Ciphertext-Policy Attribute-Based Broadcast Encryption with Small Keys

Benjamin Wesolowski[1] and Pascal Junod[2](✉)

[1] EPFL, Lausanne, Switzerland
benjamin.wesolowski@epfl.ch
[2] University of Applied Sciences and Arts Western Switzerland
(HES-SO/HEIG-VD), Yverdon-les-Bains, Switzerland
pascal.junod@heig-vd.ch

Abstract. Broadcasting is a very efficient way to securely transmit information to a large set of geographically scattered receivers, and in practice, it is often the case that these receivers can be grouped in sets sharing common characteristics (or attributes). We describe in this paper an efficient ciphertext-policy attribute-based broadcast encryption scheme (CP-ABBE) supporting negative attributes and able to handle access policies in conjunctive normal form (CNF). Essentially, our scheme is a combination of the Boneh-Gentry-Waters broadcast encryption and of the Lewko-Sahai-Waters revocation schemes; the former is used to express attribute-based access policies while the latter is dedicated to the revocation of individual receivers. Our scheme is the first one that involves a public key and private keys having a size that is independent of the number of receivers registered in the system. Its selective security is proven with respect to the Generalized Diffie-Hellman Exponent (GDHE) problem on bilinear groups.

Keywords: Attribute-based encryption · Broadcast encryption

1 Introduction

Broadcast channels allow transmitting information to a large set of geographically scattered receivers in a very efficient way. When this information is of high value, such as a high-definition Pay-TV stream or when delivered by a military geolocation system, for instance, one needs technical ways to enforce the signal reception by authorized receivers only. More than twenty years ago, the problem of securing a broadcast channel has began to attract cryptographers: the first works were the ones of Berkovits [2] and of Fiat and Naor [15], who coined the term "broadcast encryption". The underlying idea is that the broadcasting center sends an encrypted message to a set of non-revoked receivers, which is a

This work was supported by the EUREKA-Celtic+ H2B2VS project and by the University of Applied Sciences and Arts Western Switzerland (HES-SO). It was performed while the first author was working at HES-SO/HEIG-VD.

S. Kwon and A. Yun (Eds.): ICISC 2015, LNCS 9558, pp. 53–68, 2016.
DOI: 10.1007/978-3-319-30840-1_4

subset of all receivers. Obviously, revoked receivers (or other entities) spying the broadcast channel must not be able to decrypt a ciphertext, even if they collude together by sharing their private key material.

Precisely, if we denote by $\mathcal{U}$, with $n = |\mathcal{U}|$, the set of users (or receivers) and by $\mathcal{R}$, with $\ell = |\mathcal{R}|$, the set of revoked receivers, respectively, a *broadcast encryption scheme* is often meant to allow the secure transmission of information to an arbitrary set of receivers, *i.e.*, when $n - \ell \ll n$, while *revocation systems* are designed to exclude a small set of rogue receivers, *i.e.* when $\ell \ll n$.

A key characteristic of broadcast encryption and revocation schemes is the fact that no synchronism is assumed between the broadcasting center and the receivers, besides the initial key setup procedure: one speaks from *stateless* receivers. It means that, once each receiver is provisioned with its decryption key material, all the information required to decrypt a ciphertext must be contained in that ciphertext. Many stateless broadcast encryption schemes have been proposed in the past, being in the secret-key [18,20,34]) or in the public-key settings [6–8,12,13,17,27,37], while a large body of literature tackling the same problem, but for *stateful* receivers, this time, is available; we refer the reader to [9] and the references therein.

Attribute-Based Encryption. In practice, it is often the case that the receivers in a system can be grouped by common characteristics (or *attributes*). If we stick to a scenario around Pay-TV, receivers could be categorized by geographical location ("receivers located in California", "receivers located in a rural zone"), by technical capabilities ("receivers supporting HD content", "receivers supporting 4 K content", "receivers having an OS with patch level 3.14.159"), by subscription type ("receivers having access to the XYZ sport channels package", "receivers having access to the FGH adult channels package"), etc. Ideally, a broadcaster might then be willing to grant access to receivers according to a complicated access equation, such as to all "receivers having access to XYZ sport channels package, having an OS with patch level 3.14.159, but *not* located in California".

The idea of attribute-based encryption (ABE) has been proposed by Sahai and Waters in [41], as a generalization of identity-based encryption [5,42]; it was then formalized by Goyal and his co-authors in [19], who proposed the concepts of *ciphertext-policy (CP-ABE)* and *key-policy (KP-ABE)* encryption schemes. In the CP-ABE and KP-ABE models, the access policies are embedded in the ciphertext and in the private key, respectively. Since then, numerous variants of CP- and KP-ABE schemes have been published; see for instance [3,10,16,21,22,26,28,29,35,38,40,43].

Attribute-Based Broadcast Encryption. Transforming an ABE encryption scheme for using it in a broadcast scenario is a natural question, as in practice, broadcasters are most of the time addressing sets of receivers sharing the same characteristics, instead of individual ones. An exception where a receiver might be addressed individually is when a key update is necessary, for example. This operation is rather costly in terms of bandwidth, as synchronism comes into play. It means that the individual key update messages have to be broadcast sufficiently many times on a sufficiently long period to guarantee their reception with high probability. This explains why addressing individual receivers is not possible in practice to enforce access equations in a broadcast setting and why efficient stateless broadcast encryption schemes are so useful.

The key difference between an *attributed-based broadcast encryption (ABBE)* scheme and an ABE one is the additional possibility to revoke individual receivers in an efficient way. Given an ABE scheme, it is possible to create a revocation system by defining a dedicated unique attribute for each receiver and to specify an access policy which rejects the revoked receivers. Unfortunately, this is in general not efficient, since in an ABE scheme, the length of the keys or ciphertexts depend often in a linear way from the number of attributes. This can become unpractical when the number of receivers is large. Concretely, one could use an ABE supporting negative attributes, such as [35], and assign individual attributes to each receivers. A ciphertext can then be sent to the non-revoked receiver identities by conjunctively adding the AND of negations of revoked receivers attributes to the access policy. Implementing this idea with [35], this would imply an acceptable overhead of $O(\ell)$ group elements in the ciphertext, with $\ell = |\mathcal{R}|$, but the private key would involve $O(n)$ attributes, where n is the total number of receivers. Furthermore, this scheme would not be dynamic in the sense of [12], i.e., one cannot easily add receivers in the system without sending individual messages to the receivers, which is, as mentionned above, costly in terms of bandwidth in a broadcast setting.

In a context where the number of receivers is way larger than the number of attributes, one is therefore interested in splitting the revocation system from the access structure. Motivated by this fact, a line of research has focused on designing ABE schemes allowing to efficiently revoke individual receivers. In other words, revoking a receiver is implemented conjunctively, meaning that even if that receiver possesses compatible attributes for a given access equation, but it belongs to the revoked receivers set $\mathcal{R}$, it will not be able to correctly decrypt the ciphertext.

