar{k}]$ for $\bar{k} = 1, \ldots, N$ are negligible in λ.*

We say that an Augmented R-CP-ABE scheme is *selectively* index-hiding if we add an **Init** stage before **Setup** where the adversary commits to the challenge access policy $\mathbb{A}^*$.

3.2 The Reduction of Traceable R-CP-ABE to Augmented R-CP-ABE

Let $\Sigma_{\mathsf{A}} = (\mathsf{Setup}_{\mathsf{A}}, \mathsf{KeyGen}_{\mathsf{A}}, \mathsf{Encrypt}_{\mathsf{A}}, \mathsf{Decrypt}_{\mathsf{A}})$ be an AugR-CP-ABE, define $\mathsf{Encrypt}(\mathsf{PP}, M, R, \mathbb{A}) = \mathsf{Encrypt}_{\mathsf{A}}(\mathsf{PP}, M, R, \mathbb{A}, 1)$, then $\Sigma = (\mathsf{Setup}_{\mathsf{A}}, \mathsf{KeyGen}_{\mathsf{A}}, \mathsf{Encrypt}, \mathsf{Decrypt}_{\mathsf{A}})$ is a R-CP-ABE derived from Σ_{A}. In the following, we show that if Σ_{A} is message-hiding and index-hiding, then Σ is secure (w.r.t. Definition 1). Furthermore, we propose a tracing algorithm Trace for Σ and show that if Σ_{A} is message-hiding and index-hiding, then Σ (equipped with Trace) is traceable (w.r.t. Definition 2).

Theorem 1. *If Σ_{A} is message-hiding and index-hiding (resp. selectively index-hiding), then Σ is secure (resp. selectively secure).*

Proof. Due to page limitation, the proof details are omitted here and can be found in the full version [19].

We now propose a tracing algorithm Trace, which uses a general tracing method previously used in [3–5,8,16,20], and show that equipped with Trace, Σ is traceable (w.r.t. Def. 2).

$\mathsf{Trace}^{\mathcal{D}}(\mathsf{PP}, R_{\mathcal{D}}, \mathbb{A}_{\mathcal{D}}, \epsilon) \rightarrow \mathbb{K}_T \subseteq [N]$: Given a policy-specific decryption blackbox $\mathcal{D}$ associated with a (revocation list, access policy) pair $(R_{\mathcal{D}}, \mathbb{A}_{\mathcal{D}})$ and probability $\epsilon > 0$, the tracing algorithm works as follows:

1. For $k = 1$ to $N + 1$, do the following:
   (a) Repeat the following $8\lambda(N/\epsilon)^2$ times:
      i Sample M from the message space at random.
      ii Let $CT_{R_{\mathcal{D}},\mathbb{A}_{\mathcal{D}}} \leftarrow \mathsf{Encrypt}_{\mathsf{A}}(\mathsf{PP}, M, R_{\mathcal{D}}, \mathbb{A}_{\mathcal{D}}, k)$.
      iii Query oracle $\mathcal{D}$ on input $CT_{R_{\mathcal{D}},\mathbb{A}_{\mathcal{D}}}$, and compare the output of $\mathcal{D}$ with M.
   (b) Let $\hat{p}_k$ be the fraction of times that $\mathcal{D}$ decrypted the ciphertexts correctly.
2. Let $\mathbb{K}_T$ be the set of all $k \in [N]$ for which $\hat{p}_k - \hat{p}_{k+1} \geq \epsilon/(4N)$. Output $\mathbb{K}_T$.

Theorem 2. *If Σ_{A} is message-hiding and index-hiding (resp. selectively index-hiding), then Σ is traceable (resp. selectively traceable).*

Proof. Due to page limitation, the proof details are omitted here and can be found in the full version [19].

4 An Efficient Augmented R-CP-ABE

We propose an AugR-CP-ABE scheme which is highly expressive and efficient with sub-linear overhead in the number of users in the system. It is also *large universe*, where attributes do not need to be enumerated during setup, and the public parameter size is independent of the attribute universe size. We prove that this AugR-CP-ABE scheme is message-hiding and selectively index-hiding in the standard model.

Combining this AugR-CP-ABE with the results in Sect. 3.2, we obtain a large universe R-CP-ABE which is selectively secure and traceable, and for a fully collusion-resistant blackbox traceable CP-ABE, the resulting R-CP-ABE achieves the most efficient level to date, with sub-linear overhead.

To obtain this practical CP-ABE scheme supporting traitor tracing, revocation and large universe, we borrow ideas from the Blackbox Traceable CP-ABE of [16], the Trace and Revoke scheme of [8] and the Large Universe CP-ABE of [23], but the work is not trivial as a straightforward combination of the ideas would result in a scheme which is inefficient, insecure, or is not able to achieve strong traceability. Specifically, by incorporating the ideas from [8,23] into the Augmented CP-ABE of [16], we can obtain a large universe AugR-CP-ABE which is message-hiding, but proving the index-hiding property is a challenging task. The proof techniques for index-hiding in [16] only work if the attribute universe size is polynomial in the security parameter and the parameters of attributes have to be enumerated during setup. They are not applicable to large universe. The proof techniques in [23] are applicable to large universe, but work only for message-hiding, while not applicable to index-hiding. To prove index-hiding in the large universe setting, we introduce a new assumption that the index-hiding of our large universe AugR-CP-ABE can be based on. In particular, in the underlying q-1 assumption of [23] on bilinear groups $(p, \mathbb{G}, \mathbb{G}_T, e)$, the challenge term $T \in \mathbb{G}_T$ is $e(g,g)^{ca^{q+1}}$ or a random element, and such a term in the target group could be used to prove the message-hiding as the message space is $\mathbb{G}_T$. To prove the index-hiding, which is based on the ciphertext components in the source group $\mathbb{G}$, we need the challenge term to be in the source group $\mathbb{G}$ so that

