

An Efficient Public-Key Attribute-Based Broadcast Encryption Scheme Allowing Arbitrary Access Policies*

Pascal Junod
HEIG-VD
Yverdon-les-Bains, Switzerland
pascal.junod@heig-vd.ch

Alexandre Karlov
Nagravision SA, Cheseaux-sur-Lausanne,
Switzerland
EPFL, Lausanne, Switzerland
alexandre.karlov@nagra.com

ABSTRACT

We describe a new public-key and provably secure attribute-based broadcast encryption scheme which supports complex access policies with **AND**, **OR** and **NOT** gates. Our scheme, especially targetting the implementation of efficient Pay-TV systems, can handle conjunctions of disjunctions by construction and disjunctions of conjunctions by concatenation, which are the most general forms of Boolean expressions. It is based on a modification of the Boneh-Gentry-Waters broadcast encryption scheme in order to achieve attribute collusion resistance and to support complex Boolean access policies. The security of our scheme is proven in the generic model of groups with pairings. Finally, we compare our scheme to several other Attribute-based Broadcast Encryption designs, both in terms of bandwidth requirements and implementation costs.

Categories and Subject Descriptors

E.3 [Data Encryption]: Public-Key Cryptosystems; D.4.6 [Operating Systems]: Security and Protection—*cryptographic controls*

General Terms

Algorithms, Security

Keywords

Attribute-based encryption, broadcast encryption, pairing-based cryptography

1. INTRODUCTION

Securing a broadcast channel has always been an interesting and challenging task for cryptographers and has been discussed for the first time by Berkovits [3] and Fiat and

*© ACM, (2010). This is the author's version of the work. It is posted here by permission of ACM for your personal use. Not for redistribution.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

DRM'10, October 4, 2010, Chicago, Illinois, USA.

Copyright 2010 ACM 978-1-4503-0091-9/10/10 ...\$10.00.

Naor [13]. In this setting, the broadcasting center can send an encrypted message to a set of privileged, i.e., non-revoked users which is a subset of the set of all possible receivers. We can distinguish between two receiver models: in the *stateless* receiver model it is not possible, or too costly in terms of bandwidth, to guarantee synchronism with the broadcasting center. For the *stateful* receiver model [8, 9, 22, 25, 27], one assumes that synchronism is guaranteed between the receivers and the broadcasting center, with help of a feedback channel, for instance. In this paper, we will assume to find ourselves in a pure stateless scenario, more precisely in the Pay-TV setting, where bandwidth issues are of uttermost importance.

Attribute-Based Encryption.

It is noteworthy that in certain scenarios, like in Pay-TV systems, for instance, the receivers can frequently be arranged according to some natural characteristics, or *attributes*: one can mention the receiver's geographical location based on a ZIP code, their subscription to certain packages or their current firmware version. Intuitively, the broadcaster should be able to broadcast in a bandwidth-efficient way to receivers satisfying a set of these properties in a more or less complex manner, often modeled by a *Boolean access policy*. For instance, the broadcaster might desire to enforce an access policy by sending the content only to receivers which are in (“New York”) **OR** (“New Jersey”) **AND** (“with a receiver's firmware not older than 2.1.1”). Another appealing and direct application of attribute-based encryption in a broadcast setting is the direct mapping of families of Pay-TV channels to a single attribute (we might call this attribute a *product*): for instance, we can imagine mapping all the TV channels targetting kids to an attribute named “Family TV”. In some circumstances, for example in football games, it is required that only receivers in the specific geographical region are able to decrypt the content (for instance, anywhere but around the stadium, in order to encourage local people to physically go to the game). This operation is called a “blackout” and is easy to realize if there exists a geographical attribute per receiver. Another example is the one of promotional packages: subscribers who have their birthday in the current month can watch a given channel package for free. Finally, with the recent deployment of High-Definition (HD) video content, we might also imagine that the HD content can be decrypted only by the newer (and more secure) receivers holding a corresponding attribute. In summary, the broadcaster might not be interested in (or does not know) all the receivers which are able to

access the content, but merely wants to describe the authorized set of receivers in terms of some descriptive attributes using a Boolean access policy and to efficiently broadcast the allowed receivers a symmetric session key encrypting the multimedia content.

Direct Revocation.

Another explicit requirement in the broadcast setting is that it should be possible to directly revoke individual receivers without impacting non-revoked users and this in a bandwidth-efficient way. For instance, the fact that an individual receiver does not pay anymore its subscription fees should not impact other receivers. In particular, such an event should not imply any re-keying operation, for those operations are either impossible or very costly in terms of bandwidth in a pure broadcasting scenario.

Flexibility of Attributes Organisation.

In practice, broadcasters tend to frequently change the structure of their products, depending on their current business model. For instance, they might add a new channel to an existing product, or move one or several channels from one product to another. Hence, it should be easy, bandwidth-efficient and seamless for the receivers to change the structure of products.

1.1 Related Work

Attribute-Based Encryption.

The notion of attribute-based encryption (ABE) was introduced by Sahai and Waters in [23] as a generalization of ID-based encryption called *Fuzzy IBE*. Their scheme is a threshold ABE system where ciphertexts are labeled by a certain set of attributes and users' private keys are associated with a set of attributes along with a threshold parameter. At least k attributes must overlap between the two sets in order to be able to decrypt a ciphertext. Goyal et al. [16] formalized the concepts of key-policy ABE (KP-ABE) and ciphertext-policy ABE (CP-ABE) and provided a construction for the former with a security proof in the generic bilinear group model. In the KP-ABE model, the access structure is specified in the private key, while in the CP-ABE one, it is specified in the ciphertext, those two forms being complementary to each other. Bethencourt, Sahai and Waters proposed the first construction of a CP-ABE in [4]. Their scheme can handle **AND** and **OR** gates using so-called *access trees*. Later on, Ostrovsky, Sahai and Waters [21] extended both schemes to handle any non-monotone access structures, including the possibility of using negated clauses in access policies. Recently Goyal et al. [15] proposed a CP-ABE scheme supporting any access policies of bounded polynomial size, notably with a security proof in the standard model. Another prominent CP-ABE construction is the one proposed by Waters [26] and based on the concept of *linear secret sharing scheme (LSSS)* [2]. It is quite similar to the Bethencourt, Sahai and Waters construction, except that the security is proven in the standard model and that it is fully expressive. Chase, in [10], proposed a multi-authority attribute-based encryption scheme where attribute keys are issued by multiple authorities and which can achieve conjunction in a single authority setting over a pre-determined number of clauses. Müller, Katzenbeisser and Eckert [19] give a construction supporting DNF

policies and which shares the idea of blinding the private key.