Lubicz and Sirvent [33] have proposed a scheme allowing to express access policies in disjunctive normal form (DNF), *i.e.*, with disjunctions (OR) of conjunctions (AND), and able to handle negative attributes (NOT). Then, Attrapadung and Imai [1] proposed another approach, namely using a separate broadcast encryption scheme on the top of an ABE scheme, and they constructed both ciphertext-policy and key-policy variants. Since then, other designs have been published as well, see e.g. [24,32,45].

Finally, we note that attribute-based broadcast encryption schemes have numerous applications besides the Pay-TV or the geolocation satellites scenarios mentionned above. For instance, applications involving ABBE have been proposed in the context of secure storage of personal health records [31], of securing smart grids [14], and, more generally, in any data outsourcing systems requiring privacy [23].

Our Contributions. In this paper, we describe an efficient ciphertext-policy attribute-based broadcast encryption scheme (CP-ABBE) able to handle access policies in conjunctive normal form (CNF), *i.e.*, as conjunctions of disjunctions of attributes, and supporting negative attributes. Essentially, our scheme is a combination of the Boneh-Gentry-Waters broadcast encryption scheme [6] and of the Lewko-Sahai-Waters revocation system [27]. The former is used to express attribute-based access policies while the latter is dedicated to the revocation of individual receivers.

Denoting by $\mathcal{B}$ the set of attributes, our scheme requires a public key and private keys of size $O(N)$, where $N = |\mathcal{B}|$ is the total number of attributes. Ciphertexts are of size $O(\bar{\nu}+\ell)$, where $\ell = |\mathcal{R}|$ is the number of revoked receivers and $\bar{\nu}$ is the number of clauses in the access policy. We note that $\bar{\nu}$, N and ℓ are quantities independent of the number n of receivers registered in the system. As a consequence, and to the best of our knowledge, our proposal is the first ABBE scheme whose public and private key sizes *do not depend on the number of receivers in the system*, while the ciphertext length keeps linear in the size of the access policy and in the number of revoked receivers. This property is especially important in scenarios involving large numbers of users, such as large-scale Pay-TV or cloud-based storage systems, for instance.

Eventually, we prove the selective security of our scheme with respect to the Generalized Diffie-Hellman Exponent (GDHE) problem on bilinear groups [4], and we derive security bounds in the generic group model.

This paper is organized as follows: in Sect. 2, we recall the formal definition of attribute-based broadcast encryption schemes, their underlying security model as well as other mathematical preliminaries. Then, we describe our new scheme Sect. 3 and we prove its security in Sect. 4. Finally, we compare its characteristics to other existing ABBE schemes and we discuss some of its practical aspects in Sect. 6.

2 Mathematical Preliminaries

Let $\mathcal{U}$ denote a set of receivers (or users), $\mathcal{R} \subset \mathcal{U}$ the set of revoked receivers and $\mathcal{B}$ a set of attributes. Furthermore, let λ be a security parameter. A ciphertext-policy attribute-based broadcast encryption (CP-ABBE) scheme consists of the following four algorithms:

- $\mathsf{Setup}(\lambda) \to (\mathsf{pk}, \mathsf{msk})$ is a randomized algorithm which takes a security parameter λ as input and outputs the public key pk and a master key msk.
- $\mathsf{KeyGen}(u, \omega, \mathsf{msk}, \mathsf{pk}) \to \mathsf{dk}_u$ is a randomized algorithm that takes as input a receiver $u \in \mathcal{U}$, a set of attributes $\omega \subset \mathcal{B}$, the master key msk and the public key pk. It outputs a private, individual decryption key $\mathsf{dk}_{(u,\omega)}$ for the receiver u. $\mathsf{dk}_{(u,\omega)}$ will be simply denoted dk_u if it is clear from the context that u has set of attributes ω.
- $\mathsf{Encrypt}(\mathcal{R}, \mathbb{A}, \mathsf{pk}) \to (\mathsf{hdr}, \mathsf{k})$ is a randomized algorithm that takes as input a set of revoked receivers $\mathcal{R} \subset \mathcal{U}$, a Boolean access policy $\mathbb{A}$ expressed in conjonctive normal form and the public key pk. It outputs a header hdr as well as a session key k.

- $\mathsf{Decrypt}(\mathsf{hdr}, (\mathcal{R}, \mathbb{A}), \mathsf{dk}_{(u,\omega)}, (u, \omega), \mathsf{pk}) \rightarrow \mathsf{k}$ or $\perp$ is an algorithm taking as input a header hdr, a set of revoked receivers $\mathcal{R}$, an access policy $\mathbb{A}$, a decryption key $\mathsf{dk}_{(u,\omega)}$ for receiver u equipped with attributes ω as well as the public key pk. It outputs the session key k if and only if ω satisfies $\mathbb{A}$ and u is not in $\mathcal{R}$; otherwise, it outputs $\perp$.

The *selective security* notion for CP-ABBE is defined by the following probabilistic game:

- **Setup.** The adversary chooses a distribution of attributes $\mathfrak{B} : \mathcal{U} \rightarrow \mathcal{P}(\mathcal{B})$, declares a set of revoked receivers $\mathcal{R}^* \subset \mathcal{U}$ and an access policy $\mathbb{A}^*$. The challenger runs the Setup algorithm and gives the public key pk to the adversary $\mathcal{A}$.
- **Query phase 1.** The adversary is allowed to (adaptively) issue queries to the challenger for private keys dk_u for receivers $u \in \mathcal{U}$ such that either $u \in \mathcal{R}^*$ or $\mathfrak{B}(u)$ does not satisfy the policy $\mathbb{A}^*$, *i.e.*, receivers not able to decrypt a ciphertext.
- **Challenge.** After having run the encryption algorithm $\mathsf{Encrypt}(\mathcal{R}^*, \mathbb{A}^*, \mathsf{pk})$, the challenger gets a header hdr and a session key k. Next, he draws a bit b uniformly at random, sets $\mathsf{k}_b = \mathsf{k}$ and picks k_{1-b} uniformly at random in the space of possible session keys. He finally gives the triple $(\mathsf{hdr}, \mathsf{k}_0, \mathsf{k}_1)$ to the adversary.
- **Query phase 2.** The adversary is again allowed to (adaptively) issue queries for private keys dk_u for receivers $u \in \mathcal{U}$ such that either $u \in \mathcal{R}^*$ or $\mathfrak{B}(u)$ does not satisfy the policy $\mathbb{A}^*$.
- **Guess.** The adversary outputs a guess bit b'.