the simulator can embed the challenge term into these ciphertext components. Inspired by the Source Group q-Parallel BDHE Assumption in [15], which is a close relative to the (target group) Decisional Parallel BDHE Assumption in [26], we modify the q-1 assumption to its source group version where the challenge term is $g^{ca^{q+1}}$ or a random element in $\mathbb{G}$. Based on this new assumption and with a new crucial proof idea, we prove the index-hiding property for our large universe AugR-CP-ABE. We prove that this new assumption holds in the generic group model.

4.1 Preliminaries

Linear Secret-Sharing Schemes (LSSS). An LSSS is a share-generating matrix A whose rows labeled by attributes via a function ρ. An attribute set S satisfies the LSSS access matrix (A, ρ) if the rows labeled by the attributes in S have the *linear reconstruction* property, namely, there exist constants $\{\omega_i\}$ such that, for any valid shares $\{\lambda_i\}$ of a secret s, we have $\sum_i \omega_i \lambda_i = s$. The formal definitions of access structures and LSSS can be found in the full version [19].

Bilinear Groups. Let $\mathcal{G}$ be a group generator, which takes a security parameter λ and outputs $(p, \mathbb{G}, \mathbb{G}_T, e)$ where p is a prime, $\mathbb{G}$ and $\mathbb{G}_T$ are cyclic groups of order p, and $e : \mathbb{G} \times \mathbb{G} \to \mathbb{G}_T$ is a map such that: (1) (Bilinear) $\forall g, h \in \mathbb{G}, a, b \in \mathbb{Z}_p, e(g^a, h^b) = e(g, h)^{ab}$, (2) (Non-Degenerate) $\exists g \in \mathbb{G}$ such that $e(g, g)$ has order p in $\mathbb{G}_T$. We refer to $\mathbb{G}$ as the *source group* and $\mathbb{G}_T$ as the *target group*. We assume that group operations in $\mathbb{G}$ and $\mathbb{G}_T$ as well as the bilinear map e are efficiently computable, and the description of $\mathbb{G}$ and $\mathbb{G}_T$ includes a generator of $\mathbb{G}$ and $\mathbb{G}_T$ respectively.

Complexity Assumptions. Besides the Decision 3-Party Diffie-Hellman s Assumption (D3DH) and the Decisional Linear Assumption (DLIN) that are used in [8] to achieve traceability in broadcast encryption, the index-hiding property of our AugR-CP-ABE construction will rely on a new assumption, which is similar to the Source Group q-Parallel BDHE Assumption [15] and is closely related to the q-1 assumption in [23]. We refer to it as the Extended Source Group q-Parallel BDHE Assumption. Here we only review this new assumption, and refer to the full version [19] for the details of the D3DH and DLIN.

The Extended Source Group q-Parallel BDHE Assumption *Given a group generator $\mathcal{G}$ and a positive integer q, define the following distribution:*

$$
\begin{aligned}
&(p, \mathbb{G}, \mathbb{G}_T, e) \xleftarrow{R} \mathcal{G}(\lambda), \quad g \xleftarrow{R} \mathbb{G}, \ a, c, d, b_1, \ldots, b_q \xleftarrow{R} \mathbb{Z}_p, \\
&D = \big((p, \mathbb{G}, \mathbb{G}_T, e), \ g, g^{cd}, g^{d}, g^{da^q}, \\
&\qquad g^{a^i}, \ g^{b_j}, \ g^{a^i b_j}, \ g^{a^i / b_j^2}, \ g^{cdb_j} \quad \forall i, j \in [q], \\
&\qquad g^{a^i / b_j} \qquad\qquad \forall i \in [2q] \setminus \{q+1\}, j \in [q], \\
&\qquad g^{a^i b_{j'} / b_j^2} \qquad\qquad \forall i \in [2q], j, j' \in [q] \ s.t. \ j' \neq j, \\
&\qquad g^{cda^i b_{j'} / b_j}, g^{cda^i b_{j'} / b_j^2} \qquad \forall i \in [q], j, j' \in [q] \ s.t. \ j \neq j' \ \big), \\
&T_0 = g^{ca^{q+1}}, T_1 \xleftarrow{R} \mathbb{G}.
\end{aligned}
$$

The advantage of an algorithm $\mathcal{A}$ in breaking the Extended Source Group q-Parallel BDHE Assumption is $Adv^q_{\mathcal{G},\mathcal{A}}(\lambda) := |\Pr[\mathcal{A}(D, T_0) = 1] - \Pr[\mathcal{A}(D, T_1) = 1]|$.