Broadcast Encryption.

The notion of broadcast encryption was introduced by Berkovits [3], quickly followed by the important work of Fiat and Naor [13]. Since then, several stateless broadcast encryption schemes have been proposed in the literature [6, 7, 11, 12, 14, 17, 20]. In such schemes the broadcasting center can dynamically specify a privileged subset of authorized receivers among ℓ receivers that can decrypt selected ciphertexts.

Attribute-Based Broadcast Encryption.

CP-ABE scheme supporting negated clauses allows a direct revocation of individual receivers by conjunctively adding the AND of negations of revoked user identities (where each identity is mapped to an individual attribute), however this solution lacks efficiency in bandwidth terms. For instance, if we use Ostrovsky et al. [21] CP-ABE scheme, the revocation of users would add an overhead of $O(r)$ group elements to the ciphertext, where r is the cardinality of the revoked receivers set. While in traditional attribute-based encryption schemes the revocation can be performed solely based on attributes, an *attribute-based broadcast encryption (ABBE)* scheme should allow individual receivers to be directly revoked as well in an efficient way. In [18], Lubicz and Sirvent propose an ABBE scheme allowing to express access policies in disjunctive normal form (i.e. disjunction - **OR** of conjunctions - **AND**), with the **OR** function provided by ciphertext concatenation, and being able to handle attribute negations (**NOT**) as well. In their scheme, the authors however use an individual receiver-specific attribute and the disjunction is obtained by concatenation of several instances of the encryption scheme. Attrapadung and Imai [1] propose another approach, namely using a separate broadcast encryption scheme on the top of an ABE construction, to construct ciphertext-policy and key-policy variants. In both papers, the receiver revocation is conjunctive, meaning that even if the receiver possesses all the necessary attributes for a given clause, but belongs to the non-authorized set, it will not be able to decrypt the ciphertext correctly.

1.2 Application to Pay-TV

It was argued in [1, 16] that Pay-TV is a natural application of KP-ABE schemes, i.e., broadcasting multimedia content holding a set of properties to receivers storing a private key generated according a pre-defined access policy. From the attribute-based broadcast encryption perspective in a stateless scenario, we are certainly more interested in CP-ABE. As a matter of fact, the roles are inverted in currently deployed Pay-TV systems: the content comes with an attached access policy and the receivers, depending on the attributes they have at their disposal, are able or not to decrypt the content (i.e., clearly following a CP-ABE philosophy). Indeed, let us assume that you attach several TV channels to a single attribute (we might call this attribute a "product"). Then, the access policy defines which products give an access to a given channel (or content). In practice, broadcasters tend to frequently change the structure of their products, depending on their current business model. For instance, they might add a new channel to an existing product. Hence, in a KP-ABE scenario, changing the structure

of products would imply sending individual messages to each receiver containing a new, individualized access policy. In a stateless broadcast scenario, where guaranteeing synchronism between the broadcasting center and the receiver is extremely costly in terms of bandwidth, this is a practically impossible task to perform if the number of users is large. Accordingly, we are firmly convinced that CP-ABE is much more flexible and better suited for management of Pay-TV contents. As a final remark, we would like to emphasize that bandwidth needs are likely the most important feature looked at when comparing encryption schemes in the Pay-TV world. Indeed, the computational capacities of modern receivers tend to follow Moore's law in a quite natural way, while increasing bandwidth capacities in a pure broadcast setting is extremely costly.

1.3 Our Contributions

In this paper, we describe a new public-key and provably secure attribute-based broadcast encryption scheme which supports complex access policies with **AND**, **OR** and **NOT** gates. One of our goals, besides obtaining a high flexibility for the definition of access policies, was to optimize the bandwidth requirements (i.e., the ciphertext size) as much as possible, somewhat sacrificing the size of private keys and the encryption/decryption costs. Our scheme can handle conjunctions of disjunctions (CNF) by construction and disjunctions of conjunctions (DNF) by concatenation; furthermore, it supports direct revocation of individual receivers as well. Our construction is based on a modification of the Boneh-Gentry-Waters broadcast encryption scheme [6] to achieve attribute collusion resistance and to support complex Boolean access policies, the attribute collusion attack being likely the principal reason why broadcast encryption primitives cannot be directly used to build ABE and ABBE schemes. The security of our scheme is proven in the generic model of groups with pairings. Finally, we compare our scheme to several other ABBE designs, both in terms of bandwidth requirements and implementation costs.

2. ATTRIBUTE-BASED BROADCAST ENCRYPTION

2.1 Mathematical Preliminaries

To begin, we briefly review necessary facts about bilinear maps and bilinear map groups. Let \mathbb{G} and \mathbb{G}_T be two cyclic groups of prime order p , whose operation will be multiplicatively written. Let g be a generator of \mathbb{G} and let $e : \mathbb{G} \times \mathbb{G} \rightarrow \mathbb{G}_T$ be a non-degenerate bilinear map, namely such that for all $x, y \in \mathbb{G}$ and $a, b \in \mathbb{Z}/p\mathbb{Z}$, we have $e(x^a, y^b) = e(x, y)^{ab}$ and $e(g, g) \neq 1$. \mathbb{G} will be called a *bilinear group* if the group action in \mathbb{G} can efficiently be computed and if there exists a group \mathbb{G}_T and an efficiently computable bilinear map $e(\cdot, \cdot)$ defined as above.

