Attribute-Based Broadcast Encryption Scheme Made Efficient

David Lubicz[1,2] and Thomas Sirvent[1,2]

[1] DGA-CELAR, Bruz, France
[2] IRMAR, Université de Rennes 1, France
david.lubicz@univ-rennes1.fr, thomas.sirvent@m4x.org

Abstract. In this paper, we describe a new broadcast encryption scheme for stateless receivers. The main difference between our scheme and the classical ones derived from the complete subtree paradigm is that the group of privileged users is described by attributes. Actually, some real applications have been described where the use of a more adaptable access structure brings more efficiency and ease of deployment. On the other side, the decryption algorithm in so far existing attribute-based encryption schemes adapted for broadcast applications is time-consuming for the receiver, since it entails the computation of a large number of pairings. This is a real drawback for broadcast applications where most of the technological constraints are on the receiver side.

Our scheme can be viewed as a way to benefit at the same time from the performance of decryption of the classical broadcast schemes and the management easiness provided by the use of a more adaptable data structure based on attributes. More precisely, our scheme allows one to select or revoke users by sending ciphertexts of linear size with respect to the number of attributes, which is in general far less than the number of users. We prove that our scheme is fully collusion secure in the generic model of groups with pairing.

Keywords: Public-key broadcast encryption, Attribute-based encryption, Generic model of groups with pairing.

1 Introduction

A broadcast encryption scheme [FN93] is used whenever an emitter wants to send messages to several recipients using an unsecured channel. Such a scheme actually allows the broadcaster to choose dynamically a subset of privileged users inside the set of all possible recipients and to send a ciphertext, readable only by the privileged users. This kind of schemes is helpful in numerous commercial applications such as the broadcast of multimedia content or pay-per-view television.

Many schemes have been suggested to solve this problem regarding two main settings. The first one deals with almost fixed sets of privileged users. In this case the encryption is efficient but modifying the set of privileged users entails the sending of a long message. The second setting is aimed at the management

S. Vaudenay (Ed.): AFRICACRYPT 2008, LNCS 5023, pp. 325–342, 2008.

of very large or very small sets of privileged users. Schemes designed for that purpose allow one to change at no cost the set of privileged users but the size of the encryption grows linearly with the size of the set of revoked users.

In this paper, we consider the real application where an emitter produces different kinds of content for different categories of users. This is a natural problem to deal with for a broadcaster which proposes to its customers several subscription packages, or for different broadcasters using the same asymmetrical broadcast encryption scheme. In this case, it is very possible that the set of privileged users has to be changed dramatically along with the type of content. As this set can not be considered as being particularly small or large, this situation is not covered by usual broadcast encryption schemes.

Recently, a notion of attribute-based encryption has been introduced in [SW05]. This notion seems to address that kind of problem. In [GPSW06], the authors present a declination of these ideas with applications in "targeted" broadcast encryption. In ciphertext-policy schemes, which is our concern here, each user is associated with a set of attributes and its decryption key depends on this set. A ciphertext contains an access policy based on these attributes: only users satisfying this policy may obtain the plaintext, and even a collusion of other users can not obtain it. In broadcast applications, the main drawback of this family of schemes is that the decryption may require large computations which cannot be quickly achieved by low-cost decoders.

Our Contribution. In this paper, we propose a broadcast encryption scheme, with attribute-based mechanisms: it allows the broadcaster to select or to revoke not only single users, but groups of users defined by their attributes. This scheme can be seen as an attribute-based encryption scheme, with efficient decryption and restriction of access policy: the restriction of access policy (using AND and NOT functions) is enough to provide broadcast encryption since the OR function can be simulated using concatenation, exactly like in the Subset-Cover framework.

The idea behind this scheme is the ability to compute a specific greatest common divisor of polynomials. Each receiver is associated with a polynomial (with roots depending on its attributes), and a ciphertext is associated with another polynomial (with roots depending on required attributes and revoked attributes). A receiver in the access policy defined by a ciphertext computes the greatest common divisor of its polynomial and of the polynomial associated with the ciphertext: this divisor is the same for all receivers in the access policy. A receiver not in this access policy can not compute this specific polynomial.

In this scheme, the size of the decryption key given to a receiver is linear in the number of attributes associated with this receiver. The size of a ciphertext is linear in the number of attributes used in the access policy. The public encryption key is quite long: its size is linear in the total number of attributes used in the scheme. This is not a real drawback for realistic situations where anyway the broadcaster must have a database containing the list of users together with their attributes. Moreover, a broadcaster which intends to use only a small set of

attributes requires only an encryption key with a size linear in the size of this small set.

This scheme has a new design, since it is not based on secret sharing like previous attribute-based schemes. This design allows the decryption algorithm to use only a fixed number (3) of pairing computations. As a broadcast encryption scheme, it uses the Subset-Cover framework suggested in [NNL01]. We prove the security of this scheme against full collusions in the generic model of groups with pairings. Another interesting feature in this scheme is that new decryption keys can be built without any modification of previously distributed decryption keys: adding new decryption keys requires only to extend the public key to take new attributes into account.

1.1 Related Work

Stateful Broadcast Schemes. The first broadcast schemes were based upon stateful receivers, which means that the receivers have a memory that can store some information about the past messages. Such receivers have the possibility to refresh their decryption key using information given in broadcasted messages. This is the case of "Logical Key hierarchy" (LKH) presented independently in [WGL98] and in [WHA99]: users have assigned positions as leaves in a tree, and have keys corresponding to nodes on the path from user's leaf to the root. The key corresponding to the root is used to encrypt messages to users. When users are revoked or when a new user joins, a rekey occurs, using keys corresponding to internal nodes. These techniques have been later improved in [CGI+99, CMN99, PST01].

These schemes are aimed at practical applications where the set of privileged users is updated rarely and in a marginal way. The ciphertexts are very short and are computed from a key known by all current users. In return, changing the set of privileged users (add or exclude a user) is bandwidth-consuming and must be done on a per user basis: each change entails the distribution of a new global key to privileged users. Moreover, this can only be done if all users are on-line which is a strong limitation in some applications. The frequent and important changes in the set of privileged users make these schemes inappropriate for the previously mentioned applications.

Stateless Broadcast Schemes.A different kind of broadcast schemes have been introduced later on: the goal is to avoid frequent rekeys. In [KRS99, GSW00], users have different decryption keys, and each decryption key is known by a well-chosen set of users. When the broadcaster wants to exclude a given set of users, it builds ciphertexts corresponding to decryption keys that these specific users do not know. Rekey occurs only after large permanent modifications of the privileged set of users. The ciphertexts are longer than with the LKH schemes mentioned in the previous paragraph.

Stateless receivers extend this last case: in [NNL01], the broadcaster can choose any set of privileged users without any rekey, i.e. the receivers can keep the same decryption keys during the whole life of the broadcast system. These

schemes, called Complete Subtree (CS) and Subset Difference (SD) are based on a binary tree structure, where users are placed in the leaves. They have subsequently been improved in [HS02, GST04], and an efficient extension to the public-key case based on hierarchical identity-based encryption has been proposed in [DF02]. This extension has been confirmed in [BBG05] with the first hierarchical identity-based encryption with constant-size ciphertexts.