The adversary wins the game if $b = b'$ and its advantage is defined as

$$\mathrm{Adv}^{\mathrm{ind}}(\lambda, \mathcal{U}, \mathcal{B}, \mathcal{A}) = |2\Pr[b = b'] - 1| .$$

The set of receivers u for which $\mathcal{A}$ requested the private keys is the set of *colluding receivers*. Hence, selective security ensures semantic security against colluding receivers if the advantage of the adversary is negligible.

We note that in the selective security model, the attacker must output the access policy *before* seeing the public parameters. A stronger model, named full security, has been proposed in [30]. While selective security is not the strongest model one might hope for our scheme, we think that it is stronger than what one could expect in practice, as the list of revoked nodes and the access equations are typically defined by the broadcaster.

Now, let us recall the notion of bilinear group. Let $\mathbb{G}$ and $\mathbb{G}_T$ be two (multiplicative) cyclic groups, and g a generator of $\mathbb{G}$. A map $e : \mathbb{G} \times \mathbb{G} \rightarrow \mathbb{G}_T$ is a *symmetric, non-degenerate pairing* if it is bilinear, *i.e.* for any $u, v \in \mathbb{G}$ and $a, b \in \mathbb{Z}$, we have $e(u^a, v^b) = e(u, v)^{ab}$, and if it is non-degenerate, *i.e.* $e(g, g) \neq 1$. Endowed with such a pairing, $\mathbb{G}$ is called a *bilinear group*. For practical purposes, let us further assume that in a bilinear group $\mathbb{G}$, both the action of $\mathbb{G}$ and the pairing e are efficiently computable. Finally, we recall the *Generalized Diffie-Hellman Exponent (GDHE) Problem* [4].

Definition 1 (GDHE Decisional Problem). *Let $\mathbb{G}$ and $\mathbb{G}_T$ be two groups of prime order p, g a generator of $\mathbb{G}$, and $e : \mathbb{G} \times \mathbb{G} \to \mathbb{G}_T$ a non-degenerate bilinear map. Let $f \in \mathbb{F}_p[X_1, \ldots, X_n]$ be a polynomial in n variables over $\mathbb{F}_p$, the finite field with p elements, and $P, Q \subset \mathbb{F}_p[X_1, \ldots, X_n]$ be two sets of polynomials, both containing 1. Choose $x_1, \ldots, x_n \in \mathbb{F}_p$ and $U \in \mathbb{G}_T$ uniformly at random. Given the elements*

$$g^{\pi(x_1,\ldots,x_n)} \text{ and } e(g,g)^{\rho(x_1,\ldots,x_n)}$$

for each $\pi \in P$ and $\rho \in Q$, the Generalized Diffie-Hellman Exponent (GDHE) Decisional Problem *is the problem of distinguishing $e(g,g)^{f(x_1,\ldots,x_n)}$ from U.*

Observe that in this setting, the classical Decisional Diffie-Hellman (DDH) problem reduces to an easy instance of the GDHE Decisional problem: let $P = \{1, a, b\}$, $Q = \{1\}$ and $f = ab$. Given g^a and g^b, we can distinguish g^{ab} from a uniform random element $h \in \mathbb{G}$ by observing that $e(g^a, g^b) = e(g^{ab}, g)$. This fact justifies the following definition, as in this example, (P, Q) and f are *dependent functions.*

Definition 2 (Dependent Functions). *A function f is said to be* dependent *on the sets P and Q if there exist constants $a_{\pi,\pi'}$ with $\pi, \pi' \in P$ and c_ρ with $\rho \in Q$ such that*

$$f = \sum_{\pi,\pi' \in P} a_{\pi,\pi'} \pi\pi' + \sum_{\rho \in Q} c_\rho \rho.$$

With this independence notion, it is proven that the (P, Q, f)-GDHE Decisional Problem is difficult in the generic group model.

Theorem 1 (Boneh et al. [4, Theorem A.2]**).** *Let*

$$d = \max\left\{2\deg(\pi), \deg(\rho), \deg(f) \mid \pi \in P, \rho \in Q\right\},$$

and $s = \max\{|P|, |Q|\}$ If f is independent of P and Q, then for any adversary $\mathcal{A}$ that makes a total of at most q queries to the oracle computing the group operations in $\mathbb{G}$, $\mathbb{G}_T$ and the pairing e, we have

$$|2\Pr[\mathcal{A} \textit{ outputs } 0] - 1| \leq \frac{(q + 2s + 2)^2 \cdot d}{p}.$$

3 The New Scheme

Basically, our new scheme is a secure combination of the Boneh-Gentry-Waters (BGW) broadcast encryption scheme [6] and the Lewko-Sahai-Waters (LSW) [27] revocation system. This design strategy, which is similar to the one of Junod and Karlov [24], is motivated as follows.

3.1 High-Level Description

The BGW scheme targets arbitrary sets of priviledged receivers and involves ciphertexts with a constant size, if, as customary, one omits bandwidth consumed by the description of the set of priviledged receivers to be addressed; its public and private keys have a size depending on the number of receivers; note that, with the BGW scheme, one needs the public key to decrypt. Hence, we use it to express arbitrary access equations, that typically depend on a small number of attributes when compared to the total number of receivers. On its side, the LSW revocation scheme has ciphertexts whose size depends on the number of revoked receivers; however, its encryption and decryption keys are independant of the total number of users in the system. In systems potentially involving millions of receivers, this is a decisive practical advantage.

Given an access structure in CNF form $\mathbb{A} = \beta_1 \wedge \cdots \wedge \beta_N$ and a revocation set $\mathcal{R}$, our idea is to associate to each clause β_i a fragment of the session key k_i which can be computed only by a receiver satisfying the corresponding clause, and a fragment k_0 computable by non-revoked receivers. Then, the session key k can be derived out of the k_i's.

This alone would not resist to an attack from colluding receivers: if receiver u is revoked but satisfies $\mathbb{A}$, he can compute k_i for $i = 1, \ldots, N$, and v is not revoked but does not satisfy $\mathbb{A}$, he can compute k_0; together, u and v can compute k. To prevent this, we do not allow a receiver u to compute k_i directly, but rather an blinded value $\mathsf{k}_i^{\varepsilon_u}$ thereof, where ε_u is a secret exponent unique for each receiver u. Then, k can be derived from any collection $(\mathsf{k}_i^{\varepsilon_u})_{i=1}^{n}$. If u can compute $\mathsf{k}_i^{\varepsilon_u}$ for $i = 1, \ldots, N$ and v can compute $\mathsf{k}_0^{\varepsilon_v}$, they cannot derive k.