The security of our system will be proved in the generic model of groups with pairings. In [5], Boneh, Boyen and Goh introduced the *Generalized Diffie Hellman Exponent (GDHE)* assumption which covers a large number of assumptions in the generic bilinear group model. Let $f \in \mathbb{F}_p[X_1, \dots, X_n]$ be a polynomial over \mathbb{F}_p and

$$P, Q \in \mathbb{F}_p[X_1, \dots, X_n]^s$$

be two s -tuples of polynomials. We write $P = (p_1, \dots, p_s)$ and $Q = (q_1, \dots, q_s)$ and we require that $p_1 = q_1 = 1$. For a

function $\varphi : \mathbb{F}_p[X_1, \dots, X_n] \rightarrow \Omega$, we write

$$\varphi(P(X_1, \dots, X_n)) = (\varphi(p_1(X_1, \dots, X_n)), \dots, \varphi(p_s(X_1, \dots, X_n))).$$

In what follows, we briefly recall the decisional version of the Generalized Diffie-Hellman Exponent Problem as introduced in [5], the concept of dependent functions and the definition of the degree of a set of multivariate polynomials over $\mathbb{F}_p[X_1, \dots, X_n]^s$.

DEFINITION 1 (GDHE DECISIONAL PROBLEM). *Given a generator $g \in \mathbb{G}$, $h = e(g, g)$ and the vector*

$$(g^{P(X_1, \dots, X_n)}, h^{Q(X_1, \dots, X_n)}) \in \mathbb{G}^s \times \mathbb{G}_T^s$$

distinguish $h^{f(X_1, \dots, X_n)}$ from a random value $U \in_R \mathbb{G}_T$.

DEFINITION 2 (DEPENDENT FUNCTIONS). *A function f is said to be dependent on the sets P and Q if there exist $s^2 + s$ constants $\{a_{i,j}\}_{i,j=1}^s, \{b_k\}_{k=1}^s$ such that*

$$f = \sum_{i,j=1}^s a_{i,j} p_i p_j + \sum_{k=1}^s b_k q_k$$

A function which is not dependent on (P, Q) is said to be *independent* of (P, Q) .

DEFINITION 3. *For a set $P \subseteq \mathbb{F}_p[X_1, \dots, X_n]^s$, the degree of P is $\deg(P) = \max_{f \in P} \deg(f)$, where $\deg(f)$ is the total degree of polynomial $f \in \mathbb{F}_p[X_1, \dots, X_n]^s$.*

The following result of Boneh, Boyen and Goh [5], expressed in the framework of generic groups [24], gives a complexity upper bound on the security of the decisional version of the Generalized Diffie-Hellman Exponent Problem in the generic bilinear group model. One considers two random encodings ξ and ξ_T of the additive group \mathbb{Z}_p^+ , i.e., injective maps $\xi, \xi_T : \mathbb{Z}_p^+ \rightarrow \{0, 1\}^m$. Let furthermore $\mathbb{G} = \{\xi(x) : x \in \mathbb{Z}_p^+\}$ and $\mathbb{G}_T = \{\xi_T(x) : x \in \mathbb{Z}_p^+\}$. The adversary is given oracles to compute the induced group action on \mathbb{G} and \mathbb{G}_T as well as an oracle to compute a non-degenerate bilinear map $e : \mathbb{G} \times \mathbb{G} \rightarrow \mathbb{G}_T$. Those oracles hence hide the groups structure to the adversary.

THEOREM 1 (BONEH, BOYEN AND GOH [5]). *Let*

$$d = \max(2 \deg(P), \deg(Q), \deg(f)).$$

If f is independent of (P, Q) , then for any adversary \mathcal{A} that makes a total of at most q queries to the oracles computing the group operations in \mathbb{G} , \mathbb{G}_T and the pairing e , we have:

$$\Pr \left[\mathcal{A} \left(\begin{array}{c} p, \xi(P(X_1, \dots, X_n)), \\ \xi_T(Q(X_1, \dots, X_n)), \\ \xi_T(t_0), \xi_1(t_1) \end{array} \right) = b : \begin{array}{c} X_1, \dots, X_n, y \stackrel{R}{\leftarrow} \mathbb{Z}/p\mathbb{Z}, \\ b \stackrel{R}{\leftarrow} \{0, 1\}, \\ t_b \leftarrow f(X_1, \dots, X_n), \\ t_{1-b} \leftarrow y \end{array} \right) - \frac{1}{2} \right] \leq \frac{(q + 2s + 2)^2 \cdot d}{2^p}$$

2.2 Boolean Access Policies

We now discuss the concept of Boolean access policies and the associated notations we will use. Let us denote by $\mathcal{U} = \{u_1, u_2, \dots, u_\ell\}$ the set of cardinality ℓ of all users within the system and that might be allowed to receive some confidential information. A group of users is then simply defined as a non-empty set $\mathcal{G} \subseteq \mathcal{U}$, while $\mathcal{B}(u)$, for a user

$u \in \mathcal{U}$, is the set of all groups the user belongs to. For instance, if $\mathcal{U} = \{u_1, u_2, u_2\}$ and $\mathcal{G}_1 = \{u_1, u_2\}$, $\mathcal{G}_2 = \{u_2\}$ and $\mathcal{G}_3 = \{u_1, u_3\}$, then $\mathfrak{B}(u_1) = \{\mathcal{G}_1, \mathcal{G}_3\}$, $\mathfrak{B}(u_2) = \{\mathcal{G}_1, \mathcal{G}_2\}$ and $\mathfrak{B}(u_3) = \{\mathcal{G}_3\}$. For ease of notation, we will assign an *attribute* A_i to a user belonging to group \mathcal{G}_i , and, accordingly, assign a *negated attribute* \bar{A}_i to a user *not* belonging to that group. We define the attribute repartition for user u_i as $\mathfrak{B}(u_i)$ for $i = 1, \dots, \ell$. Practically, groups of users can be organized according to some property or characteristic they have in common. This can be their geographic location, their adherence to some subscription package, the version of firmware they are running or a property of any other nature. For ease of understanding, we will denote by \mathcal{B} and $\bar{\mathcal{B}}$ the sets of positive attributes $\mathcal{B} = \{A_1, \dots, A_r\}$ and of negative attributes $\bar{\mathcal{B}} = \{A_{r+1}, \dots, A_{r+s}\}$, respectively.