Definition 5. *$\mathcal{G}$ satisfies the Extended Source Group q-Parallel BDHE Assumption if $Adv^q_{\mathcal{G},\mathcal{A}}(\lambda)$ is a negligible function of λ for any PPT algorithm $\mathcal{A}$.*

This new assumption is closely related to the q-1 assumption in [23], except that the challenge term $g^{ca^{q+1}}$ remains in the source group, all the input terms (in D) replace c with cd, and additional input terms g^d and g^{da^q} are given to the adversary. The relation between this assumption and the q-1 assumption is analogous to that between the Source Group q-Parallel BDHE Assumption [15] and the Decisional Parallel BDHE Assumption [26], i.e. the challenge term changes from a term in the target group (i.e. $e(g,g)^{ca^{q+1}}$) to a term in the source group (i.e. $g^{ca^{q+1}}$), and the input terms are modified accordingly (i.e. replacing c with cd, and adding g^d). The main difference is that in this new assumption, there is an additional input term g^{da^q}. Note that giving the term g^{da^q} does not pose any problem in the generic group model. Intuitively, there are two ways that the adversary may make use of the term g^{da^q}: (1) pairing g^{da^q} with the challenge term: since the pairing result of any two input terms would not be $e(g,g)^{cda^{2q+1}}$, the adversary cannot break this new assumption in this way; (2) pairing the challenge term with another input term whose exponent contains d: however, the result could be a random element or one of $\{\ e(g,g)^{c^2da^{q+1}}$, $e(g,g)^{cda^{q+1}}$, $e(g,g)^{c^2db_ja^{q+1}}$, $e(g,g)^{c^2da^{q+1+i}b_{j'}/b_j}$, $e(g,g)^{c^2da^{q+1+i}b_{j'}/b_j^2}\}$, and as there is no input term which can be paired with g^{da^q} to obtain any of these terms, the adversary cannot break this new assumption by this way either. In the full version [19] we prove that this assumption holds in the generic group model. It is worth mentioning that Liu et al. [18] modified the Source Group q-Parallel BDHE Assumption [15] by adding g^{da^q} to and removing $g^{a^{q+2}}, \ldots, g^{a^{2q}}$ from the input terms.

Notations. Suppose that the number of users N in the system equals to m^2 for some m. In practice, if N is not a square, we can add some "dummy" users until it pads to the next square. We arrange the users in an $m \times m$ matrix and uniquely assign a tuple (i,j), where $i,j \in [m]$, to each user. A user at position (i,j) of the matrix has index $k = (i-1) * m + j$. For simplicity, we directly use (i,j) as the index where $(i,j) \geq (\bar{i},\bar{j})$ means that $((i > \bar{i}) \vee (i = \bar{i} \wedge j \geq \bar{j}))$. Let $[m,m]$ be the set $\{(i,j)|i,j \in [m]\}$. The use of pairwise notation (i,j) is purely a notational convenience, as $k = (i-1)*m+j$ defines a bijection between $\{(i,j)|i,j \in [m]\}$ and $[N]$. For a given vector $\boldsymbol{v} = (v_1, \ldots, v_d)$, by $g^{\boldsymbol{v}}$ we mean the vector $(g^{v_1}, \ldots, g^{v_d})$. Furthermore, for $g^{\boldsymbol{v}} = (g^{v_1}, \ldots, g^{v_d})$ and $g^{\boldsymbol{w}} = (g^{w_1}, \ldots, g^{w_d})$, by $g^{\boldsymbol{v}} \cdot g^{\boldsymbol{w}}$ we mean the vector $(g^{v_1+w_1}, \ldots, g^{v_d+w_d})$, i.e. $g^{\boldsymbol{v}} \cdot g^{\boldsymbol{w}} = g^{\boldsymbol{v}+\boldsymbol{w}}$, and by $e_d(g^{\boldsymbol{v}}, g^{\boldsymbol{w}})$ we mean $\prod_{k=1}^{d} e(g^{v_k}, g^{w_k})$, i.e. $e_d(g^{\boldsymbol{v}}, g^{\boldsymbol{w}}) = e(g,g)^{(\boldsymbol{v}\cdot\boldsymbol{w})}$, where $(\boldsymbol{v} \cdot \boldsymbol{w})$ is the inner product of $\boldsymbol{v}$ and $\boldsymbol{w}$. Given a prime p, one can randomly choose $r_x, r_y, r_z \in \mathbb{Z}_p$, and set $\boldsymbol{\chi}_1 = (r_x, 0, r_z)$, $\boldsymbol{\chi}_2 = (0, r_y, r_z)$, $\boldsymbol{\chi}_3 = \boldsymbol{\chi}_1 \times \boldsymbol{\chi}_2 = (-r_y r_z, -r_x r_z, r_x r_y)$. Let $span\{\boldsymbol{\chi}_1, \boldsymbol{\chi}_2\} = \{\nu_1\boldsymbol{\chi}_1 + \nu_2\boldsymbol{\chi}_2 | \nu_1, \nu_2 \in \mathbb{Z}_p\}$ be the subspace spanned by $\boldsymbol{\chi}_1$ and $\boldsymbol{\chi}_2$. We can see that $\boldsymbol{\chi}_3$ is orthogonal to the subspace $span\{\boldsymbol{\chi}_1, \boldsymbol{\chi}_2\}$

and $\mathbb{Z}_p^3 = span\{\boldsymbol{\chi}_1, \boldsymbol{\chi}_2, \boldsymbol{\chi}_3\} = \{\nu_1\boldsymbol{\chi}_1 + \nu_2\boldsymbol{\chi}_2 + \nu_3\boldsymbol{\chi}_3 | \nu_1, \nu_2, \nu_3 \in \mathbb{Z}_p\}$. For any $\boldsymbol{v} \in span\{\boldsymbol{\chi}_1, \boldsymbol{\chi}_2\}$, $(\boldsymbol{\chi}_3 \cdot \boldsymbol{v}) = 0$, and for random $\boldsymbol{v} \in \mathbb{Z}_p^3$, $(\boldsymbol{\chi}_3 \cdot \boldsymbol{v}) \neq 0$ happens with overwhelming probability.