3.2 Formal Definitions

Let us write $\mathcal{B}^* = \mathcal{B} \cup \neg\mathcal{B}$ the set of all attributes $\mathcal{B}$ and their negations $\neg\mathcal{B}$. Let $\mathfrak{B} : \mathcal{U} \to \mathcal{P}(\mathcal{B}^*)$ be a *distribution of attributes*, *i.e.*, a map such that for any receiver $u \in \mathcal{U}$ and attribute $a \in \mathcal{B}$, either $a \in \mathfrak{B}(u)$ or $\neg a \in \mathfrak{B}(u)$, but not both. Let $\mathrm{id} : \mathcal{U} \to (\mathbb{Z}/p\mathbb{Z})^*$ be a public injection, and $\imath : \mathcal{B}^* \to \{2, 4, 6, \ldots, t-1\}$ be a public bijection where $t = 4N + 1$.

$\mathsf{Setup}(\lambda) \to (\mathsf{pk}, \mathsf{msk})$ According to the security parameter λ, choose two groups $\mathbb{G}$ and $\mathbb{G}_T$ of prime order $p > 2^\lambda$ as well as a non-degenerate bilinear pairing $e : \mathbb{G} \times \mathbb{G} \to \mathbb{G}_T$. Additionnaly, choose two non-zero elements $g, h = g^\xi \in \mathbb{G}$ and seven random exponents $\alpha, \gamma, b, \beta, \delta, r$ and r' in $\mathbb{Z}/p\mathbb{Z}$. Finally, let $g_i = g^{\alpha^i}$. The public key pk consists of the elements of $\mathbb{G}$ g, $g_n^{\gamma r'}$, g^r, $g_{n+1}^{rr'}$, $g_{n+1}^{rr'b}$, $g_{n+1}^{rr'b^2}$, $h^{b\alpha^{n+1}r'r}$, $g^{\delta r}$, g_n, $\left(g_{\imath(a)}^r\right)_{a \in \mathcal{B}^*}$, and the two elements of $\mathbb{G}_T$ $e(g_1, g_n)^{rr'\beta\gamma}$ and $e(g_1, g_n)^{r\beta}$. The authority keeps the exponents secret.

$\mathsf{KeyGen}(u, \mathfrak{B}(u), \mathsf{msk}, \mathsf{pk}) \to \mathsf{dk}_u$ Let $u \in \mathcal{U}$. Choose two random elements $\sigma_u, \varepsilon_u \in \mathbb{Z}/p\mathbb{Z}$. Define

$$D_{u,0} = \left(g^\gamma g^{b^2\sigma_u}\right)^{\varepsilon_u}, D_{u,1} = \left(g^{b\cdot\mathrm{id}(u)}h\right)^{\sigma_u\varepsilon_u}, D_{u,2} = g^{-\sigma_u\varepsilon_u}, D_{u,3} = g_1^{r(\beta+\varepsilon_u)}.$$

The private key of receiver u is

$$\mathsf{dk}_u = \left((D_{u,k})_{k=0}^{3}, \left(g_{\imath(a)}^{\varepsilon_u}\right)_{a \in \mathcal{B}^*}, \left(g_{n+1+\imath(a)}^{\varepsilon_u}\right)_{a \in \mathcal{B}^*}, \left(g_{\imath(a)}^{\delta\varepsilon_u}\right)_{a \in \mathfrak{B}(u)} \right).$$

$\mathsf{Encrypt}(\mathcal{R}, \mathbb{A}, \mathsf{pk}) \rightarrow (\mathsf{hdr}, \mathsf{k})$ Given an access policy $\mathbb{A} = \beta_1 \wedge \ldots \wedge \beta_N$, with $\beta_i = \beta_{i,1} \vee \ldots \vee \beta_{i,M_i}$ (modeled as $\beta_{i,j} \subseteq \mathcal{B} \cup \neg\mathcal{B}$) and a revocation set $\mathcal{R} \subset \mathcal{U}$, one chooses $s_0, \ldots, s_N \in \mathbb{Z}/p\mathbb{Z}$ at random and one defines $s = \gamma r' s_0 + \sum_{i=1}^{N} s_i$ (which needs not be computed). Also, one splits $s_0 = \sum_{u \in \mathcal{R}} s_u$. Let us define

$$C = g_n^s = \left(g_n^{\gamma \cdot r'}\right)^{s_0} g_n^{\left(\sum_{i=1}^{N} s_i\right)}.$$

For all $i = 1, \ldots, N$, one defines the elements $C_{i,0} = g^{r s_i}$ and

$$C_{i,1} = \left(g^{r\delta} \prod_{a \in \beta_i} g_{n+1-\imath(a)}^{r} \right)^{s_i},$$

as well as the corresponding N parts of the header $\mathsf{hdr}_i = (C_{i,0}, C_{i,1})$. One defines $C_0 = g_{n+1}^{r r' s_0}$, and for each $u \in \mathcal{R}$,

$$C_{u,1} = g_{n+1}^{r r' b s_u} \text{ and } C_{u,2} = \left(g^{b^2 \mathrm{id}(u)} h^b\right)^{\alpha^{n+1} r r' s_u}.$$

Let $\mathsf{hdr}_0 = (C_0, (C_{u,1})_{u \in \mathcal{R}}, (C_{u,2})_{u \in \mathcal{R}})$ and $\mathsf{hdr} = (C, \mathsf{hdr}_0, \ldots, \mathsf{hdr}_N)$. The global session key k is given by

$$\mathsf{k} = e(g_1, g_n)^{r\beta s} = \left(e(g_1, g_n)^{r r' \beta \gamma}\right)^{s_0} \cdot e\left(g_1^r, g_n^\beta\right)^{\left(\sum_{i=1}^{N} s_i\right)}$$

$\mathsf{Decrypt}(\mathsf{hdr}, (\mathcal{R}, \mathbb{A}), \mathsf{dk}_u, (u, \omega), \mathsf{pk}) \rightarrow \mathsf{k}$ or $\perp$ If $u \in \mathcal{R}$ or if there exists $i \in \{1, \ldots, N\}$, such that $\beta_i \cap \mathfrak{B}(u) = \emptyset$, return $\perp$. For $i = 1, \ldots, N$, choose one satisfying attribute $a \in \beta_i \cap \mathfrak{B}(u)$ and compute

$$\mathsf{k}_i^{\varepsilon_u} = \frac{e(g_{\imath(a)}^{\varepsilon_u}, C_{i,1})}{e\left(g_{\imath(a)}^{\delta\varepsilon_u} \prod_{a' \in \beta_i \setminus \{a\}} g_{n+1-\imath(a')+\imath(a)}^{\varepsilon_u}, C_{i,0}\right)}.$$

Also compute $\mathsf{k}_0^{\varepsilon_u}$ as

$$e(D_{u,0}, C_0) e\left(D_{u,1}, \prod_{u' \in \mathcal{R}} C_{u',1}^{1/(\mathrm{id}(u)-\mathrm{id}(u'))}\right)^{-1} e\left(D_{u,2}, \prod_{u' \in \mathcal{R}} C_{u',2}^{1/(\mathrm{id}(u)-\mathrm{id}(u'))}\right)^{-1}.$$

We have $\mathsf{k}_0^{\varepsilon_u} = e(g_1, g_n)^{r r' s_0 \varepsilon_u \gamma}$ and $\mathsf{k}_i^{\varepsilon_u} = e(g_1, g_n)^{r s_i \varepsilon_u}$ for $i = 1, ..., N$. Eventually, we can recover k as

$$\mathsf{k} = \frac{e(D_{u,3}, C)}{\prod_{i=0}^{N} \mathsf{k}_i^{\varepsilon_u}} = e(g_1, g_n)^{r\beta s}.$$

One can observe that the public-key size depends only on the total number of attributes defined in the system, and that the same holds for the decryption keys. The header size linearly depends only on the number of revoked rogue receivers.