The concept of *Boolean access policy* is central in ABBE schemes: it defines which groups are allowed to decrypt or not a given ciphertext. For instance, the expression $\mathbb{A} = \bar{A}_1 \wedge (A_2 \vee A_3)$ is a Boolean access policy which would allow all users being either in \mathcal{G}_2 or \mathcal{G}_3 , but not in \mathcal{G}_1 , to decrypt the ciphertext. Boolean access policies can virtually be any kind of Boolean expressions, however, we will be interested in specific forms of expressions, like the *disjunctive normal form (DNF)* or the *conjunctive normal form (CNF)*: an expression in DNF will be written as $\bigvee_{i=1}^n \bigwedge_{j=1}^m \alpha_{i,j}$, while an expression in CNF is written as $\bigwedge_{i=1}^n \bigvee_{j=1}^m \alpha_{i,j}$, where the literals $\alpha_{i,j}$ can be negated or not. Those two forms are universal, since every Boolean expression can be written in CNF and DNF; however, a conversion from one of those two forms to another might result in an exponential blow-up of the number of clauses. In the following, we write $\mathfrak{B}(u_i) \sim \mathbb{A}$ (respectively $\mathfrak{B}(u_i) \approx \mathbb{A}$) to mean that the attribute set $\mathfrak{B}(u_i)$ is (not) compatible with the access policy \mathbb{A} .

The formal definition of an attribute-based broadcast encryption scheme consists of three randomised algorithm:

1. **Setup**($1^\lambda, \ell, \mathfrak{B}(u_i)_{1 \leq i \leq \ell}$): This algorithm takes a security parameter λ , the total number ℓ of users within the system, and the attribute repartition $\mathfrak{B}(u_i)$ for each user u_i . It returns an encryption key \mathbf{ek} and ℓ decryption keys \mathbf{dk}_i which will be distributed to each respective receiver.
2. **Encrypt**(\mathbf{ek}, \mathbb{A}): This algorithm takes the encryption key \mathbf{ek} and an access policy \mathbb{A} in input, and it returns a header \mathbf{hdr} as well as a session key $\mathbf{SK} \in \mathcal{K}$, where \mathcal{K} is a finite set of message encryption keys.
3. **Decrypt**($\mathbb{A}, \mathbf{hdr}, \mathbf{dk}_i$): This algorithm takes a decryption key \mathbf{dk}_i , a header \mathbf{hdr} and an access policy \mathbb{A} ; it returns the session key \mathbf{SK} if and only if $\mathfrak{B}(u_i) \sim \mathbb{A}$ and otherwise, it outputs the symbol \perp .

Such a system has obviously to be correct, namely that for all possible access policies \mathbb{A} and all possible attribute repartitions $\mathfrak{B}(u_i)_{1 \leq i \leq \ell}$, if

$$(\mathbf{ek}, \mathbf{dk}_1, \dots, \mathbf{dk}_\ell) = \mathbf{Setup}(1^\lambda, \ell, \mathfrak{B}(u_i)_{1 \leq i \leq \ell})$$

and $(\mathbf{hdr}, \mathbf{SK}) = \mathbf{Encrypt}(\mathbf{ek}, \mathbb{A})$, then

$$\mathbf{Decrypt}(\mathbb{A}, \mathbf{hdr}, \mathbf{dk}_i) = \mathbf{SK}$$

for the u_i 's such that $\mathfrak{B}(u_i) \sim \mathbb{A}$ and

$$\mathbf{Decrypt}(\mathbb{A}, \mathbf{hdr}, \mathbf{dk}_i) = \perp$$

for the u_i 's with $\mathfrak{B}(u_i) \not\sim \mathbb{A}$.

2.3 Security Model

2.3.1 Semantic Security

In this paper, we will consider a slightly more general version of the model considered by Lubicz and Sirvent in [18] which they called *semantic security with full static collusions*. Contrarily to [18], we allow the adversary to fix the attributes repartition $\mathfrak{B}(u_i)$ for all users i . An ABBE scheme will be considered secure within this model if given a header and all the decryption keys of revoked users, it is not possible for an adversary to infer any information about the session key. More formally, let us consider the following game:

1. The challenger and the adversary \mathcal{A} are given a system consisting of n attributes.
2. The adversary \mathcal{A} outputs a Boolean policy \mathbb{A} as well as a repartition $\mathfrak{B}(u_i)_{1 \leq i \leq \ell}$ which he intends to attack.
3. The challenger runs the algorithm **Setup**($1^\lambda, \ell, \mathfrak{B}(u_i)_{1 \leq i \leq \ell}$) and gives to \mathcal{A} the public key \mathbf{ek} and the decryption keys \mathbf{dk}_i corresponding to the users u_i that the adversary may control, i.e.,

$$\{u_i : \mathfrak{B}(u_i) \approx \mathbb{A}\}.$$

4. The challenger runs the algorithm **Encrypt**(\mathbf{ek}, \mathbb{A}) and obtains a header \mathbf{hdr} and a session key \mathbf{SK} . Next, the challenger draws a bit b uniformly at random, set $\mathbf{SK}_b = \mathbf{SK}$, $\mathbf{SK}_{1-b} \in_R \mathcal{K}$ and finally gives

$$(\mathbf{hdr}, \mathbf{SK}_b, \mathbf{SK}_{1-b}) \text{ to } \mathcal{A}.$$

5. The adversary \mathcal{A} outputs a guess bit b' .

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

$$\text{Adv}^{\text{ind}}(\lambda, n, \mathfrak{B}(u_i)_{1 \leq i \leq \ell}, \mathcal{A}) = |2 \Pr[b = b'] - 1|,$$

where the probability is taken over the random bit b and all the bits used in the simulation of the algorithms **Setup**(.) and **Encrypt**(.). Then, semantic security against full static collusions is defined as follows.