4.2 Augmented R-CP-ABE Construction

Now we propose a large universe Augmented R-CP-ABE, where the attribute universe is $\mathcal{U} = \mathbb{Z}_p$.

$\mathsf{Setup}_\mathsf{A}(\lambda, N = m^2) \rightarrow (\mathsf{PP}, \mathsf{MSK})$. The algorithm calls the group generator $\mathcal{G}(\lambda)$ to get $(p, \mathbb{G}, \mathbb{G}_T, e)$, and sets the attribute universe to $\mathcal{U} = \mathbb{Z}_p$. It then randomly picks $g, h, f, f_1, \ldots, f_m, G, H \in \mathbb{G}$, $\{\alpha_i, r_i, z_i \in \mathbb{Z}_p\}_{i\in[m]}$, $\{c_j \in \mathbb{Z}_p\}_{j\in[m]}$, and outputs the public parameter PP and master secret key MSK

$$\mathsf{PP} = \Big((p, \mathbb{G}, \mathbb{G}_T, e),\ g,\ h,\ f,\ f_1, \ldots, f_m,\ G,\ H, \\ \{E_i = e(g,g)^{\alpha_i}, G_i = g^{r_i},\ Z_i = g^{z_i}\}_{i\in[m]},\ \ \{H_j = g^{c_j}\}_{j\in[m]} \Big),$$

$$\mathsf{MSK} = \Big(\alpha_1, \ldots, \alpha_m,\ \ r_1, \ldots, r_m,\ \ c_1, \ldots, c_m \Big).$$

A counter $ctr = 0$ is implicitly included in MSK.

$\mathsf{KeyGen}_\mathsf{A}(\mathsf{PP}, \mathsf{MSK}, S \subseteq \mathbb{Z}_p) \rightarrow \mathsf{SK}_{(i,j),S}$. The algorithm first sets $ctr = ctr + 1$ and computes the corresponding index in the form of (i, j) where $1 \leq i, j \leq m$ and $(i-1) * m + j = ctr$. Then it picks random exponents $\sigma_{i,j} \in \mathbb{Z}_p$, $\{\delta_{i,j,x} \in \mathbb{Z}_p\}_{\forall x\in S}$, and outputs a secret key $\mathsf{SK}_{(i,j),S} = \big((i,j), S,\ K_{i,j}, K'_{i,j}, K''_{i,j}, \{\bar{K}_{i,j,j'}\}_{j'\in[m]\setminus\{j\}}, \{K_{i,j,x}, K'_{i,j,x}\}_{x\in S}\big)$ where

$$K_{i,j} = g^{\alpha_i} g^{r_i c_j} (f f_j)^{\sigma_{i,j}},\ K'_{i,j} = g^{\sigma_{i,j}},\ K''_{i,j} = Z_i^{\sigma_{i,j}}, \\ \{\bar{K}_{i,j,j'} = f_{j'}^{\sigma_{i,j}}\}_{j'\in[m]\setminus\{j\}},\ \{K_{i,j,x} = g^{\delta_{i,j,x}},\ K'_{i,j,x} = (H^x h)^{\delta_{i,j,x}} G^{-\sigma_{i,j}}\}_{x\in S}.$$

$\mathsf{Encrypt}_\mathsf{A}(\mathsf{PP}, M, R, \mathbb{A} = (A, \rho), (\bar{i}, \bar{j})) \rightarrow CT_{R,(A,\rho)}$. $R \subseteq [m, m]$ is a revocation list. $\mathbb{A} = (A, \rho)$ is an LSSS matrix where A is an $l \times n$ matrix and ρ maps each row A_k of A to an attribute $\rho(k) \in \mathcal{U} = \mathbb{Z}_p$. The encryption is for recipients whose (index, attribute set) pairs $((i,j), S_{(i,j)})$ satisfy $((i,j) \in [m,m] \setminus R) \wedge \big(S_{(i,j)}\ satisfies\ (A,\rho)\big) \wedge ((i,j) \geq (\bar{i}, \bar{j}))$. Let $\bar{R} = [m,m] \setminus R$ and for $i \in [m]$, $\bar{R}_i = \{j' | (i,j') \in \bar{R}\}$, that is, $\bar{R}$ is the non-revoked index list, and $\bar{R}_i$ is the set of non-revoked column index on the i-th row. The algorithm randomly chooses $\kappa, \tau, s_1, \ldots, s_m, t_1, \ldots, t_m \in \mathbb{Z}_p$, $\boldsymbol{v}_c, \boldsymbol{w}_1, \ldots, \boldsymbol{w}_m \in \mathbb{Z}_p^3$, $\xi_1, \ldots, \xi_l \in \mathbb{Z}_p$, and $\boldsymbol{u} = (\pi, u_2, \ldots, u_n) \in \mathbb{Z}_p^n$. In addition, it randomly chooses $r_x, r_y, r_z \in \mathbb{Z}_p$, and sets $\boldsymbol{\chi}_1 = (r_x, 0, r_z)$, $\boldsymbol{\chi}_2 = (0, r_y, r_z)$, $\boldsymbol{\chi}_3 = \boldsymbol{\chi}_1 \times \boldsymbol{\chi}_2 = (-r_y r_z, -r_x r_z, r_x r_y)$. Then it randomly chooses $\boldsymbol{v}_i \in \mathbb{Z}_p^3\ \forall i \in \{1, \ldots, \bar{i}\}$, $\boldsymbol{v}_i \in span\{\boldsymbol{\chi}_1, \boldsymbol{\chi}_2\}\ \forall i \in \{\bar{i}+1, \ldots, m\}$, and computes a ciphertext $\langle R, (A,\rho),\ (\boldsymbol{R}_i, \boldsymbol{R}'_i, Q_i, Q'_i,\ Q''_i,\ T_i)_{i=1}^m, (\boldsymbol{C}_j, \boldsymbol{C}'_j)_{j=1}^m, (P_k, P'_k, P''_k)_{k=1}^l \rangle$ as follows:

1. For each row $i \in [m]$:
   - if $i < \bar{i}$: randomly chooses $\hat{s}_i \in \mathbb{Z}_p$, and sets
   $$\boldsymbol{R}_i = g^{\boldsymbol{v}_i}, \quad \boldsymbol{R}'_i = g^{\kappa \boldsymbol{v}_i},$$
   $$Q_i = g^{s_i}, \quad Q'_i = (f \prod_{j' \in \bar{R}_i} f_{j'})^{s_i} Z_i^{t_i} f^{\pi}, \quad Q''_i = g^{t_i}, \quad T_i = E_i^{\hat{s}_i}.$$
   - if $i \geq \bar{i}$: sets
   $$\boldsymbol{R}_i = G_i^{s_i \boldsymbol{v}_i}, \quad \boldsymbol{R}'_i = G_i^{\kappa s_i \boldsymbol{v}_i}, Q_i = g^{\tau s_i (\boldsymbol{v}_i \cdot \boldsymbol{v}_c)},$$
   $$Q'_i = (f \prod_{j' \in \bar{R}_i} f_{j'})^{\tau s_i (\boldsymbol{v}_i \cdot \boldsymbol{v}_c)} Z_i^{t_i} f^{\pi}, Q''_i = g^{t_i}, \quad T_i = M \cdot E_i^{\tau s_i (\boldsymbol{v}_i \cdot \boldsymbol{v}_c)}.$$
2. For each column $j \in [m]$:
   - if $j < \bar{j}$: randomly chooses $\mu_j \in \mathbb{Z}_p$, and sets $\boldsymbol{C}_j = H_j^{\tau(\boldsymbol{v}_c + \mu_j \boldsymbol{\chi}_3)} \cdot g^{\kappa \boldsymbol{w}_j}, \boldsymbol{C}'_j = g^{\boldsymbol{w}_j}$.
   - if $j \geq \bar{j}$: sets $\boldsymbol{C}_j = H_j^{\tau \boldsymbol{v}_c} \cdot g^{\kappa \boldsymbol{w}_j}, \boldsymbol{C}'_j = g^{\boldsymbol{w}_j}$.
3. For each $k \in [l]$: sets $P_k = f^{A_k \cdot \boldsymbol{u}} G^{\xi_k}, \quad P'_k = (H^{\rho(k)} h)^{-\xi_k}, \quad P''_k = g^{\xi_k}$.

$\mathsf{Decrypt}_{\mathsf{A}}(\mathsf{PP}, CT_{R,(A,\rho)}, \mathsf{SK}_{(i,j),S}) \rightarrow M$ or $\perp$. If $(i,j) \in R$ or S does not satisfy (A,ρ), the algorithm outputs $\perp$, otherwise:

1. Since S satisfies (A,ρ), the algorithm can efficiently compute constants $\{\omega_k \in \mathbb{Z}_p\}$ such that $\sum_{\rho(k) \in S} \omega_k A_k = (1, 0, \ldots, 0)$, then compute
$$D_P = \prod_{\rho(k) \in S} \Big(e(K'_{i,j}, P_k) \cdot e(K_{i,j,\rho(k)}, P'_k) \cdot e(K'_{i,j,\rho(k)}, P''_k) \Big)^{\omega_k}$$
$$= \prod_{\rho(k) \in S} \big(e(g^{\sigma_{i,j}}, f^{A_k \cdot \boldsymbol{u}}) \big)^{\omega_k} = e(g^{\sigma_{i,j}}, f)^{\sum_{\rho(k) \in S} \omega_k (A_k \cdot \boldsymbol{u})} = e(g^{\sigma_{i,j}}, f)^{\pi}.$$
Note that if S does not satisfy (A,ρ), no such constants $\{\omega_k\}$ would exist.
2. Since $(i,j) \in \bar{R} (= [m,m] \setminus R)$ implies $j \in \bar{R}_i$, the algorithm can compute
$$\bar{K}_{i,j} = K_{i,j} \cdot (\prod_{j' \in \bar{R}_i \setminus \{j\}} \bar{K}_{i,j,j'}) = g^{\alpha_i} g^{r_i c_j} (f f_j)^{\sigma_{i,j}} \cdot (\prod_{j' \in \bar{R}_i \setminus \{j\}} f_{j'}^{\sigma_{i,j}})$$
$$= g^{\alpha_i} g^{r_i c_j} \cdot (f \prod_{j' \in \bar{R}_i} f_{j'})^{\sigma_{i,j}}.$$
Note that if $(i,j) \in R$ (implying $j \notin \bar{R}_i$), the algorithm cannot produce such a $\bar{K}_{i,j}$. The algorithm then computes
$$D_I = \frac{e(\bar{K}_{i,j}, Q_i) \cdot e(K''_{i,j}, Q''_i)}{e(K'_{i,j}, Q'_i)} \cdot \frac{e_3(\boldsymbol{R}'_i, \boldsymbol{C}'_j)}{e_3(\boldsymbol{R}_i, \boldsymbol{C}_j)}.$$
3. Computes $M = T_i / (D_P \cdot D_I)$ as the output message. Suppose that the ciphertext is generated from message M' and encryption index $(\bar{i}, \bar{j})$, it can be verified that only when $(i > \bar{i})$ or $(i = \bar{i} \wedge j \geq \bar{j})$, $M = M'$. This is because for $i > \bar{i}$, we have $(\boldsymbol{v}_i \cdot \boldsymbol{\chi}_3) = 0$ (since $\boldsymbol{v}_i \in span\{\boldsymbol{\chi}_1, \boldsymbol{\chi}_2\}$), and for $i = \bar{i}$, we have that $(\boldsymbol{v}_i \cdot \boldsymbol{\chi}_3) \neq 0$ happens with overwhelming probability (since $\boldsymbol{v}_i$ is randomly chosen from $\mathbb{Z}_p^3$). The correctness details can be found in the full version [19].