The efficiency of these schemes are only proved when few users are revoked, but the binary tree structure presented in [NNL01] and its following improvements may be used to characterize groups of users by attributes: for example, the left subtrees of the internal nodes at a given level may correspond to users with a given attribute, and the right subtrees to users with this attribute missing. This seems doable, even if the tree structure constrains the organization of the attributes (the binary tree must be balanced to keep a good efficiency, so every attribute must concern about half of the users). The Figure 1.1 shows that the selection of users with a given attribute, or the revocation of users without this attribute, is efficient if the attribute corresponds to a high level in the tree, but very inefficient when the attribute is near the leafs. As a consequence, the use of these schemes for selection or revocation of users regarding to their attributes i s not practical, since the size of ciphertexts may be linear in the number of revoked users.

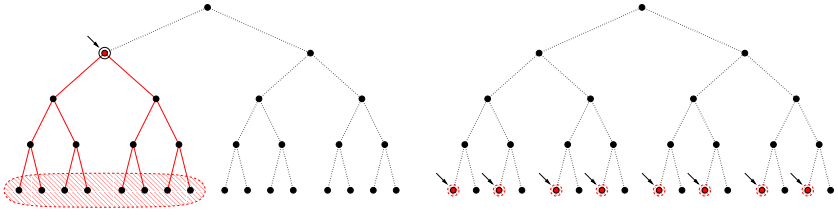

Fig. 1. Selection with CS/SD scheme: first attribute versus last attribute

New public-key broadcast schemes with constant-size ciphertexts have been proposed in [BGW05] (scheme 1) and in [DPP07] (scheme 2). In these schemes, a receiver needs however the exact knowledge of the set of privileged users, which means the transmission of an information with non-constant size, which is not mentioned in the ciphertexts.

These schemes require moreover decryption keys of size linear in the number of users (this is clearly stated in [DPP07]; in [BGW05], a receiver has a constant-size private key, but needs the encryption key to perform a decryption). This storage may be excessive for low-cost devices.

Broadcast Scheme from HIBE with Wildcards. Management of attributes can be performed by the combination of the scheme given in [DF02] with a hierarchical identity-based encryption scheme with wildcards, like presented in [ACD+06, BDNS06]. The resulting scheme would allow the selection of users with given attributes, i.e. build ciphertexts addressed to intersections of groups.

The revocation of all users with a fixed attribute from the SD technique is however unclear, and its use is not efficient since the size of the ciphertexts is not constant in the hierarchical identity based encryption (see [BBG05]).

Attribute-Based Encryption. Attribute-based encryption has been suggested in [SW05], and later developed in [GPSW06]. In a first version (later called key-policy attribute-based encryption), the goal is to define access policies, and to allow a user to obtain some information if the access policy associated with this user is valid for this content. In this way, the decryption key given to a user depends on an access policy, and the encryption of a content relies on attributes, which are used in the evaluation of an access policy. Even a collusion of users with invalid access policies for a given ciphertext should not be able to obtain the corresponding plaintext.

Later, in [BSW07], a new scheme is proposed, but with an inversion: the access policy is defined with the content, and attributes are used to build decryption keys given to users. These ciphertext-policy attribute-based encryption schemes have direct applications for broadcast: the access policy defines a set of privileged users. With a relevant distribution of attributes, any set of privileged users may be described by an access policy.

In these schemes, an access policy is build using secret sharing techniques, like Shamir's one based on polynomials. An access policy is defined by a tree, where leaves correspond to the presence of an attribute (the evaluation of a leaf is true if the corresponding attribute is used) and internal nodes are threshold functions (in particular, these nodes may be AND, or OR functions). With such structure, revocation is quite difficult, since adding attributes can only provide a larger access to the content.

This problem is solved in [OSW07], where the access policy may be non-monotonic: the use of NOT functions becomes possible. Combining results from [BSW07, OSW07] gives rise to a ciphertext-policy attribute-based encryption which can be used for broadcast applications. The design of these schemes requires however a receiver to perform a large number of pairing computations (linear in the number of attributes used in the access policy). A low-cost receiver may not be able to compute so much pairings in complex access policies.

Our scheme has a completely different design, and it allows only very specific access policies. An access policy in this scheme is a disjunction (OR function, using the Subset-Cover framework) of conjunctions (AND functions) of attributes and of negations of attributes. Such access policy is more restrictive, but it is enough for practical broadcast applications. In return, a receiver performs only 3 pairing computations whatever the access policy is.

Dynamic Broadcast Encryption Scheme. The notion of dynamic schemes has been defined in [DPP07]. In such schemes, new users can be added without modification of previously distributed decryption keys. The encryption key has only to be slightly extended. This feature seems to be very useful in practical applications. The dynamic schemes suggested in [DPP07] requires ciphertexts of size linear in the number of revoked users. This feature is quite rare in broadcast schemes, but common in attribute-based encryption schemes.

1.2 Organization

The paper is organized as follows. In Section 2, we give a formal definition of groups of users, and an associated definition of attribute-based broadcast encryption schemes. In Section 3, we describe our scheme and prove its correctness. In Section 4, we prove the security of this scheme.

2 Preliminaries

We give a formal definition of groups of users and an associated definition of attribute-based broadcast encryption schemes deduced from the definition given in [BGW05]. We present then the security model. The last part explains how to define groups of users in concrete applications.

2.1 Groups of Users

In our applications, we have a large number of users, and a large number of groups (in practice, we need for each user a group containing this single user). Each user belongs to a few groups of users. We choose a description which takes advantage of this fact.

Let $\mathcal{U}$ be the set of all users. We represent an element of $\mathcal{U}$ by an integer in $\{1, \ldots, n\}$. A group of users is a subset $\mathcal{G}$ of $\mathcal{U}$. From the inverse point of view, for a fixed number l of groups of users, we can associate with a user $u \in \mathcal{U}$ the set of groups he belongs to: $\mathcal{B}(u) = \{i \in \{1, \ldots, l\} \,/\, u \in \mathcal{G}_i\} \subset \{1, \ldots, l\}$.

2.2 Attribute-Based Broadcast Encryption Schemes

In this part, we give a formal definition of an attribute-based broadcast encryption scheme. This model does not take into account the fact that the scheme could be dynamic, like in [DPP07], even if our scheme seems dynamic. The following definitions are just a slight adaptation of [BGW05, BSW07] to deal with groups of users.

A public-key attribute-based broadcast encryption scheme with security parameter λ is a tuple of three randomized algorithms:

- ***Setup***$(\lambda, n, (\mathcal{B}(u))_{1 \leq u \leq n})$: takes as input the security parameter λ, the number of users n, and groups of users. It outputs an encryption key EK, and n decryption keys $(\mathrm{dk}_u)_{1 \leq u \leq n}$.
- ***Encrypt***$(\mathrm{EK}, \mathcal{B}^N, \mathcal{B}^R)$: takes as input the encryption key EK and two sets of groups $\mathcal{B}^N$ and $\mathcal{B}^R$. It outputs a header hdr and a message encryption key $K \in \mathcal{K}$, where $\mathcal{K}$ is a finite set of message encryption keys.
- ***Decrypt***$(\mathrm{dk}_u, \mathrm{hdr})$: takes as input a decryption key given to a user u and a header hdr. If the header hdr comes from an encryption using $(\mathcal{B}^N, \mathcal{B}^R)$ such that $\mathcal{B}^N \subset \mathcal{B}(u)$ and $\mathcal{B}(u) \cap \mathcal{B}^R = \emptyset$, then it outputs a message encryption key $K \in \mathcal{K}$. In the other case, it outputs $\perp$.