DEFINITION 4. *An ABBE scheme is semantically secure against full static collusions if for all randomised polynomial-time adversaries \mathcal{A} and for all access policies involving at most n attributes defined by $\mathfrak{B}(u_i)_{1 \leq i \leq \ell}$,*

$$\text{Adv}^{\text{ind}}(\lambda, n, \mathfrak{B}(u_i)_{1 \leq i \leq \ell}, \mathcal{A})$$

is a negligible^a function of λ when n and ℓ are at most polynomial in λ .

2.3.2 Attributes Collusion Attack

An important security property of attribute-based encryption schemes is resistance against attribute collusions, that is: if a user u_1 has attribute A_1 and a user u_2 has attribute A_2 then they should not be able to decrypt a header which has access policy $A_1 \wedge A_2$. We note that a simple combination of broadcast encryption systems with every key being an attribute is trivially prone to this kind of attack.

^aA function $f : \mathbb{N} \rightarrow \mathbb{R}^+$ is called *negligible* if for any polynomial p there exists an integer x_0 such that $x \geq x_0 \implies f(x) < \frac{1}{p(x)}$.

3. CONSTRUCTION

As before, denote by \mathcal{U} the set of all users, with $|\mathcal{U}| = \ell$. In a natural way, any broadcast encryption system is disjunctive (i.e. is an *OR*-protocol): only non-revoked users $u \in \mathcal{S} \subseteq \mathcal{U}$ are able to decrypt a broadcasted message. For instance, the broadcasting center can enforce the fact that only users i_1, i_2 and i_3 receive the content, that is $u_{i_1} \vee u_{i_2} \vee u_{i_3}$ would be able to decrypt the session key. Let $\mathcal{B} \cup \bar{\mathcal{B}} = \{A_1, A_2, \dots, A_n\}$ be the set of all attributes. Each user has one or several attributes, that is $\mathfrak{B}(u_i) = \{j \in \{1, \dots, n\} \mid u_i \text{ has attribute } A_j\}$ and hence one or several users are associated with a given attribute A_j .

Consider now a generic broadcast encryption system. By associating the decryption keys with attributes and distributing those keys to the users according to the user-attribute relation, we obtain a very simple ABBE scheme that is able to broadcast to a disjunction of attributes, i.e. every user associated with an attribute A_i will have the decryption key for this attribute. The main issue with this approach is that it does not guarantee attribute collusion resistance. In order to address this problem, we chose to modify the underlying scheme by using private key blindings and a final key derivation in order to support complex access policies along with attribute collusion resistance.

3.1 Achieving Attribute Collusion Resistance

We show now how to modify the Boneh-Gentry-Waters public-key broadcast encryption scheme [6] to obtain the attribute collusion-resistance property. In our scheme, every private key is unique to a given user u_i , with $1 \leq i \leq \ell$. Below, n is the total number of attributes in the system, that is $|\mathcal{B} \cup \bar{\mathcal{B}}| = n$. We now formally define the three algorithms, namely **Setup**(\cdot), **Encrypt**(\cdot) and **Decrypt**(\cdot).

Setup($1^\lambda, \ell, \mathfrak{B}(u_i)_{1 \leq i \leq \ell}$).

We choose two cyclic groups \mathbb{G} and \mathbb{G}_T of prime order p according to the security parameter λ . Let g be a generator of \mathbb{G} and let $e : \mathbb{G} \times \mathbb{G} \rightarrow \mathbb{G}_T$ be a non-degenerate bilinear map. Like in the Boneh-Gentry-Waters scheme, this algorithm picks a random generator $g \in_R \mathbb{G}$, two random values $\alpha, \gamma \in_R \mathbb{Z}/p\mathbb{Z}$ and for $i = 1, \dots, n, n+2, \dots, 2n$, it computes $g_i = g^{\alpha^i} \in \mathbb{G}$ and $v = g^\gamma$. It generates also two new secret values $\beta, r \in_R \mathbb{Z}/p\mathbb{Z}$. The encryption key ek is public and it is given by $\text{ek} = (g_1^r, \dots, g_n^r, g_{n+2}^r, \dots, g_{2n}^r, v^r, g_n^\beta, g_n)$. To compute the decryption key of a user u which has the N_1 positive attributes $A_{i_1}, \dots, A_{i_{N_1}}$ and the N_2 negative attributes $A_{j_1}, \dots, A_{j_{N_2}}$, the setup algorithm generates a random value $s_u \in_R \mathbb{Z}/p\mathbb{Z}$ and computes

$$\text{dk}_u = (g_1^{r(\beta+s_u)}, g_1^{s_u}, \dots, g_n^{s_u}, g_{n+2}^{s_u}, \dots, g_{2n}^{s_u}, d_{i_1}, d_{i_2}, \dots, d_{i_{N_1}}, d_{j_1}, \dots, d_{j_{N_2}})$$

where the d_i values are defined as $d_i = g_i^{\gamma \cdot s_u}$. Note that each attribute has its positive and negative version associated with two different keys.

Encrypt(ek, \mathbb{A}).