If the number of attributes does not change during the lifetime of the system, we note that our new ABBE scheme is fully dynamic in the sense of [12]. Indeed, the deployment of new receivers does not imply to change the encryption or the decryption keys of other receivers, which is a desirable property for a stateless scheme.

At first sight, the system of attributes might look a bit less flexible in the sense that all receivers decryption keys include elements depending on all positive and negative attributes defined in the system. It means that the definition of new attributes after the system start arrives with the necessity of transmitting them to all receivers in a individual way, which comes with significant bandwidth issues in a system involving millions of receivers. However, this burden keeps acceptable if one considers the fact that one can define sufficiently many attributes at the start of the system and thus easily keep the set of attributes completely static during the system lifetime.

4 Security Analysis

To prove the security of our scheme, and similarly to the approach taken in [12], we show that the CP-ABBE selective security of this scheme reduces to an instance of a (P, Q, f)-GDHE problem [4]. We then prove that (P, Q) and f are independent, which implies in particular that the corresponding problem is difficult in the generic group model. This leads to a security reduction in the standard model, and a proof of security in the generic group model. Thereafter, all the polynomials considered are from the polynomial ring

$$\mathbb{F}_p[\alpha, \beta, \gamma, \delta, \xi, b, r, r', s_i, s_u, \sigma_u, \varepsilon_u : i \in \mathbb{N}, u \in \mathcal{U}].$$

Let $\mathcal{A}$ be an adversary for the CP-ABBE selective security game. It declares a distribution of attributes $\mathfrak{B} : \mathcal{U} \to \mathcal{P}(\mathcal{B}^*)$, an access structure $\mathbb{A}$ and a set $\mathcal{R}$ of revoked receivers. Let $\mathcal{C}$ be the set of all receivers which do not satisfy the policy $\mathbb{A}$, and/or are revoked. Let P be the list of polynomials consisting of 1, and all the following elements corresponding to the information in pk, hdr, and dk_u for all the receivers $u \in \mathcal{C}$.

1. Contribution of pk: the set P_{pk} of polynomials 1, $\alpha^n \gamma r'$, r, $\alpha^{n+1} r r'$, $\alpha^{n+1} r r' b$, $\alpha^{n+1} r r' b^2$, $\xi b \alpha^{n+1} r r'$, δr, α^n and for $a \in \mathcal{B}^*$, the element $\alpha^{\imath(a)} r$.
2. Contribution of dk_u, for any $u \in \mathcal{C}$: the set P_{dk_u} of polynomials $\varepsilon_u(\gamma + b^2 \sigma_u)$, $\sigma_u \varepsilon_u (b \cdot \mathrm{id}(u) + \xi)$, $\sigma_u \varepsilon_u$, $\alpha r(\beta + \varepsilon_u)$, for each $a \in \mathcal{B}^*$, $\alpha^{\imath(a)} \varepsilon_u$, $\alpha^{n+1+\imath(a)} \varepsilon_u$, and for each $a \in \mathfrak{B}(u)$, $\alpha^{\imath(a)} \delta \varepsilon_u$;
3. Contribution of hdr: the set P_{hdr} of polynomials $\alpha^n s$, $\alpha^{n+1} r r' s_0$, for each $i = 1, \ldots, N$, $r s_i$, $r s_i \left(\delta + \sum_{a \in \beta_i} \alpha^{n+1-\imath(a)}\right)$, and for each revoked receiver $u \in \mathcal{R}$, $\alpha^{n+1} r r' b s_u$, $\alpha^{n+1} r r' s_u (b^2 \cdot \mathrm{id}(u) + \xi b)$.

The list Q is simply $(1, \alpha^{n+1}rr'\beta\gamma, \alpha^{n+1}r\beta)$ and $f = \alpha^{n+1}rs\beta$.

Lemma 1. *If the adversary $\mathcal{A}$ solves the CP-ABBE selective security game with advantage ε, then a simulator can be constructed to solve the (P, Q, f)-GDHE problem with advantage ε in polynomial time, with one oracle call to $\mathcal{A}$.*

Proof. The proof is available in the full version of this paper [44].

According to Lemma 1, an adversary for the CP-ABBE selective security game gives rise to an adversary for the (P, Q, f)-GDHE problem. It now needs to be justified that the (P, Q, f)-GDHE problem is difficult. The end of Sect. 2 explains that we can suppose this problem to be difficult when (P, Q) and f are independent: it is proven to be difficult in the generic group model, and assumed to remain difficult in cryptographic bilinear groups. Thus, it remains to show that (P, Q) and f are indeed independent.

Lemma 2. *(P, Q) and f are independent.*

Proof. The proof is available in the full version of this paper [44].

We are now able to derive a bound on the security of our new scheme in the generic group model.

Theorem 2. *For any probabilistic algorithm $\mathcal{A}$ that totalizes at most q queries to the oracle performing group operations in $(\mathbb{G}, \mathbb{G}_T)$ and evaluations of $e(\cdot, \cdot)$, and declaring a set of revoked receivers of size at most η, as well as an access policy with at most N clauses ($\mathbb{A} = \beta_1 \wedge \cdots \wedge \beta_N$), then $\mathrm{Adv}^{\mathrm{ind}}(\lambda, \mathcal{U}, \mathcal{B}, \mathcal{A})$ is smaller or equal to*

$$\frac{(q + 4(N + N + \eta) + 22 + |\mathcal{U}|(10N + 8))^2(8N + 3)}{2^{\lambda - 1}}.$$

Proof. This is a direct consequence of Lemmas 1 and 2, and Theorem 1, with $|P_{\mathsf{pk}}| = 9 + 2N$, $|P_{\mathsf{dk}_u}| = 4 + 5N$, $|P_{\mathsf{hdr}}| = 2 + 2N + 2\ell$ and $d = 16N + 6$.