In the encryption process, a message M is encrypted with a key K and the resulting ciphertext C is sent together with the header hdr. Users in all groups mentioned in $\mathcal{B}^N$ (needed groups) and outside all groups mentioned in $\mathcal{B}^R$ (revoked groups) can compute K from the header hdr and their decryption key dk_u. Using the key K, a user recovers M from C.

Note that in these definitions, the decryption key and the header are the only elements that a user needs in the computation of the key K. The encryption key and the knowledge of the set of privileged users are not necessary for decryption. The header corresponds then exactly to the cost of the broadcast scheme in terms of transmission. In fact, in our scheme, the knowledge of the set of privileged users is implicitly included in the header, encoded in the attributes corresponding to the required and revoked groups.

In this description, we do not allow an encryption for an arbitrary set of privileged users, which is the usual definition of a broadcast encryption scheme. Any set of privileged users can however be represented by a union of sets used in this "basic encryption" for well-chosen groups of users (in fact, it is enough that each user belongs to a group containing only this single user). Different basic encryptions are then used to encrypt a common key, instead of a message. The full message can then be sent, using this common key.

2.3 Security Model

We consider semantic security of attribute-based broadcast encryption schemes. The adversary is assumed static, as in previous models: the only difference with standard definitions is that the groups of users are given to the adversary before the beginning of the game played by the challenger and the adversary $\mathcal{A}$:

- The challenger and the adversary are given l fixed groups of users, defined by $(\mathcal{B}(u))_{1\leq u\leq n}$.
- The adversary $\mathcal{A}$ outputs two sets of groups $\mathcal{B}^N$ and $\mathcal{B}^R$ corresponding to a configuration it intends to attack.
- The challenger runs $\mathit{Setup}(\lambda, n, (\mathcal{B}(u))_{1\leq u\leq n})$ and gives to $\mathcal{A}$ the encryption key EK, and the decryption keys dk_u corresponding to users that the adversary may control, i.e. such that $\mathcal{B}^N \cap \mathcal{B}(u) \neq \mathcal{B}^N$ or $\mathcal{B}^R \cap \mathcal{B}(u) \neq \emptyset$.
- The challenger runs $\mathit{Encrypt}(\mathrm{EK}, \mathcal{B}^N, \mathcal{B}^R)$, and obtains a header hdr and a key $K \in \mathcal{K}$. Next, the challenger draws a random bit b, sets $K_b = K$, picks up randomly K_{1-b} in $\mathcal{K}$, and gives (hdr, K_0, K_1) to the adversary $\mathcal{A}$.
- The adversary $\mathcal{A}$ outputs a bit b'.

The adversary $\mathcal{A}$ wins the previous game when $b' = b$. The advantage of $\mathcal{A}$ in this game, with parameters $(\lambda, n, (\mathcal{B}(u))_{1\leq u\leq n})$, is $|2\Pr[b' = b] - 1|$, where the probability is taken over the choices of b and all the random bits used in the simulation of the *Setup* and *Encrypt* algorithms:

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

An attribute-based broadcast encryption scheme is semantically secure against full static collusions if for all randomized polynomial-time (in λ) adversary $\mathcal{A}$ and

for all groups of users $(\mathcal{B}(u))_{1\leq u\leq n}$ with at most l groups, $\mathrm{Adv}^{\mathrm{ind}}(\lambda, n, (\mathcal{B}(u)), \mathcal{A})$ is a negligible function in λ when n and l are at most polynomials in λ.

From such semantically secure schemes, we can build schemes secure in a stronger model: the use of generic transformations, like the ones presented in [FO99a, FO99b, OP01] has a negligible cost, and we obtain chosen-ciphertext security in the random oracle model. This explains why our security model is limited to chosen-plaintexts attacks.

2.4 Well-Chosen Groups of Users

In real broadcast applications, one has often to deal with obvious groups of users, because users are classified for instance by subscription package or subscription period. These groups are easily managed by an attribute-based broadcast encryption scheme, by simply using one attribute for each obvious group of users.

In some circumstances, it may happen that the group of privileged users does not fit easily with a description based on these obvious groups of users. Even if rare, it is preferable to be able to deal with such situations.

A solution consists in adding some extra attributes to the set of attributes corresponding to obvious groups. These new attributes describe a binary tree structure over the users, and allows the same management of users as in the SD-scheme. More precisely, we place users in the leaves of a binary tree, each node corresponds to a new attribute and each user receives the attributes of its parent nodes. At most $2n$ new attributes are added, and a user belongs to at most $\lceil \log_2(n) \rceil + 1$ new groups.

With this setting, there is an attribute for each user and this simple fact guarantees that any subset of users can be described by attributes. Moreover, basic encryption with privileged users corresponding to members of one group, excluding members of another group give at least the same sets as in the SD-method presented in [NNL01]. The efficiency of the attribute-based broadcast encryption scheme is then at least as good as in the SD-method, for any set of privileged users.

3 Construction

In this section, we first present bilinear maps. We describe next the *Setup*, *Encrypt* and *Decrypt* algorithms of a public-key attribute-based broadcast encryption scheme based on groups with a bilinear map. The correctness can then be verified.

3.1 Bilinear Maps

In the following definitions, we consider the symmetric setting of bilinear maps, like in [Jou00, BF01]. Let $\mathbb{G}_1$ and $\mathbb{G}_2$ be two cyclic groups of prime order p. The group laws in $\mathbb{G}_1$ and $\mathbb{G}_2$ are noted additively. Let g_1 be a generator of $\mathbb{G}_1$. Let $e : \mathbb{G}_1 \times \mathbb{G}_1 \to \mathbb{G}_2$ be a non-degenerate pairing:

– for all $a, b \in (\mathbb{Z}/p\mathbb{Z})$, $e(a\,g_1, b\,g_1) = ab.e(g_1, g_1)$,
– let $g_2 = e(g_1, g_1)$, g_2 is a generator of $\mathbb{G}_2$.

We make the assumption that the group laws in $\mathbb{G}_1$ and $\mathbb{G}_2$, and the bilinear map e can be computed efficiently.

3.2 Setup Algorithm

From the security parameter λ, the first step of the setup consists in constructing a tuple $(\mathbb{G}_1, \mathbb{G}_2, g_1, g_2, e, p)$, where:

– p is a prime, the length of which is λ,
– $\mathbb{G}_1$ and $\mathbb{G}_2$ are two cyclic groups of prime order p,
– e is a non-degenerate pairing from $\mathbb{G}_1 \times \mathbb{G}_1$ into $\mathbb{G}_2$,
– g_1 is a generator of $\mathbb{G}_1$ and $g_2 = e(g_1, g_1)$.