4.3 Augmented R-CP-ABE Security

The following theorem states that the AugR-CP-ABE proposed above is message-hiding. Then in Theorem 4, we state that the AugR-CP-ABE is also selectively index-hiding.

Theorem 3. *No PPT adversary can win* $\mathsf{Game}^{\mathsf{A}}_{\mathsf{MH}}$ *with non-negligible advantage.*

Proof. The argument for message-hiding in $\mathsf{Game}^{\mathsf{A}}_{\mathsf{MH}}$ is straightforward since an encryption to index $N+1$ (i.e. $(m+1,1)$) contains no information about the message. The simulator simply runs $\mathsf{Setup}_{\mathsf{A}}$ and $\mathsf{KeyGen}_{\mathsf{A}}$ and encrypts M_b under the challenge (revocation list, access policy) pair $(R^*, \mathbb{A}^*)$ and index $(m+1,1)$. Since for all $i=1$ to m, $T_i = E_i^{\hat{s}_i}$ contains no information about the message, the bit b is perfectly hidden and $\mathsf{MH}^{\mathsf{A}}\mathsf{Adv}_{\mathcal{A}} = 0$.

Theorem 4. *Suppose that the D3DH, the DLIN and the Extended Source Group q-Parallel BDHE Assumption hold. Then no PPT adversary can selectively win* $\mathsf{Game}^{\mathsf{A}}_{\mathsf{IH}}$ *with non-negligible advantage, provided that the challenge LSSS matrix's size* $l \times n$ *satisfies* $l, n \leq q$.

Proof. It follows Lemmas 1 and 2 below.

Lemma 1. *If the D3DH and the Extended Source Group q-Parallel BDHE Assumption hold, then for* $\bar{j} < m$, *no PPT adversary can selectively distinguish between an encryption to* $(\bar{i}, \bar{j})$ *and* $(\bar{i}, \bar{j}+1)$ *in* $\mathsf{Game}^{\mathsf{A}}_{\mathsf{IH}}$ *with non-negligible advantage, provided that the challenge LSSS matrix's size* $l \times n$ *satisfies* $l, n \leq q$.

Proof. In $\mathsf{Game}^{\mathsf{A}}_{\mathsf{IH}}$ with index $(\bar{i}, \bar{j})$, let $(R^*, (A^*, \rho^*))$ be the challenge (revocation list, access policy) pair, the restriction is that the adversary $\mathcal{A}$ does not query a secret key for (index, attribute set) pair $((i,j), S_{(i,j)})$ such that $((i,j) = (\bar{i}, \bar{j})) \wedge ((i,j) \in [m,m] \setminus R^*) \wedge \big(S_{(i,j)}\ satisfies\ (A^*, \rho^*)\big)$. Under this restriction, there are two ways for $\mathcal{A}$ to take:

Case I: In Phase 1 and Phase 2, $\mathcal{A}$ does not query a secret key with index $(\bar{i}, \bar{j})$.
Case II: In Phase 1 or Phase 2, $\mathcal{A}$ queries a secret key with index $(\bar{i}, \bar{j})$. Let $S_{(\bar{i},\bar{j})}$ be the corresponding attribute set. **Case II** has the following sub-cases:

1. $(\bar{i}, \bar{j}) \notin [m,m] \setminus R^*$, $S_{(\bar{i},\bar{j})}$ satisfies (A^*, ρ^*).
2. $(\bar{i}, \bar{j}) \notin [m,m] \setminus R^*$, $S_{(\bar{i},\bar{j})}$ does not satisfy (A^*, ρ^*).
3. $(\bar{i}, \bar{j}) \in [m,m] \setminus R^*$, $S_{(\bar{i},\bar{j})}$ does not satisfy (A^*, ρ^*).

If $\mathcal{A}$ is in **Case I**, **Case II.1** or **Case II.2**, it follows the restrictions in the index-hiding game for Augmented Broadcast Encryption (AugBE) in [8], where the adversary does not query the key with index $(\bar{i}, \bar{j})$ or $(\bar{i}, \bar{j})$ is not in the receiver list $[m,m] \setminus R^*$. **Case II.3** captures the index-hiding requirement of Augmented R-CP-ABE in that even if a user has a key with index $(\bar{i}, \bar{j})$ and $(\bar{i}, \bar{j}) \notin R^*$, the user cannot distinguish between an encryption to $(R^*, (A^*, \rho^*), (\bar{i}, \bar{j}))$ and

$(R^*, (A^*, \rho^*), (\bar{i}, \bar{j}+1))$ if $S_{(\bar{i},\bar{j})}$ does not satisfy (A^*, ρ^*). This is the most challenging part of proving the index-hiding when we attempt to *securely intertwine* the tracing techniques of broadcast encryption (e.g. [8]) into the large universe CP-ABE (e.g. [23]). Compared to the proof of [16], the challenge here is to prove the index-hiding in the large universe setting, as discussed previously.