5 Optimizing the Bandwidth and Computational Overheads

As the number of revoked receivers grows, the computation of $\mathsf{k}_0^{\varepsilon_u}$ can become expensive for the receivers. The heavy computations are the products

$$\prod_{u' \in \mathcal{R}} C_{u',i}^{1/(\mathrm{id}(u) - \mathrm{id}(u'))}$$

for $i = 1, 2$, which require $O(\ell)$ exponentiations. This could be optimized if the $C_{u',1}$'s and $C_{u',2}$'s did not change from a message to another: those products could be computed the first time and reused, and any new revoked receiver would only require one exponentiation and multiplication for each of the receivers.

To do so, the broadcaster chooses a random $s_{u'}$ for every revoked receiver $u' \in \mathcal{R}$, and reuses it for all the following communications, thus generating the same $C_{u',1}$'s and $C_{u',2}$'s.

This optimization requires a new proof of security. We can show that even if the adversary is given access to old ciphertexts $\mathsf{hdr}^{(1)}, \ldots, \mathsf{hdr}^{(m)}$, (in addition to the challenge hdr) for which the sets of revoked receivers are subsets $\mathcal{R}^{(j)}$ of the set of revoked receivers $\mathcal{R}$ for hdr, and the access policies have $N^{(j)}$ clauses denoted $\beta_i^{(j)}$, for each $j = 1, \ldots, m$, the underlying (P, Q, f)-GDHE is still difficult (*i.e.*, (P, Q) and f are independent). We need to suppose $N^{(j)} > 0$ for each $j = 1, \ldots, m$.

This technique reduces the computational cost, but in a fully stateless situation, the broadcaster still needs to send all the $C_{u',1}$'s and $C_{u',2}$'s with each message. In a context where it is possible to maintain a synchronized state, via a two-way connection with a possibly very limited bandwidth, it is possible for the broadcaster to send with each ciphertext only the $C_{u',1}$'s and $C_{u',2}$'s for the newly revoked receivers. Then, the ciphertexts' lengths drop from $O(N + \ell)$ to $O(N + |\Delta\mathcal{R}|)$ (where $\Delta\mathcal{R}$ is the set of newly revoked receivers, for example those revoked during the last day or the last week).

The only thing we have to change from the setting of the original security proof is to add to P the contribution of the ciphertexts $\mathsf{hdr}^{(1)}, \ldots, \mathsf{hdr}^{(m)}$, where the secret exponents of $\mathsf{hdr}^{(j)}$ are denoted $s^{(j)}, s_0^{(j)}, s_i^{(j)}$ and $s_{u'}^{(j)}$ for $i = 1, \ldots, N^{(j)}$ and $u' \in \mathcal{R}^{(j)}$. This contribution consists, for each $j = 1, \ldots, m$, of the polynomials $\alpha^n s^{(j)}, \alpha^{n+1} r r' s_0^{(j)}$ and for each $i = 1, \ldots, N^{(j)}$, the polynomials

$$r s_i^{(j)}, r s_i^{(j)} \left(\delta + \sum_{a \in \beta_i^{(j)}} \alpha^{n+1-\imath(a)} \right).$$

Only a few observations are needed to adapt the original security proof to this new setting. The first thing is to notice that we now have new terms with a factor $\alpha^{n+1}\beta$. Those are, for any $j = 1, \ldots, M$ and $u \in \mathcal{C}$, $\alpha^{n+1} r s^{(j)} (\beta + \varepsilon_u)$. But those terms cannot have a non-zero coefficient in the linear combination forming f, because for each j, $\alpha^{n+1} r s^{(j)} (\beta + \varepsilon_u)$ is the only term containing the monomial $\alpha^{n+1} r s_1^{(j)} (\beta + \varepsilon_u)$, thus the later could not be canceled by any other linear combination of terms (here we use our assumption that $N^{(j)} > 0$).

The second thing to notice is that the terms which can cancel the monomials $\alpha^{n+1} r \varepsilon_u r' \gamma s_v$ for $v \in \mathcal{R}$ are now not only $\alpha^{n+1} r r' s_0 \varepsilon_u (\gamma + b^2 \sigma_u)$, but also the terms $\alpha^{n+1} r r' s_0^{(j)} \varepsilon_u (\gamma + b^2 \sigma_u)$ for all the j's such that $v \in \mathcal{R}^{(j)}$. We can then deduce that there is a linear combination of those terms such that the resulting coefficient of the monomial $\alpha^{n+1} r \varepsilon_u r' \gamma s_v$ is non-zero, and this coefficient is the same as the one of $\alpha^{n+1} r r' s_v \varepsilon_u b^2 \sigma_u$, which therefore is also non-zero. The end of the proof, consisting in showing that this coefficient of $\alpha^{n+1} r r' s_v \varepsilon_u b^2 \sigma_u$ cannot be canceled, remains unchanged. In conclusion, one can safely reuse the secret exponents s_u.

6 Practical Aspects

In this section, we compare the practical properties of our scheme to the other existing ABBE schemes listed in Table 1.

Table 1. Bandwidth and key storage complexity comparison. Denoting the set of all receivers by $\mathcal{U}$, the set of all attributes by $\mathcal{B}$, the set of revoked receivers by $\mathcal{R}$, then k_u is the number of attributes assigned to a receiver $u \in \mathcal{U}$, ν the length of the access structure, $\bar{\nu}$ the number of clauses in a CNF access structure, $N = |\mathcal{B}|$, $n = |\mathcal{U}|$ and $\ell = |\mathcal{R}|$.

Scheme	Access structure	Size of pk	Size of dk_u	Size of hdr
Attrapadung-Imai [1]	Monotone	$O(N+n)$	$O(N+n)$	$O(\nu)$
Lubicz-Sirvent [33]	AND & NOT	$O(N+n)$	$O(k_u)$	$O(\nu+\ell)$
Junod-Karlov [24]	CNF	$O(N+n)$	$O(N+n)$	$O(\bar{\nu})$
Zhou-Huang [45]	AND & NOT	$O(N+\log n)$	$O(N+\log n)$	$\approx O(\log n)$
Li-Zhang [32]	Monotone	$O(N+n)$	$O(k_u+n)$	$O(\nu)$
This paper	CNF	$O(N)$	$O(N)$	$O(\bar{\nu}+\ell)$

Size of Keys. First, we observe that our scheme is the only one where the public and private key sizes do not depend on the total number of receivers $n = |\mathcal{U}|$ registered in the system. Except for the Zhou-Huang scheme, whose dependency is of logarithmic nature, this dependence in n is linear in the competing schemes, which is highly impractical for a large scale deployment potentially involving millions of receivers, such as a Pay-TV system, for instance. The length of the keys in our scheme only depends linearly on the total number of attributes $N = |\mathcal{B}|$ defined in the system. This allows high scalability: the broadcaster can initially decide on a large set of possible receivers $\mathcal{U}$ without affecting the length of the keys. Adding new receivers to the system can be done efficiently, whereas with a key size linear in n, the broadcaster should choose the smallest possible $\mathcal{U}$ and change all the settings and keys when there are too many new receivers. This is undesirable in practice, as changing all the keys is way too expensive, especially when they are so long. In a nutshell, from the point of view of the key lengths, the Zhou-Huang scheme and our scheme are the only really practical candidates for large-scale deployment, while the Lubicz-Sirvent scheme can also be considered as acceptable since only its public key size is large, the private keys being pretty small.