Four elements (α, β, γ and δ) are randomly chosen in $(\mathbb{Z}/p\mathbb{Z})^*$. Each group of users $\mathcal{G}_i$, mentioned in $(\mathcal{B}(u))_{1\leq u\leq n}$ is then associated with an attribute μ_i randomly chosen in $(\mathbb{Z}/p\mathbb{Z})$, such that all these attributes are pairwise different and different from α. Another attribute μ_0 is chosen with the same constraints, corresponding to a virtual group containing no users. The encryption key is:

$$\mathrm{EK} = \Big(g_1\,,\, \beta\,\gamma\,\delta\,g_1\,,\, (\mu_i)_{0\leq i\leq l}\,,\, \big(\alpha^i\,g_1\big)_{0\leq i\leq l}\,,\, \big(\alpha^i\,\gamma\,g_1\big)_{0\leq i\leq l}\,,\, \big(\alpha^i\,\delta\,g_1\big)_{0\leq i\leq l}\Big).$$

For each user $u \in \mathcal{U}$, s_u is randomly chosen in $(\mathbb{Z}/p\mathbb{Z})^*$. Let $\Omega(u)$ be the set of attributes corresponding to the groups he belongs to: $\Omega(u) = \{\mu_i \in (\mathbb{Z}/p\mathbb{Z}) \,/\, i \in \mathcal{B}(u)\}$. Let $l(u)$ be the size of $\Omega(u)$, i.e. the number of groups containing u. Let $\Pi(u) = \prod_{\mu\in\Omega(u)}(\alpha - \mu)$. The decryption key of u is:

$$\mathrm{dk}_u = \Big(\Omega(u), (\beta + s_u)\,\delta\,g_1\,,\, \gamma\,s_u\,\Pi(u)\,g_1\,,\, \big(\alpha^i\,\gamma\,\delta\,s_u\,g_1\big)_{0\leq i<l(u)}\Big).$$

3.3 Encryption Algorithm

If $\mathcal{B}^N \cap \mathcal{B}^R \neq \emptyset$, the encryption algorithm aborts and returns $\perp$, since a user can not be simultaneously inside and outside a given group of users. Otherwise, let $\Omega^N = \{\mu_i \,/\, i \in \mathcal{B}^N\}$ and $\Omega^R = \{\mu_i \,/\, i \in \mathcal{B}^R\}$. Let $l^N = |\mathcal{B}^N|$ be the number of required groups and $l^R = |\mathcal{B}^R|$ be the number of revoked groups[1]. Let $\Pi^N = \prod_{\mu\in\Omega^N}(\alpha - \mu)$, let $\Pi^R = \prod_{\mu\in\Omega^R}(\alpha - \mu)$ and let $\Pi^{NR} = \Pi^N\Pi^R$. Let z be randomly chosen in $(\mathbb{Z}/p\mathbb{Z})^*$. The result of the encryption is:

$$\mathrm{hdr} = \Big(\Omega^N, \Omega^R, z\,\Pi^{NR}\,g_1\,,\, \gamma\,z\,\Pi^N\,g_1\,,\, \big(\alpha^i\,\delta\,z\,g_1\big)_{0\leq i<l^R}\Big),\ K = \beta\,\gamma\,\delta\,z\,\Pi^N\,g_2.$$

All these elements can be computed using only the encryption key EK.

[1] A slight modification occurs when $\mathcal{B}^R$ is empty: in such case, the encryption considers that the virtual group containing no users is revoked and then $\Omega^R = \{\mu_0\}$, $l^R = 1$.

3.4 Decryption Algorithm

We consider here the decryption of a header hdr with a decryption key dk_u:

$$\begin{cases} \mathrm{dk}_u = (\Omega(u), \mathrm{dk}_1, \mathrm{dk}_2, \mathrm{dk}_{3,0}, \ldots, \mathrm{dk}_{3,l(u)-1}) , \\ \mathrm{hdr} = (\Omega^N, \Omega^R, \mathrm{hdr}_1, \mathrm{hdr}_2, \mathrm{hdr}_{3,0}, \ldots, \mathrm{hdr}_{3,l^R-1}) . \end{cases}$$

The receiver u is valid for this header if $\Omega(u)$ contains Ω^N and if the intersection between Ω^R and $\Omega(u)$ is empty. To decrypt the header, the valid receiver u uses the extended Euclidean algorithm over the polynomials $\prod_{\mu\in(\Omega^N\cup\Omega^R)}(X-\mu)$ and $\prod_{\mu\in\Omega(u)}(X-\mu)$. It obtains two unitary polynomials, $V(X)=\sum_{0\le i<l(u)} v_i X^i$ and $W(X)=\sum_{0\le i<l^R} w_i X^i$, in $(\mathbb{Z}/p\mathbb{Z})[X]$, such that:

$$V(X) \prod_{\mu\in(\Omega^N\cup\Omega^R)} (X-\mu) + W(X) \prod_{\mu\in\Omega(u)} (X-\mu) = \prod_{\mu\in\Omega^N} (X-\mu).$$

From these polynomials, the receiver computes the key:

$$K(\mathrm{dk}_u, \mathrm{hdr}) = e(\mathrm{dk}_1, \mathrm{hdr}_2) - e\left(\sum_{i=0}^{l(u)-1} v_i \, \mathrm{dk}_{3,i} \, , \, \mathrm{hdr}_1 \right) - e\left(\mathrm{dk}_2 \, , \, \sum_{i=0}^{l^R-1} w_i \, \mathrm{hdr}_{3,i} \right).$$

3.5 Proof of Correctness

If dk_u is the valid decryption key given to a user u, if hdr is a header built using the encryption and if u is a valid user for hdr, then the decryption gives:

$$K(\mathrm{dk}_u, \mathrm{hdr}) = (\beta + s_u)\, \gamma\, \delta\, z\, \Pi^N\, g_2 - \gamma\, \delta\, z\, s_u\, V(\alpha)\, \Pi^{NR}\, g_2 - \gamma\, \delta\, z\, s_u\, W(\alpha)\, \Pi(u)\, g_2.$$

By definition of the two polynomials V and W, we have the following relation: $V(\alpha)\,\Pi^{NR} + W(\alpha)\Pi(u) = \Pi^N$. The computed key is then exactly the key associated with the header in the encryption:

$$K(\mathrm{dk}_u, \mathrm{hdr}) = (\beta + s_u)\, \gamma\, \delta\, z\, \Pi^N\, g_2 - \gamma\, \delta\, z\, s_u\, \Pi^N\, g_2 = \beta\, \gamma\, \delta\, z\, \Pi^N\, g_2.$$

4 Security of the Scheme

The previous scheme can be proved in different ways. The usual strategy is first to define some security assumption and to prove this assumption in the generic model of groups with pairing. The reduction of the security of the scheme to this assumption concludes the proof. Following this strategy, we need a new security assumption which is an extension of the decisional version of the General Diffie-Hellman Exponent (GDHE) problem, precisely studied in the full version of [BBG05]. For the sake of simplicity, we prefer here a more direct proof in the generic model of groups with pairing.

In this section, we define the decisional problem upon which our broadcast encryption mechanism is built. We assess its security in the framework of the generic model of groups with pairing.

4.1 A Decisional Problem

Let $\mathbb{G}_1$ and $\mathbb{G}_2$ be two cyclic groups of prime order p and e be a non-degenerate pairing from $\mathbb{G}_1 \times \mathbb{G}_1$ into $\mathbb{G}_2$. Let g_1 be a generator of $\mathbb{G}_1$ and $g_2 = e(g_1, g_1)$. Let α, β, γ, δ, z be elements of $(\mathbb{Z}/p\mathbb{Z})^*$. For all $i \in \{0, \ldots, l\}$, let μ_i be an element of $(\mathbb{Z}/p\mathbb{Z})$ different from α and from μ_j where $j < i$.