This algorithm takes the encryption key ek and an access policy \mathbb{A} in input, and it returns a header hdr as well as a session key SK . We distinguish between two cases:

- The access policy is expressed in CNF $\mathbb{A} = \beta_1 \wedge \beta_2 \wedge \dots \wedge \beta_i \wedge \dots \wedge \beta_N$. Let $t_1, \dots, t_N \in_R \mathbb{Z}/p\mathbb{Z}$ and $t =$

$\sum_{i=1}^N t_i \bmod p$. The header of the message will consist of $N+1$ parts and will be computed as $\text{hdr} = (g_n^t, \text{hdr}_1, \dots, \text{hdr}_N)$. Each clause is implicitly related to a session key $\text{SK}_i = e(g_{n+1}, g)^{rt_i}$. The formula to compute the N parts of the header is similar to the BGW scheme, i.e.,

$$\text{hdr}_i = \left(g^{rt_i}, \left(v^r \prod_{j \in \beta_i} g_{n+1-j}^r \right)^{t_i} \right) \in \mathbb{G}^2, \quad (1)$$

while the global session key of the header is given by $\text{SK} = e(g_1^r, g_n^\beta)^t = e(g, g)^{\beta r \alpha^{n+1} t}$.

- Provided an access policy expressed in DNF $\mathbb{A} = \beta_1 \vee \beta_2 \vee \dots \vee \beta_i \vee \dots \vee \beta_N$, the header of message will consist of N parts $\text{hdr} = (\text{hdr}^{(1)}, \dots, \text{hdr}^{(N)})$ where the part $\text{hdr}^{(i)}$ corresponds to the clause β_i :

$$\text{hdr}^{(i)} = \left(g_n^{t^{(i)}}, \text{hdr}_{i,1}, \dots, \text{hdr}_{i,M} \right)$$

Each clause is then related to a global session key $\text{SK}^{(i)} = \left(\prod_{j=1}^N \text{SK}_{i,j} \right)^\beta$ with $i = 1, \dots, N$ and $t^{(i)} = \sum_{j=1}^M t_j^{(i)}$. Since it is a DNF access policy, it is enough to have any of the clause β_i to be fulfilled, that is any of the global session keys $\text{SK}^{(i)}$ can decrypt the message. The part $\text{hdr}_{i,j}$ is derived exactly as in (1) except that only one attribute (only one decryption key) A_ϕ will be targeted, i.e.:

$$\text{hdr}_{i,j} = \left(g^{rt_j^{(i)}}, (v^r \cdot g_{n+1-\phi})^{t_j^{(i)}} \right) \in \mathbb{G}^2$$

Decrypt($\mathbb{A}, \text{hdr}, \text{dk}_i$).

This algorithm takes a decryption key dk_i , a header hdr and an access policy \mathbb{A} and returns the session key SK if the decryption key dk_i is allowed to decrypt the ciphertext. As for the encryption operation, we distinguish two cases:

- Provided the header

$$\text{hdr} = (\text{hdr}_0, \text{hdr}_1, \dots, \text{hdr}_N)$$

in CNF with $\text{hdr}_i = (C_0, C_1)$, $1 \leq i \leq N$, for each clause β_i which contains the attribute A_k , a receiver which has this attribute can compute

$$\text{SK}_i^{s_u} = \frac{e(g_k^{s_u}, C_1)}{e \left(d_k \cdot \prod_{j \in \beta_i, j \neq k} g_{n+1-j+k}^{s_u}, C_0 \right)}$$

The global session key SK is then given by

$$\text{SK} = \frac{e(\text{hdr}_0, g_1^{r(\beta+s_u)})}{\prod_{i=1}^N \text{SK}_i^{s_u}}$$

- Provided the header

$$\text{hdr} = ((\text{hdr}_{0,1}, \text{hdr}_{1,1}, \dots, \text{hdr}_{1,M}), \dots, (\text{hdr}_{N,1}, \text{hdr}_{N,1}, \dots, \text{hdr}_{N,M}))$$

expressed in DNF with $\text{hdr}_{i,j} = (C_0, C_1)$, $1 \leq j \leq M$, a receiver can compute

$$\left(\text{SK}_j^{(i)} \right)^{s_u} = \frac{e(g_k^{s_u}, C_1)}{e \left(d_k \cdot g_{n+1-\phi+k}^{s_u}, C_0 \right)}$$

for an attribute A_ϕ that it has. The global session key (among N valid session keys) is then given by

$$\text{SK}^{(i)} = \frac{e(\text{hdr}_{i,0}, g_1^{r(\beta+s_u)})}{\prod_{j=1}^M (\text{SK}_j^{(i)})^{s_u}}$$

We show in §B.1 that the encryption is sound for policies expressed in CNF, while the DNF case is similar.

Direct revocation.

The direct revocation of the receiver i is efficiently achieved by using its unique identifier, which in this case will be represented by the attribute A_{id_i} proper only to this receiver. In the CNF case, the final policy will be $\mathbb{A} = \mathbb{A}_{\text{CNF}} \wedge (A_{\text{id}_{i_1}} \vee A_{\text{id}_{i_2}} \vee \dots \vee A_{\text{id}_{i_m}})$. In the DNF case, the final policy will be $\mathbb{A} = (\mathbb{A}_{\text{DNF}}) \wedge (A_{\text{id}_{i_1}} \vee A_{\text{id}_{i_2}} \vee \dots \vee A_{\text{id}_{i_m}})$. It is important to note that in both cases, the only way to achieve the direct revocation is conjunctively, i.e. only receivers identified by $A_{\text{id}_{i_1}}, A_{\text{id}_{i_2}}, \dots, A_{\text{id}_{i_m}}$ AND satisfying \mathbb{A}_{DNF} (respectively \mathbb{A}_{CNF}) will be able to decrypt the content. Additionally, in the DNF case, each valid session key among N must be mixed with the direct revocation session key using a one-way function, for example. Note also that, in both cases, the direct revocation requires only one additional ciphertext element.