To prove this lemma, we flip a random coin $c \in \{0,1\}$ as our guess on which case that $\mathcal{A}$ is in. In particular, if $c = 0$, we guess that $\mathcal{A}$ is in **Case I**, **Case II.1** or **Case II.2**, and make a reduction that uses $\mathcal{A}$ to solve a D3DH problem instance, using a proof technique similar to that of [8]. Actually, in this proof, we reduce from our AugR-CP-ABE to the AugBE in [8]. If $c = 1$, we guess that $\mathcal{A}$ is in **Case I**, **Case II.2** or **Case II.3**, and use $\mathcal{A}$ to solve an Extended Source Group q-Parallel BDHE problem instance, which is where the main novelty resides among all the proofs in this work. The proof details are provided in the full version [19].

Lemma 2. *If the D3DH, the DLIN and the Extended Source Group q-Parallel BDHE Assumption hold, then for $1 \leq \bar{i} \leq m$, no PPT adversary can selectively distinguish between an encryption to $(\bar{i}, m)$ and $(\bar{i}+1, 1)$ in $\mathsf{Game}^{\mathsf{A}}_{\mathsf{IH}}$ with non-negligible advantage, provided that the challenge LSSS matrix's size $l \times n$ satisfies $l, n \leq q$.*

Proof. Similar to the proof of Lemma 6.3 in [8], to prove this lemma we define the following hybrid experiment: H_1: encrypt to $(\bar{i}, \bar{j} = m)$; H_2: encrypt to $(\bar{i}, \bar{j} = m+1)$; and H_3: encrypt to $(\bar{i}+1, 1)$. This lemma follows Claims 1 and 2 below.

Claim 1. *If the D3DH and the Extended Source Group q-Parallel BDHE Assumption hold, then no PPT adversary can selectively distinguish between experiment H_1 and H_2 with non-negligible advantage, provided that the challenge LSSS matrix's size $l \times n$ satisfies $l, n \leq q$.*

Proof. The proof is identical to that for Lemma 1.

Claim 2. *If the D3DH and the DLIN hold, then no PPT adversary can distinguish between experiment H_2 and H_3 with non-negligible advantage.*

Proof. Note that $(\bar{i}, m+1) \notin [m, m]$ implies that for any revocation list $R^* \subseteq [m, m]$, we have $(\bar{i}, m+1) \notin \bar{R}^* (= [m, m] \setminus R^*)$, i.e., the adversaries for distinguishing H_2 and H_3 will not be in **Case II.3**. Thus, similar to that of case $c = 0$ in the proof of Lemma 1, in this proof we reduce from our AugR-CP-ABE to the AugBE in [8]. In the proof of index-hiding for an AugBE scheme Σ_{AugBE} in [8, Lemma 6.3], two hybrid experiments were defined and proven indistinguishable via a sequence of hybrid sub-experiments.

- H_2^{AugBE}: Encrypt to $(\bar{i}, m+1)$, (i.e. H_2 in [8])
- H_3^{AugBE}: Encrypt to $(\bar{i}+1, 1)$, (i.e. H_5 in [8])

By following [8, Lemma 6.3], *if the D3DH and the DLIN hold, no PPT adversary can distinguish between* H_2^{AugBE} *and* H_3^{AugBE} *for* Σ_{AugBE} *with non-negligible advantage.* Suppose there is a PPT adversary $\mathcal{A}$ that can distinguish between H_2 and H_3 for our AugR-CP-ABE scheme with non-negligible advantage. We can build a reduction, which is similar to that of case $c = 0$ in the proof of Lemma 1, to use $\mathcal{A}$ to distinguish between H_2^{AugBE} and H_3^{AugBE} for Σ_{AugBE} with non-negligible advantage.

5 Conclusion

In this paper, we proposed the first practical CP-ABE that simultaneously supports (1) traitor tracing, (2) revocation and (3) large universe. The scheme is highly expressive in supporting any monotonic access structures. Besides achieving fully collusion-resistant blackbox traceability and direct revocation, it is also efficient with the overhead in $O(\sqrt{N})$ only. Furthermore, it supports large attribute universe and does not need to fix the values of attributes during the system setup. The scheme was proven selectively secure and traceable in the standard model.