The encryption key is:

$$\mathrm{EK} = \Big(g_1 \,,\, \beta\,\gamma\,\delta\, g_1 \,,\, (\mu_i)_{0\le i\le l} \,,\, \big(\alpha^i\, g_1\big)_{0\le i\le l} \,,\, \big(\alpha^i\,\gamma\, g_1\big)_{0\le i\le l} \,,\, \big(\alpha^i\,\delta\, g_1\big)_{0\le i\le l} \Big).$$

For each user $u \in \mathcal{U}$, $\Omega(u)$ is a subset of $\{\mu_1, \ldots, \mu_l\}$. Let $l(u) = |\Omega(u)|$ and let $\Pi(u) = \prod_{\mu\in\Omega(u)}(\alpha - \mu)$. The decryption key dk_u of the user u is:

$$\mathrm{dk}_u = \Big(\Omega(u), (\beta + s_u)\,\delta\, g_1 \,,\, \gamma\, s_u\, \Pi(u)\, g_1 \,,\, \big(\alpha^i\,\gamma\,\delta\, s_u\, g_1\big)_{0\le i<l(u)} \Big).$$

Let Ω^N be a subset of $\{\mu_1, \ldots, \mu_l\}$, let Ω^R be a non-empty subset of $\{\mu_0, \ldots, \mu_l\}$ such that $\Omega^N \cap \Omega^R = \emptyset$, let $l^R = |\Omega^R|$. Let $\mathcal{R}$ be the set of revoked users for these sets:

$$\mathcal{R} = \big\{ u \in \mathcal{U} \,/\, \Omega(u) \cap \Omega^N \neq \Omega^N \ \text{ or } \ \Omega(u) \cap \Omega^R \neq \emptyset \big\}.$$

Let $\Pi^N = \prod_{\mu\in\Omega^N}(\alpha - \mu)$, let $\Pi^R = \prod_{\mu\in\Omega^R}(\alpha - \mu)$ and let $\Pi^{NR} = \Pi^N \Pi^R$. The header hdr and the key K are defined by:

$$\mathrm{hdr} = \Big(\Omega^N, \Omega^R, z\,\Pi^{NR}\, g_1 \,,\, \gamma\, z\, \Pi^N\, g_1 \,,\, \big(\alpha^i\,\delta\, z\, g_1\big)_{0\le i<l^R} \Big), \ K = \beta\,\gamma\,\delta\, z\, \Pi^N\, g_2.$$

Let b be a bit, let K_{1-b} be an element of $(\mathbb{Z}/p\mathbb{Z})^*$, let $K_b = K$. The decisional problem is the following: guess b from the knowledge of EK, hdr, K_0, K_1 and all the dk_u, where $u \in \mathcal{R}$.

4.2 Interpretation in the Generic Model

In this section, we use the notations of the full version of [BBG05] in order to assess the difficulty of the preceding decisional problem in the generic model of groups with pairing model. This extends the classical model of generic groups presented in [Nec93, Sho97].

The first part of the proof consists in showing that there exists no formula giving the key from the header, the encryption key, and the decryption keys corresponding to revoked users. The second part details why an adversary can not distinguish the key from a random element in the generic model of groups with pairing.

No Formula. Let $\mathcal{P}$ be the ring of polynomials over the variables A, B, C, D, Z and $\{S_u, u \in \mathcal{R}\}$. Each of these variables represent an element picked at random in the decisional problem and not explicitly unveiled: A is used for α, B for β, C for γ, D for δ, Z for z and for all $u \in \mathcal{U}$, S_u is used for s_u.

Let $\mathcal{D}$ be the tuple of elements in $\mathcal{P}$, corresponding to the discrete logarithms of elements in $\mathbb{G}_1$ given to an adversary in the problem. The tuple $\mathcal{D}$ contains 1, $B\,C\,D$, $Z\,\Pi^{NR}(A)$, $C\,Z\,\Pi^{N}(A)$ and the following polynomials:

- A^i, $A^i\,C$ and $A^i\,D$ for all $i \in \{0,\ldots,l\}$,
- $(B + S_u)\,D$ and $C\,S_u\,\Pi_u(A)$, for all $u \in \mathcal{R}$,
- $A^i\,C\,D\,S_u$, for all $u \in \mathcal{R}$ and $i \in \{0,\ldots,l(u)-1\}$,
- $A^i\,D\,Z$ for all $i \in \{0,\ldots,l^R-1\}$,

where

$$\Pi^N(A) = \prod_{\mu\in\Omega^N}(A-\mu), \quad \Pi^R(A) = \prod_{\mu\in\Omega^R}(A-\mu),$$

$$\Pi_u(A) = \prod_{\mu\in\Omega(u)}(A-\mu), \quad \Pi^{NR}(A) = \Pi^N(A)\,\Pi^R(A).$$

Lemma 1. *Let $\mathcal{M}$ be the sub-$\mathbb{Z}$-module of $\mathcal{P}$ generated by all products of elements of $\mathcal{D}$. If $l^R \leq \sqrt{p}/2$ and for all u, $l(u) \leq \sqrt{p}/2$, the element $B\,C\,D\,Z\,\Pi^N(A)$ is an element of $\mathcal{M}$ with probability less than $1/\sqrt{p}$, this last probability being taken over all possible choices of the attributes μ_i in $(\mathbb{Z}/p\mathbb{Z})$.*

Proof. This lemma is proved in appendix A.1.

Indistinguishability in the Generic Model. In the generic model of groups with pairing, we consider two injective maps ξ_1 and ξ_2 from $(\mathbb{Z}/p\mathbb{Z})$ into $\{0,1\}^*$, also known as encoding functions. The additive law on $(\mathbb{Z}/p\mathbb{Z})$ induces a group law over $\xi_1(\mathbb{Z}/p\mathbb{Z})$ and $\xi_2(\mathbb{Z}/p\mathbb{Z})$, and the sets $\xi_1(\mathbb{Z}/p\mathbb{Z})$ and $\xi_2(\mathbb{Z}/p\mathbb{Z})$ together with these group laws are respectively denoted by $\mathbb{G}_1$ and $\mathbb{G}_2$. Oracles corresponding to the group law and the inverse law of each group are provided. A new law, corresponding to the pairing, is also given as an oracle: for all $x, y \in \mathbb{G}_1, e(x,y) = \xi_2({\xi_1}^{-1}(x) \times {\xi_1}^{-1}(y)) \in \mathbb{G}_2$. An algorithm computing in this model has only access to these 5 oracles, and has no information about ξ_1 and ξ_2: its computations are based on queries to these oracles.

In our case, this model means that a challenger will use randomly chosen encoding functions from $(\mathbb{Z}/p\mathbb{Z})$ into a set of p binary strings. The challenger randomly chooses α, β, γ, δ, z, $(\mu_i)_{0\leq i\leq l}$, $(s_u)_{u\in\mathcal{U}}$ following their constraints, and gives to the adversary all values $\xi_1\left(f(\alpha,\beta,\gamma,\delta,z,s_1,\ldots,s_n)\right)$, where f is in the tuple $\mathcal{D}$. The adversary receives moreover $\xi_2(\kappa_0)$ and $\xi_2(\kappa_1)$,where κ_{1-b} is chosen randomly in $(\mathbb{Z}/p\mathbb{Z})^*$ and $\kappa_b = \beta\,\gamma\,\delta\,z\Pi^N$. The adversary makes then queries to oracles and finally outputs its guess b'.