Ciphertexts Size. The overhead on the ciphertext is $O(N+\ell)$ for our scheme, which is the same as the Lubicz-Sirvent scheme. The three schemes presenting a smaller overhead of size $O(N)$ have to compensate with private keys whose size is linear in n. The Zhou-Huang scheme can in theory reach an overhead as small as $O(\log n)$. This length relies on an optimization phase, which leads to an average length in $O(\log n)$ and a worst case length in $O(n)$; the worst case however occurs with small probability. This optimization phase is a Sum-of-Product Expression (SOPE) minimization, which is known to be an NP-Hard problem, so we can only hope for approximations. Finally, we would like to emphasize that ν and $\bar{\nu}$ have a somewhat different cardinality in the case of access structures involving only AND and NOT gates or in the case of complete CNF formulas. In the first case, ν represents the number of atomic variables, *i.e.*, the number of attributes or their negation, while in the case of a complete CNF formula, $\bar{\nu}$ represents the number of clauses, and it is independent of the number of atomic variables in the clauses. Hence, $\bar{\nu}$ is always smaller or equal, if not significantly smaller, than ν.

Overall Comparison. As mentionned before large-scale deployments rule out the schemes with a private key of length linear in $n = |\mathcal{U}|$. Remain the Lubicz-Sirvent and the Zhou-Huang schemes, which we will compare to ours. Compared to the Lubicz-Sirvent scheme, our scheme allows a much shorter public key; our private keys can be slightly larger, but still bounded by $O(N)$, which should not make a significant difference as long as the set of attributes remains reasonably small. The ciphertext overhead is the same. Our scheme allows a more flexible access control model via CNF formulas. The Lubicz-Sirvent only allows AND and NOT gates; one can also add OR gates, allowing access control by CNF formulas, via ciphertext concatenation, but the ciphertext overhead is then multiplied by the number of clauses. Note that, similarly to the Junod-Karlov scheme, our scheme allows to implement access policies in DNF form by concatenation as well. Overall, as long as $N = \mathcal{B}$ is of reasonable size, our scheme is more flexible and efficient than the Lubicz-Sirvent one. Compared to the Zhou-Huang scheme, the lengths of the public and private keys are similar; even though there is this additional term $\log n$ in the Zhou-Huang's scheme, there is no difference under the reasonable assumption that $N = O(\log n)$. As for the Lubicz-Sirvent scheme, the Zhou-Huang scheme only allows AND and NOT gates, and OR gates via ciphertext concatenation and a ciphertext overhead multiplied by the number of clauses. Furthermore, as mentioned above, the ciphertext overhead depends on the SOPE minimization phase, which is a NP-hard problem.

Practical Performances. We have implemented our new scheme using the C programming language and with help of the PBC library[1] for the elliptic curve arithmetic and pairings. The curve used let us work in a group of 160-bit long order and a base field of 512-bit long order, suitable for cryptographic use (it is a Type A curve, in PBC's classification). We ran an example with 5 attributes, on a 2.3 GHz Intel Core i7; the setup phase, including the generation of the public key takes 237 ms, generating the private key of a receiver takes 75 ms, the decryption of a message with 3 clauses, and without new revocations takes 25 ms, and each new revocation adds 4 ms to the first decryption after the revocation.

[1] This open-source library is freely available at http://crypto.stanford.edu/pbc/.

7 Conclusion

This paper describes, to the best of our knowledge, the first attribute-based broadcast encryption scheme for which the length of the encryption and decryption keys does not depend on the total number of users, but only on the number of attributes defined in the system. This property has been achieved by combining the Boneh-Gentry-Waters broadcast encryption scheme with the Lewko-Sahai-Waters revocation system in a secure way. Our scheme requires also a modest bandwidth, as the length of the header depends only of the number of revoked rogue receivers. The access equations can be defined in conjunctive normal form, *i.e.*, as AND of clauses involving ORs of attributes, and it supports negative attributes. We have proven the security of this scheme relatively to a GDHE problem in the standard model, which additionnaly allows us to derive corresponding security bounds in the generic group model. In summary, we are convinced that our scheme is fully practical in a number of real-life scenarios, including Pay-TV or cloud-storage ones involving millions of users.

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. Berkovits, S.: How to broadcast a secret. In: Davies, D.W. (ed.) EUROCRYPT 1991. LNCS, vol. 547, pp. 535–541. Springer, Heidelberg (1991)
3. Bethencourt, J., Sahai, A., Waters, B.: Ciphertext-policy attribute-based encryption. In: Proceedings of IEEE-S&P 2007, pp. 321–334. IEEE Computer Society (2007)
4. Boneh, D., Boyen, X., Goh, E.-J.: Hierarchical identity based encryption with constant size ciphertext. In: Cramer [11], pp. 440–456
5. Boneh, D., Franklin, M.K.: Identity-based encryption from the Weil pairing. In: Kilian [25], pp. 213–229
6. Boneh, D., Gentry, C., Waters, B.: Collusion resistant broadcast encryption with short ciphertexts and private keys. In: Shoup, V. (ed.) CRYPTO 2005. LNCS, vol. 3621, pp. 258–275. Springer, Heidelberg (2005)
7. Boneh, D., Waters, B.: A fully collusion resistant broadcast, trace, and revoke system. In: Proceedings of ACM-CCS 2006, pp. 211–220. Association for Computing Machinery, New York, NY, USA (2006)
8. Boneh, D., Waters, B., Zhandry, M.: Low overhead broadcast encryption from multilinear maps. In: Garay, J.A., Gennaro, R. (eds.) CRYPTO 2014, Part I. LNCS, vol. 8616, pp. 206–223. Springer, Heidelberg (2014)
9. Burmester, M.: Group key agreement. In: van Tilborg, H.C.A., Jajodia, S. (eds.) Encyclopedia of Cryptography and Security, pp. 520–526. Springer, New York (2011)
10. Chase, M.: Multi-authority attribute based encryption. In: Vadhan, S.P. (ed.) TCC 2007. LNCS, vol. 4392, pp. 515–534. Springer, Heidelberg (2007)
11. Cramer, R. (ed.): EUROCRYPT 2005. LNCS, vol. 3494. Springer, Heidelberg (2005)