4. SECURITY

To make our security analysis more intelligible, we will process in two steps. First we prove the semantic security for the case where one user has all the revoked attributes and we show that the advantage of distinguishing the valid ciphertext from a random is negligible. In the second step we show that two attribute sets belonging to different users (i.e. blinded under different constants) cannot be combined to distinguish a ciphertext formed under these attributes from a random value. In fact, our approach can be justified by the following argument: suppose the adversary chooses $\mathbb{A} = \beta_1 \wedge \beta_2 \wedge \dots \wedge \beta_n$ as the policy he plans to attack. We have to provide the adversary with all attributes such that the above policy is not satisfied. That is, it can be provided either with attributes in $\overline{\beta_1}$, or with attributes in $\overline{\beta_2}$ and so on. However it will never get attributes in $\overline{\beta_1} \vee \overline{\beta_2} \vee \dots \vee \overline{\beta_n}$ under the same blinding. Therefore the first step will consist in showing that even if the adversary has all the attributes that do not satisfy \mathbb{A} , it will be unable to distinguish the valid ciphertext from a random value. In the second step we show that even if the adversary gets the attributes in β_1 blinded for user u_1 and attributes in β_2 blinded for user u_2 , he is still unable to distinguish the ciphertext from a random value. Finally we combine these two results to prove the semantic security of our scheme.

4.1 Single-Receiver Semantic Security

Following the security model described in §2.3, the adversary \mathcal{A} outputs the Boolean policy $\mathbb{A} = \beta_1 \wedge \beta_2 \wedge \dots \wedge \beta_n$ which he wants to attack. Each clause β_i is a set of attributes $\{A_{i_1}, A_{i_2}, \dots, A_{i_N}\}$ represented by private decryption keys $\{d_{i_1}, d_{i_2}, \dots, d_{i_N}\}$. Among these keys, one is sufficient to correctly decrypt the clause β_i . Then, the challenger runs the **Setup**(\cdot) algorithm and provides the adversary with all decryption keys corresponding to the set of attributes $\mathcal{G}_R = \mathcal{G} \setminus (\beta_1 \cap \beta_2 \dots \cap \beta_n)$ with $|\mathcal{G}_R| = R$. That

is, the adversary is provided with the private key $\text{dk}_u = (g_1^{r(\beta+s_u)}, g_1^{s_u}, \dots, g_n^{s_u}, g_{n+2}^{s_u}, \dots, g_{2n}^{s_u}, d_{i_1}, d_{i_2}, \dots, d_{i_R})$ where $d_{i_j} \in \mathcal{G} \setminus (\beta_1 \cap \beta_2 \dots \cap \beta_n)$ and $s_u \in_R \mathbb{Z}/p\mathbb{Z}$. According to the framework of Boneh et al. [5], we now describe this fact as an instance of the (P, Q, f) -GDHE problem with

$$P = \left(\begin{array}{c} 1, r, \alpha r(\beta + s_u), \alpha^n \beta, \gamma r, \gamma \alpha^k s_u \\ \alpha r, \alpha^2 r, \dots, \alpha^n r, \alpha^{n+2} r, \dots, \alpha^{2n} r \\ \alpha s_u, \alpha^2 s_u, \dots, \alpha^n s_u, \alpha^{n+2} s_u, \dots, \alpha^{2n} s_u \\ \alpha^n t, r t_1, \dots, r t_R, \alpha^n \\ t_{i_1}(\gamma r + \sum_{j \in \beta_{i_1}} \alpha^{n+1-j}), \dots, \\ t_{i_R}(\gamma r + \sum_{j \in \beta_{i_R}} \alpha^{n+1-j}) \end{array} \right)$$

$$Q = (1)$$

$$f = \alpha^{n+1} r \beta \sum_{i_j \in \mathcal{G}_R} t_{i_j}.$$

where $t = \sum_{i=1}^N t_i \pmod p$. We first need to show the independence of f and (P, Q) (according to Def. 2).

Lemma 1. If $d_{i_j} \in \mathcal{G} \setminus (\beta_1 \cap \beta_2 \dots \cap \beta_n)$, then (P, Q) are independent of f .

PROOF. The proof is given in §B.2 \square

We can now state the following result, which follows from Theorem 1 in a straightforward way.

THEOREM 2. For any probabilistic algorithm \mathcal{A} that totalizes at most q queries to the oracles performing group operations in $(\mathbb{G}, \mathbb{G}_T)$ and evaluations of $e(\cdot, \cdot)$

$$\text{Adv}^{\text{GDHE}}(\mathcal{A}) \leq \frac{(q + 2(4n + 6 + 2R) + 2)^2}{p}.$$

4.2 Attribute collusion resistance

We are now going to prove the attribute collusion resistance property. First we start by a simple case with $\mathbb{A} = \beta_1 \wedge \beta_2$ and thus having only two clauses. We will also consider two users u_1 and u_2 for the moment. As with semantic security, the collusion resistance can be described by a (P, Q, f) -GDHE problem with

$$P = \left(\begin{array}{c} 1, r, \alpha r(\beta + s_{u_1}), \alpha r(\beta + s_{u_2}), \\ \alpha^n \beta, \gamma r, \gamma \alpha^{k_1} s_{u_1}, \gamma \alpha^{k_2} s_{u_2} \\ \alpha r, \alpha^2 r, \dots, \alpha^n r, \alpha^{n+2} r, \dots, \alpha^{2n} r \\ \alpha s_{u_1}, \alpha^2 s_{u_1}, \dots, \alpha^n s_{u_1}, \alpha^{n+2} s_{u_1}, \dots, \alpha^{2n} s_{u_1} \\ \alpha s_{u_2}, \alpha^2 s_{u_2}, \dots, \alpha^n s_{u_2}, \alpha^{n+2} s_{u_2}, \dots, \alpha^{2n} s_{u_2} \\ \alpha^n t, r t_1, r t_2, \alpha^n \\ t_1(\gamma r + \sum_{j \in \beta_1} \alpha^{n+1-j}), t_2(\gamma r + \sum_{j \in \beta_2} \alpha^{n+1-j}) \end{array} \right)$$

$$Q = (1)$$

$$f = \alpha^{n+1} r \beta (t_1 + t_2)$$

As in the previous case, the key point here is to prove that f is independent of (P, Q) .