We use the following theorem, proposed and proved in the full version of [BBG05] (Theorem A.2):

Theorem 1. *Let $\mathcal{D}$ be a subset of $\mathcal{P}$ of size k and suppose that for all $f \in \mathcal{D}$, $\deg(f) \leq d$. Let ϕ be an element of $\mathcal{P}$ such that ϕ is not is the sub-$\mathbb{Z}$-module spanned by the products of any two elements of $\mathcal{D}$. We consider an adversary which receives the set $\{\xi_1\left(f(\alpha,\beta,\gamma,\delta,z,s_1,\ldots,s_n)\right) / f \in \mathcal{D}\}$, $\xi_2(\kappa_0)$ and $\xi_2(\kappa_1)$,*

where κ_{1-b} *is chosen randomly in* $(\mathbb{Z}/p\mathbb{Z})^*$ *and* $\kappa_b = \phi(\alpha, \beta, \gamma, \delta, z, s_1, \ldots, s_n)$. *All such adversary which is allowed to issue at most* q *queries to the oracles can not guess the bit* b *with a probability significantly better than* $1/2$ *:*

$$\left| Pr[b' = b] - \frac{1}{2} \right| \leq \frac{\max(2d, \deg(\phi))\,(q + 2k + 2)^2}{2p}.$$

In our context, the set $\mathcal{D}$ contains at most $nl + 3(n+l) + 7$ elements. Moreover these elements have degree less than $l+2$ and the degree of $\phi = B\,C\,D\,R\,\Pi^N(A)$ is less than $l + 4$. If ϕ is not in the span generated by the products of any two elements of $\mathcal{D}$, this lemma implies:

$$\left| Pr[b' = b] - \frac{1}{2} \right| \leq \frac{(l + 2)\,(q + 2nl + 6n + 6l + 14)^2}{p}.$$

The results of Lemma 1 and Theorem 1 give the following theorem:

Theorem 2. *In the generic model of groups with pairing, the advantage of an adversary for the problem defined in Part 2.3 of the attribute-based broadcast encryption scheme presented in Section 3, issuing at most* q *queries to the oracles is bounded by:*

$$\frac{(l + 2)\,(q + 2nl + 6n + 6l + 14)^2}{p - \sqrt{p}},$$

where n *is the number of users and* l *is the number of groups of users.*

Proof. We only have to divide the maximum probability obtained by the Theorem 1 by the factor $1 - 1/\sqrt{p}$ which is a lower bound for the probability that the polynomial ϕ is not in the sub-$\mathbb{Z}$-module generated by products of elements of $\mathcal{D}$ which is a consequence of the Lemma 1. The condition on the degrees in the Lemma 1 is verified, l being polynomial in the security parameter λ whereas p is exponential in this same parameter.

The arguments that n, q and l are at most polynomials in the security parameter λ, whereas p is exponential in λ, yield moreover that the given bound is a negligible function of the security parameter. This concludes the proof of security of our attribute-based broadcast encryption scheme.

5 Conclusion

In this paper, we have built a new public-key broadcast encryption scheme especially interesting when dealing with groups of users defined by the conjunction and exclusion of some attributes. We have described a practical application where none of previously existing broadcast or attribute-based encryption schemes behave in a suitable manner.

We have given a generic way to use attributes in order to manage groups of users in an efficient way. Finally, we have proved that our scheme is semantically secure against full static collusions in the generic model of groups with pairing.

It would be interesting to investigate the possibility to improve the access structure of our scheme by implementing efficiently the OR, or a threshold functionality. We also believe that the underlying problem of our scheme, based upon the reconstruction of the greatest common divisor of polynomials, may have some other interesting applications.

Acknowledgments. The authors would like to thank Cécile Delerablée for helpful comments on earlier drafts of this paper.

References

[ACD+06] Abdalla, M., Catalano, D., Dent, A.W., Malone-Lee, J., Neven, G., Smart, N.P.: Identity-based encryption gone wild. In: Bugliesi, M., Preneel, B., Sassone, V., Wegener, I. (eds.) ICALP 2006. LNCS, vol. 4052, pp. 300–311. Springer, Heidelberg (2006)

[BBG05] Boneh, D., Boyen, X., Goh, E.-J.: Hierarchical identity based encryption with constant size ciphertext. In: Cramer, R.J.F. (ed.) EUROCRYPT 2005. LNCS, vol. 3494, pp. 440–456. Springer, Heidelberg (2005)

[BDNS06] Birkett, J., Dent, A.W., Neven, G., Schuldt, J.: Efficient chosen-ciphertext secure identity-based encryption with wildcards. Technical Report 2006/377, Cryptology ePrint Archive (2006)

[BF01] Boneh, D., Franklin, M.: Identity-based encryption from the Weil pairing. In: Kilian, J. (ed.) CRYPTO 2001. LNCS, vol. 2139, pp. 213–229. Springer, Heidelberg (2001)

[BGW05] 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)

[BSW07] Bethencourt, J., Sahai, A., Waters, B.: Ciphertext-policy attribute-based encryption. In: Proc. of IEEE Symposium on Security and Privacy, pp. 321–334 (2007)

[CGI+99] Canetti, R., Garay, J., Itkis, G., Micciancio, D., Naor, M., Pinkas, B.: Multicast security: A taxonomy and efficient constructions. In: IEEE Infocom 1999, vol. 2, pp. 708–716 (1999)

[CMN99] Canetti, R., Malkin, T., Nissim, K.: Efficient communication-storage tradeoffs for multicast encryption. In: Stern, J. (ed.) EUROCRYPT 1999. LNCS, vol. 1592, pp. 459–474. Springer, Heidelberg (1999)

[DF02] 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)

[DPP07] Delerablee, C., Paillier, P., Pointcheval, D.: Fully collusion secure dynamic broadcast encryption with constant-size ciphertexts and decryption keys. Technical report, Prepublication accepted in Pairing 2007 (2007)

[FN93] Fiat, A., Naor, M.: Broadcast encryption. In: Stinson, D.R. (ed.) CRYPTO 1993. LNCS, vol. 773, pp. 480–491. Springer, Heidelberg (1994)

[FO99a] Fujisaki, E., Okamoto, T.: How to enhance the security of public-key encryption at minimum cost. In: Imai, H., Zheng, Y. (eds.) PKC 1999. LNCS, vol. 1560, pp. 53–68. Springer, Heidelberg (1999)

[FO99b] Fujisaki, E., Okamoto, T.: Secure integration of asymmetric and symmetric encryption schemes. In: Wiener, M.J. (ed.) CRYPTO 1999. LNCS, vol. 1666, pp. 537–554. Springer, Heidelberg (1999)

[GPSW06] Goyal, V., Pandey, O., Sahai, A., Waters, B.: Attribute-based encryption for fine-grained access control of encrypted data. In: Proc. of ACM-CCS 2006, pp. 89–98 (2006)

[GST04] 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)

[GSW00] Garay, J.A., Staddon, J., Wool, A.: Long-lived broadcast encryption. In: Bellare, M. (ed.) CRYPTO 2000. LNCS, vol. 1880, pp. 333–352. Springer, Heidelberg (2000)

[HS02] Halevy, D., Shamir, A.: The LSD broadcast encryption scheme. In: Yung, M. (ed.) CRYPTO 2002. LNCS, vol. 2442, pp. 47–60. Springer, Heidelberg (2002)

[Jou00] Joux, A.: A one round protocol for tripartite Diffie-Hellman. In: Bosma, W. (ed.) ANTS 2000. LNCS, vol. 1838, pp. 385–393. Springer, Heidelberg (2000)