References

1. Attrapadung, N., Imai, H.: Conjunctive broadcast and attribute-based encryption. In: Shacham, H., Waters, B. (eds.) Pairing 2009. LNCS, vol. 5671, pp. 248–265. Springer, Heidelberg (2009)
2. Bethencourt, J., Sahai, A., Waters, B.: Ciphertext-policy attribute-based encryption. In: IEEE Symposium on Security and Privacy, pp. 321–334 (2007)
3. Boneh, D., Franklin, M.K.: An efficient public key traitor tracing scheme. In: Wiener, M. (ed.) Advances in Cryptology – CRYPTO 1999. LNCS, vol. 1666, pp. 338–353. Springer, Heidelberg (1999)
4. Boneh, D., Sahai, A., Waters, B.: Fully collusion resistant traitor tracing with short ciphertexts and private keys. In: Vaudenay, S. (ed.) EUROCRYPT 2006. LNCS, vol. 4004, pp. 573–592. Springer, Heidelberg (2006)
5. Boneh, D., Waters, B.: A fully collusion resistant broadcast, trace, and revoke system. In: ACM Conference on Computer and Communications Security, pp. 211–220 (2006)
6. Cheung, L., Newport, C.C.: Provably secure ciphertext policy ABE. In: ACM Conference on Computer and Communications Security, pp. 456–465 (2007)
7. Deng, H., Wu, Q., Qin, B., Mao, J., Liu, X., Zhang, L., Shi, W.: Who is touching my cloud. In: Kutyłowski, M., Vaidya, J. (eds.) ESORICS 2014, Part I. LNCS, vol. 8712, pp. 362–379. Springer, Heidelberg (2014)
8. Garg, S., Kumarasubramanian, A., Sahai, A., Waters, B.: Building efficient fully collusion-resilient traitor tracing and revocation schemes. In: ACM Conference on Computer and Communications Security, pp. 121–130 (2010)
9. Goyal, V., Jain, A., Pandey, O., Sahai, A.: Bounded ciphertext policy attribute based encryption. In: Aceto, L., Damgård, I., Goldberg, L.A., Halldórsson, M.M., Ingólfsdóttir, A., Walukiewicz, I. (eds.) ICALP 2008, Part II. LNCS, vol. 5126, pp. 579–591. Springer, Heidelberg (2008)

10. Goyal, V., Pandey, O., Sahai, A., Waters, B.: Attribute-based encryption for fine-grained access control of encrypted data. In: ACM Conference on Computer and Communications Security, pp. 89–98 (2006)
11. Herranz, J., Laguillaumie, F., Ràfols, C.: Constant size ciphertexts in threshold attribute-based encryption. In: Nguyen, P.Q., Pointcheval, D. (eds.) PKC 2010. LNCS, vol. 6056, pp. 19–34. Springer, Heidelberg (2010)
12. Katz, J., Schröder, D.: Tracing insider attacks in the context of predicate encryption schemes. In: ACITA (2011). https://www.usukita.org/node/1779
13. Lewko, A.: Tools for simulating features of composite order bilinear groups in the prime order setting. In: Pointcheval, D., Johansson, T. (eds.) EUROCRYPT 2012. LNCS, vol. 7237, pp. 318–335. Springer, Heidelberg (2012)
14. Lewko, A., Okamoto, T., Sahai, A., Takashima, K., Waters, B.: Fully secure functional encryption: attribute-based encryption and (hierarchical) inner product encryption. In: Gilbert, H. (ed.) EUROCRYPT 2010. LNCS, vol. 6110, pp. 62–91. Springer, Heidelberg (2010)
15. Lewko, A., Waters, B.: New proof methods for attribute-based encryption: achieving full security through selective techniques. In: Safavi-Naini, R., Canetti, R. (eds.) CRYPTO 2012. LNCS, vol. 7417, pp. 180–198. Springer, Heidelberg (2012)
16. Liu, Z., Cao, Z., Wong, D.S.: Blackbox traceable CP-ABE: how to catch people leaking their keys by selling decryption devices on ebay. In: ACM Conference on Computer and Communications Security, pp. 475–486 (2013)
17. Liu, Z., Cao, Z., Wong, D.S.: White-box traceable ciphertext-policy attribute-based encryption supporting any monotone access structures. IEEE Trans. Inf. Forensics Secur. **8**(1), 76–88 (2013)
18. Liu, Z., Cao, Z., Wong, D.S.: Traceable CP-ABE: how to trace decryption devices found in the wild. IEEE Trans. Inf. Forensics Secur. **10**(1), 55–68 (2015)
19. Liu, Z., Wong, D.S.: Practical attribute-based encryption: Traitor tracing, revocation, and large universe. IACR Cryptology ePrint Archive 2014, 616 (2014). http://eprint.iacr.org/2014/616
20. Naor, D., Naor, M., Lotspiech, J.: Revocation and tracing schemes for stateless receivers. In: Kilian, J. (ed.) CRYPTO 2001. LNCS, vol. 2139, pp. 41–62. Springer, Heidelberg (2001)
21. Okamoto, T., Takashima, K.: Fully secure functional encryption with general relations from the decisional linear assumption. In: Rabin, T. (ed.) CRYPTO 2010. LNCS, vol. 6223, pp. 191–208. Springer, Heidelberg (2010)
22. Okamoto, T., Takashima, K.: Fully secure unbounded inner-product and attribute-based encryption. In: Wang, X., Sako, K. (eds.) ASIACRYPT 2012. LNCS, vol. 7658, pp. 349–366. Springer, Heidelberg (2012)
23. Rouselakis, Y., Waters, B.: Practical constructions and new proof methods for large universe attribute-based encryption. In: ACM Conference on Computer and Communications Security, pp. 463–474 (2013)
24. Sahai, A., Seyalioglu, H., Waters, B.: Dynamic credentials and ciphertext delegation for attribute-based encryption. In: Safavi-Naini, R., Canetti, R. (eds.) CRYPTO 2012. LNCS, vol. 7417, pp. 199–217. Springer, Heidelberg (2012)
25. Sahai, A., Waters, B.: Fuzzy identity-based encryption. In: Cramer, R. (ed.) EUROCRYPT 2005. LNCS, vol. 3494, pp. 457–473. Springer, Heidelberg (2005)
26. Waters, B.: Ciphertext-policy attribute-based encryption: an expressive, efficient, and provably secure realization. In: Catalano, D., Fazio, N., Gennaro, R., Nicolosi, A. (eds.) PKC 2011. LNCS, vol. 6571, pp. 53–70. Springer, Heidelberg (2011)