12. Delerablée, C., Paillier, P., Pointcheval, D.: Fully collusion secure dynamic broadcast encryption with constant-size ciphertexts or decryption keys. In: Takagi, T., Okamoto, E., Okamoto, T., Okamoto, T. (eds.) Pairing 2007. LNCS, vol. 4575, pp. 39–59. Springer, Heidelberg (2007)
13. Dodis, Y., Fazio, N.: Public key broadcast encryption for stateless receivers. In: Feigenbaum, J. (ed.) DRM 2002. LNCS, vol. 2696, pp. 61–80. Springer, Heidelberg (2003)
14. Fadlullah, Z.M., Kato, N., Lu, R., Shen, X., Nozaki, Y.: Toward secure targeted broadcast in smart grid. IEEE Commun. Mag. **50**(5), 150–156 (2012)
15. Fiat, A., Naor, M.: Broadcast encryption. In: Stinson, D.R. (ed.) CRYPTO 1993. LNCS, vol. 773, pp. 480–491. Springer, Heidelberg (1994)
16. Garg, S., Gentry, C., Halevi, S., Sahai, A., Waters, B.: Attribute-based encryption for circuits from multilinear maps. In: Canetti, R., Garay, J.A. (eds.) CRYPTO 2013, Part II. LNCS, vol. 8043, pp. 479–499. Springer, Heidelberg (2013)
17. Gentry, C., Waters, B.: Adaptive security in broadcast encryption systems (with short ciphertexts). In: Joux, A. (ed.) EUROCRYPT 2009. LNCS, vol. 5479, pp. 171–188. Springer, Heidelberg (2009)
18. Goodrich, M.T., Sun, J.Z., Tamassia, R.: Efficient tree-based revocation in groups of low-state devices. In: Franklin, M. (ed.) CRYPTO 2004. LNCS, vol. 3152, pp. 511–527. Springer, Heidelberg (2004)
19. 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)
20. Halevy, D., Shamir, A.: The LSD broadcast encryption scheme. In: Yung, M. (ed.) CRYPTO 2002. LNCS, vol. 2442, p. 47. Springer, Heidelberg (2002)
21. Hohenberger, S., Waters, B.: Attribute-based encryption with fast decryption. In: Kurosawa, K., Hanaoka, G. (eds.) PKC 2013. LNCS, vol. 7778, pp. 162–179. Springer, Heidelberg (2013)
22. Hohenberger, S., Waters, B.: Online/offline attribute-based encryption. In: Krawczyk, H. (ed.) PKC 2014. LNCS, vol. 8383, pp. 293–310. Springer, Heidelberg (2014)
23. Hur, J., Noh, D.K.: Attribute-based access control with efficient revocation in data outsourcing systems. IEEE Trans. Parallel Distrib. Syst. **22**(7), 1214–1221 (2011)
24. Junod, P., Karlov, A.: An efficient public-key attribute-based broadcast encryption scheme allowing arbitrary access policies. In: Proceedings of DRM, pp. 13–24. ACM (2010)
25. Kilian, J. (ed.): CRYPTO 2001. LNCS, vol. 2139. Springer, Heidelberg (2001)
26. 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)
27. Lewko, A.B., Sahai, A., Waters, B.: Revocation systems with very small private keys. In: Proceedings of IEEE S&P, pp. 273–285. IEEE (2010)
28. Lewko, A.B., Waters, B.: Decentralizing attribute-based encryption. In: Paterson [36], pp. 568–588
29. Lewko, A.B., Waters, B.: Unbounded HIBE and attribute-based encryption. In: Paterson [36], pp. 547–567
30. Lewko, A.B., Waters, B.: New proof methods for attribute-based encryption: achieving full security through selective techniques. In: Safavi-Naini and Canetti [39], pp. 180–198

31. Li, M., Yu, S., Zheng, Y., Ren, K., Lou, W.: Scalable and secure sharing of personal health records in cloud computing using attribute-based encryption. IEEE Trans. Parallel Distrib. Syst. **24**(1), 131–143 (2013)
32. Li, Q., Zhang, F.: A fully secure attribute based broadcast encryption scheme. Int. J. Netw. Secur. **17**(3), 263–271 (2015)
33. Lubicz, D., Sirvent, T.: Attribute-based broadcast encryption scheme made efficient. In: Vaudenay, S. (ed.) AFRICACRYPT 2008. LNCS, vol. 5023, pp. 325–342. Springer, Heidelberg (2008)
34. 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)
35. Ostrovsky, R., Sahai, A., Waters, B.: Attribute-based encryption with non-monotonic access structures. In: Ning, P.., De Capitani di Vimercati, S., Syverson, P. F. (eds.) Proceedings of ACM-CCS, 2007 pp. 195–203. ACM (2007)
36. Paterson, K.G. (ed.): EUROCRYPT 2011. LNCS, vol. 6632. Springer, Heidelberg (2011)
37. Phan, D.-H., Pointcheval, D., Shahandashti, S.F., Strefler, M.: Adaptive CCA broadcast encryption with constant-size secret keys and ciphertexts. In: Susilo, W., Mu, Y., Seberry, J. (eds.) ACISP 2012. LNCS, vol. 7372, pp. 308–321. Springer, Heidelberg (2012)
38. Rouselakis, Y., Waters, B.: Practical constructions and new proof methods for large universe attribute-based encryption. In: Sadeghi, A.-R., Gligor, V.D., Yung, M. (eds) Proceedings of ACM-CCS 2013, pp. 463–474. ACM (2013)
39. Safavi-Naini, R., Canetti, R. (eds.): CRYPTO 2012. LNCS, vol. 7417. Springer, Heidelberg (2012)
40. Sahai, A., Seyalioglu, H., Waters, B.: Dynamic credentials and ciphertext delegation for attribute-based encryption. In: Safavi-Naini and Canetti [39], pp. 199–217
41. Sahai, A., Waters, B.: Fuzzy identity-based encryption. In: Cramer [11], pp. 457–473
42. Shamir, A.: Identity-based cryptosystems and signature schemes. In: Blakely, G.R., Chaum, D. (eds.) CRYPTO 1984. LNCS, vol. 196, pp. 47–53. Springer, Heidelberg (1985)
43. 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)
44. Wesolowski, B., Junod, P.: Ciphertext-policy attribute-based broadcast encryption scheme with small keys. Cryptology ePrint Archive, Report 2015/836 (2015). http://eprint.iacr.org/2015/836
45. Zhou, Z., Huang, D.: On efficient ciphertext-policy attribute based encryption, broadcast encryption: extended abstract. In: Proceedings of ACM-CCS 2010, pp. 753–755. ACM (2010)