[KRS99] Kumar, R., Rajagopalan, S., Sahai, A.: Coding constructions for blacklisting problems without computational assumptions. In: Wiener, M.J. (ed.) CRYPTO 1999. LNCS, vol. 1666, pp. 609–623. Springer, Heidelberg (1999)

[Nec93] Nechaev, V.I.: Complexity of a determinate algorithm for the discrete logarithm. Mathematicheskie Zametki 55(2), 91–101 (1993)

[NNL01] Naor, M., Naor, D., Lotspiech, J.: Revocation and tracing schemes for stateless receivers. In: Kilian, J. (ed.) CRYPTO 2001. LNCS, vol. 2139, pp. 41–62. Springer, Heidelberg (2001)

[OP01] Okamoto, T., Pointcheval, D.: React: Rapid enhanced-security asymmetric cryptosystem transform. In: Naccache, D. (ed.) CT-RSA 2001. LNCS, vol. 2020, pp. 159–175. Springer, Heidelberg (2001)

[OSW07] Ostrovsky, R., Sahai, A., Waters, B.: Attribute-based encryption with non-monotonic access structures. In: Proc. of ACM-CCS 2007, pp. 195–203 (2007)

[PST01] Perrig, A., Song, D., Tygar, J.D.: Elk, a new protocol for efficient large-group key distribution. In: Proc. of IEEE Symposium on Security and Privacy, pp. 247–262 (2001)

[Sch80] Schwartz, J.T.: Fast probabilistic algorithms for verification of polynomial identities. J. Assoc. Comput. Mach. 27(4), 701–717 (1980)

[Sho97] Shoup, V.: Lower bounds for discrete logarithms and related problems. In: Fumy, W. (ed.) EUROCRYPT 1997. LNCS, vol. 1233, pp. 256–266. Springer, Heidelberg (1997)

[SW05] Sahai, A., Waters, B.: Fuzzy identity-based encryption. In: Cramer, R.J.F. (ed.) EUROCRYPT 2005. LNCS, vol. 3494, pp. 457–473. Springer, Heidelberg (2005)

[WGL98] Wong, C.K., Gouda, M., Lam, S.S.: Secure group communications using key graphs. In: Proc. of ACM-SIGCOMM 1998, pp. 68–79 (1998)

[WHA99] Wallner, D.M., Harder, E.J., Agee, R.C.: Key management for multicast: Issues and architectures. RFC 2627 (1999)

A Proof of Lemma 1

A.1 Proof of Lemma 1

Let $\mathcal{D}'$ be the set of elements of $\mathcal{P}$ which are products of pairs of elements of $\mathcal{D}$. By definition, $\mathcal{D}'$ generates $\mathcal{M}$. Suppose that $B\,C\,D\,Z\,\Pi^N(A) \in \mathcal{M}$. Then

it is a linear combination with coefficients in $\mathbb{Z}$ of elements of $\mathcal{D}'$. Considering the elements of $\mathcal{D}'$ as polynomials with respect to the variable C, we see that $B\,C\,D\,Z\,\Pi^N(A)$ can only be obtained as a linear combination of terms linear in the variable C. In the same way, it can only be obtained as a linear combination of terms linear in the variables D and Z.

All elements in $\mathcal{P}$ are homogeneous of degree 0 or 1 in the set of variables $\{B\} \cup \{S_u, u \in \mathcal{R}\}$. Elements in $\mathcal{D}'$ are then homogeneous of degree 0, 1 or 2 in the same set of variables, and the polynomial $B\,C\,D\,Z\,\Pi^N(A)$ can only be obtained as a linear combination of homogeneous terms of degree 1.

These terms of $\mathcal{D}'$ which are simultaneously linear in the variables C, D and Z, and homogeneous of degree 1 in the set of variables $\{B\} \cup \{S_u, u \in \mathcal{R}\}$ are listed in the four sets below:

$$\mathcal{D}'_1 = \left\{B\,C\,D\,Z\,\Pi^{NR}(A)\right\},$$

$$\mathcal{D}'_2 = \left\{(B + S_u)\,C\,D\,Z\,\Pi^N(A) \,/\, u \in \mathcal{R}\right\},$$

$$\mathcal{D}'_3 = \left\{A^i\,C\,D\,Z\,S_u\,\Pi_u(A) \,/\, u \in \mathcal{R}, i \in \{0, \ldots, l^R - 1\}\right\},$$

$$\mathcal{D}'_4 = \left\{A^i\,C\,D\,Z\,S_u\,\Pi^{NR}(A) \,/\, u \in \mathcal{R}, i \in \{0, \ldots, l(u) - 1\}\right\}.$$

The polynomial in $\mathcal{D}'_1$ is not $B\,C\,D\,Z\,\Pi^N(A)$, since $\Omega^R \neq \emptyset$. As B only appears in polynomials in $\mathcal{D}'_1$ and $\mathcal{D}'_2$, at least one polynomial in $\mathcal{D}'_2$ must be used in the linear combination of elements of $\mathcal{D}'$ which is equal to $B\,C\,D\,Z\,\Pi^N(A)$.

We have to cancel linearly independent terms of the form $S_u\,C\,D\,Z\,\Pi^N(A)$ appearing in the elements of $\mathcal{D}'_2$ used in the linear combination. By considering only linear terms in this specific S_u in the sets $\mathcal{D}'_3$ and $\mathcal{D}'_4$, one can see that it is necessary to build a relation of the form

$$\Pi^N(A) = \Big(\sum_{i=0}^{l^R-1} \lambda_i\,A^i\Big)\,\Pi_u(A) + \Big(\sum_{i=0}^{l(u)-1} \lambda'_i\,A^i\Big)\,\Pi^{NR}(A). \tag{1}$$

By hypothesis, the user u is revoked. We have two cases:

- Either u is in a revoked group, and $\Omega(u) \cap \Omega^R \neq \emptyset$. We consider an attribute μ in this intersection: the polynomial $A - \mu$ divides $\Pi_u(A)$ and $\Pi^{NR}(A)$, and thus it divides the right part of the equation. Since $\Omega^N \cap \Omega^R$ is empty, $A - \mu$ does not divide $\Pi^N(A)$, and the relation (1) can not exist.
- Either u is not in an imposed group, and Ω^N is not included in $\Omega(u)$. So $\Pi^N(A)$ does not divide $\Pi_u(A)$. As $\Pi^N(A)$ divides $\Pi^{NR}(A)$, it divides $(\sum_{j=0}^{l^R-1} \lambda'_j A^j)\,\Pi_u(A)$ as well. It means that we have:

$$\Big(\sum_{i=0}^{l^R-1} \lambda_i A^i\Big)\,\Pi_u(A) = \Pi^N(A)\,Q(A)\,\pi_u(A),$$

where $Q(A)$ is a strict divisor of $\sum_{i=0}^{l^R-1} \lambda_i A^i$ and $\pi_u(A)$ is divisor of $\Pi_u(A)$. So Equation (1) is equivalent to the following equation:

$$Q(A)\,\pi_u(A) + \Big(\sum_{i=0}^{l(u)-1} \lambda'_i A^i\Big)\,\Pi^R(A) = 1, \quad \text{with} \quad \deg(Q) < \deg(\Pi^R) - 1.$$

According to Lemma 2 given in next section of this appendix, such a relation does happen with probability less than $1/\sqrt{p}$.

In one case the relation (1) does not exist, in the other case such a relation exists with a probability less than $1/\sqrt{p}$. So with probability greater than $1 - 1/\sqrt{p}$ there is a contradiction with the hypothesis that $B\,C\,D\,Z\,\Pi^N(A)$ is an element of $\mathcal{M}$.

A.2 Lemma 2

Consider P_1 and P_2 two unitary polynomials of the ring $(\mathbb{Z}/p\mathbb{Z})[X]$ with $\deg P_1 = d_1$ and $\deg P_2 = d_2$. We suppose that P_1 and P_2 are relatively prime. By Bezout's Theorem, there exists V_1, V_2 in $(\mathbb{Z}/p\mathbb{Z})[X]$ unitary such that

$$V_1 P_1 + V_2 P_2 = 1, \quad \text{with} \quad \deg V_1 < d_2 \quad \text{and} \quad \deg V_2 < d_1. \tag{2}$$

The condition over the degrees determines uniquely V_1 and V_2. We are interested here in computing the probability that $\deg V_1 < d_2 - 1$. We have the following lemma:

Lemma 2. *For all $(d_1, d_2) \in (\mathbb{N}^*)^2$, for all prime p such that $p \geq (d_1 + d_2)^2$, the probability taken over all the pairs of relatively prime unitary polynomials (P_1, P_2) in $(\mathbb{Z}/p\mathbb{Z})[X]$ with degree d_1 and d_2 that the pair (V_1, V_2) of unitary polynomials defined uniquely by the relation (2) satisfies* $\deg V_1 < d_2 - 1$ *is upper bounded by $1/\sqrt{p}$.*

Proof. Let $P_1 = X^{d_1} + \sum_{k=0}^{d_1-1} \nu_k X^k$ and $P_2 = X^{d_2} + \sum_{k=0}^{d_2-1} \nu'_k X^k$ be two unitary polynomials of $(\mathbb{Z}/p\mathbb{Z})[X]$, with degrees $d_1 \in \mathbb{N}^*$ and $d_2 \in \mathbb{N}^*$. These two polynomials are relatively non primes if and only if the Sylvester determinant of dimension $d_1 + d_2$ cancels:

$$\det \begin{pmatrix}
\nu_0 & 0 & \cdots & \cdots & 0 & \nu'_0 & 0 & \cdots & 0 \\
\nu_1 & \nu_0 & \ddots & & \vdots & \nu'_1 & \nu'_0 & \ddots & \vdots \\
\vdots & \nu_1 & \ddots & \ddots & \vdots & \vdots & \nu'_1 & \ddots & 0 \\
\nu_{d_1-1} & \vdots & \ddots & \ddots & 0 & \vdots & \vdots & \ddots & \nu'_0 \\
1 & \nu_{d_1-1} & & \ddots & \nu_0 & \nu'_{d_2-1} & \vdots & & \nu'_1 \\
0 & 1 & \ddots & & \nu_1 & 1 & \nu'_{d_2-1} & & \vdots \\
\vdots & \ddots & \ddots & \ddots & \vdots & 0 & 1 & \ddots & \vdots \\
\vdots & & \ddots & \ddots & \nu_{d_1-1} & \vdots & \ddots & \ddots & \nu'_{d_2-1} \\
0 & \cdots & \cdots & 0 & 1 & 0 & \cdots & 0 & 1
\end{pmatrix} = 0.$$

Expanding this determinant, one obtains a polynomial of degree d_1+d_2-1 in the variables $\nu_0,\ldots,\nu_{d_1-1},\nu'_0,\ldots,\nu'_{d_2-1}$ over $(\mathbb{Z}/p\mathbb{Z})$. By Lemma 1 of [Sch80], the probability that this polynomial cancels is bounded by $(d_1+d_2-1)/p$, where the probability is taken over the values of $\nu_0,\ldots,\nu_{d_1-1},\nu'_0,\ldots,\nu'_{d_2-1}$. As a consequence, there is at least $(p+1-d_1-d_2)\,p^{d_1+d_2-1}$ pairs of relatively prime unitary polynomials of degree d_1 and d_2.

From now on, we suppose that P_1 and P_2 are relatively prime unitary polynomials. Let (V_1,V_2) be defined by the relation (2), we suppose that $\deg V_1 < d_2-1$. We have immediately that $\deg V_2 < d_1-1$. The relation (2) with these degree conditions in the $(\mathbb{Z}/p\mathbb{Z})$ vector space $(\mathbb{Z}/p\mathbb{Z})[X]$ implies that the following family is non free:

$$\left(1,\{P_1(X)\,X^k \,/\, k\in\{0,\ldots,d_2-2\}\},\{P_2(X)\,X^k \,/\, k\in\{0,\ldots d_1-2\}\}\right).$$

This property is captured by the cancellation of the following determinant of dimension d_1+d_2-1 depending on the coefficients of P_1 and P_2:

$$\det\begin{pmatrix}
1 & \nu_0 & 0 & \cdots & \cdots & 0 & \nu'_0 & 0 & \cdots & 0\\
0 & \nu_1 & \nu_0 & \ddots & & \vdots & \nu'_1 & \nu'_0 & \ddots & \vdots\\
\vdots & \vdots & \nu_1 & \ddots & \ddots & \vdots & \vdots & \nu'_1 & \ddots & 0\\
\vdots & \nu_{d_1-1} & \vdots & \ddots & \ddots & 0 & \vdots & \vdots & \ddots & \nu'_0\\
\vdots & 1 & \nu_{d_1-1} & & \ddots & \nu_0 & \nu'_{d_2-1} & \vdots & & \nu'_1\\
\vdots & 0 & 1 & \ddots & & \nu_1 & 1 & \nu'_{d_2-1} & & \vdots\\
\vdots & \vdots & \ddots & \ddots & \ddots & \vdots & 0 & 1 & \ddots & \vdots\\
\vdots & \vdots & & \ddots & \ddots & \nu_{d_1-1} & \vdots & \ddots & \ddots & \nu'_{d_2-1}\\
0 & 0 & \cdots & \cdots & 0 & 1 & 0 & \cdots & 0 & 1
\end{pmatrix}=0.$$

Expanding this determinant, one obtains a polynomial of degree d_1+d_2-3 in the variables $\nu_0,\ldots,\nu_{d_1-1},\nu'_0,\ldots,\nu'_{d_2-1}$ over $(\mathbb{Z}/p\mathbb{Z})$. Again by Lemma 1 of [Sch80], the probability that this polynomial cancels is bounded by $(d_1+d_2-3)/p$, where the probability is taken over the values of $\nu_0,\ldots,\nu_{d_1-1},\nu'_0,\ldots,\nu'_{d_2-1}$. As a consequence, there exists at most $(d_1+d_2-3)\,p^{d_1+d_2-1}$ pairs of relatively prime unitary polynomials of degree d_1 and d_2 such that Bezout's equation returns a unitary polynomial V_1 of degree strictly less than d_2-1.

We just have to compute the quotient of the sizes of the two aforementioned sets in order to bound the probability that a pair of relatively prime unitary polynomials verifies Bezout's equation (2) with $\deg(V_1)<d_2-1$:

$$\frac{d_1+d_2-3}{p+1-d_1-d_2}.$$

If $d_1+d_2\le\sqrt{p}$, this probability is bounded by $1/\sqrt{p}$.