Lemma 2. If $i_1 \in \beta_1$ and $i_2 \in \beta_2$, but $i_1 \notin \beta_2$ and $i_2 \notin \beta_1$, then (P, Q) are independent of f .

PROOF. The proof is given in §B.3 \square

THEOREM 3. For any probabilistic algorithm \mathcal{A} that totalizes at most q queries to the oracles performing group operations in $(\mathbb{G}, \mathbb{G}_T)$ and evaluations of $e(\cdot, \cdot)$

$$\text{Adv}^{\text{GDHE}}(\mathcal{A}) \leq \frac{(q + 2(6n + 10) + 2)^2}{p}$$

We will now generalize for an access policy $\mathbb{A} = \beta_1 \wedge \beta_2 \wedge \dots \wedge \beta_N$ consisting of N clauses and ℓ users u_1, u_2, \dots, u_ℓ . We will have

$$P = \begin{pmatrix} 1, r, \alpha r(\beta + s_{u_1}), \alpha r(\beta + s_{u_2}), \dots, \alpha r(\beta + s_{u_\ell}), \\ \alpha^n \beta, \gamma r, \gamma \alpha^{k_1} s_{u_1}, \gamma \alpha^{k_2} s_{u_2}, \dots, \gamma \alpha^{k_\ell} s_{u_\ell}, \\ \alpha r, \alpha^2 r, \dots, \alpha^n r, \alpha^{n+2} r, \dots, \alpha^{2n} r, \\ \alpha s_{u_1}, \alpha^2 s_{u_1}, \dots, \alpha^n s_{u_1}, \alpha^{n+2} s_{u_1}, \dots, \alpha^{2n} s_{u_1} \\ \alpha s_{u_2}, \alpha^2 s_{u_2}, \dots, \alpha^n s_{u_2}, \alpha^{n+2} s_{u_2}, \dots, \alpha^{2n} s_{u_2} \\ \vdots \\ \alpha s_{u_\ell}, \alpha^2 s_{u_\ell}, \dots, \alpha^n s_{u_\ell}, \alpha^{n+2} s_{u_\ell}, \dots, \alpha^{2n} s_{u_\ell} \\ \alpha^n t, r t_1, r t_2, \dots, r t_N, \alpha^n \\ t_1(\gamma r + \sum_{j \in \beta_1} \alpha^{n+1-j}), \dots, t_N(\gamma r + \sum_{j \in \beta_N} \alpha^{n+1-j}) \end{pmatrix}$$

$$Q = (1)$$

$$f = \alpha^{n+1} r \beta \sum_{i=1}^n t_i$$

Lemma 3. For every user u_i , $i \in [1, \ell]$, if $\exists j : k_i \notin \beta_j$, then (P, Q) are independent of f .

PROOF. The proof is given in §B.4. \square

We can now establish the following result, stating that even users colluding will have only a negligible advantage when trying to distinguish a ciphertext from a random value.

THEOREM 4. For any probabilistic algorithm \mathcal{A} that totalizes at most q queries to the oracles performing group operations in $(\mathbb{G}, \mathbb{G}_T)$ and evaluations of $e(\cdot, \cdot)$

$$\text{Adv}^{\text{GDHE}}(\mathcal{A}) \leq \frac{(q + 4n\ell + 4n + \ell + 6 + 2N)^2}{p}$$

Now, thanks to Lemmas 1 and 3, we can state the following theorem that proves the semantic security of our ABE scheme according to the model defined in §2.3.

THEOREM 5 (SEMANTIC SECURITY). Let \mathbb{G} be a bilinear group of prime order p . For any positive integers n, ℓ, N and R ($R < n$) our ABE scheme is semantically secure assuming the GDHE assumption holds. Moreover, the advantage of any probabilistic algorithm \mathcal{A} totalizing at most q queries to the oracles in distinguishing a valid ABE ciphertext from a random value is bounded by

$$\text{Adv}^{\text{ABBE}}(\mathcal{A}) \leq \frac{(q + 8n + 12 + 4R + 2)^2 + (q + 4n\ell + 4n + \ell + 6 + 2N)^2}{p}$$

Finally, it should be noted that we have proved the security of our ABE scheme for CNF expressions. The proof extends naturally to the DNF access policies, since in that case there is a concatenation of several independent instances of our scheme.

5. EFFICIENCY AND PRACTICAL ASPECTS

In this section, we discuss the complexity of our scheme and compare it against several other ciphertext-policy attribute-based (broadcast) encryption methods.

5.1 Complexity

There are several schemes implementing ciphertext-policy attribute-based (broadcast) encryption. For instance, in schemes such as [26], the access policy is expressed using a so-called linear secret sharing (LSSS) matrix M . Since linear secret sharing schemes [2] are described in terms of *authorized sets of attributes*, meaning that either set S_1 , or set S_2 , and so on, can decrypt the ciphertext, we note that it is more natural to talk in terms of DNF policies in that case. Moreover, in the two CP-ABE schemes described in [1], there

is a possibility to revoke individual users via an additional revocation method mathematically coupled with the main CP-ABE scheme (which also relies on linear secret sharing matrix M to describe the access policy) in order to make the global construction collusion-resistant. In [18], Lubicz and Sirvent propose an ABE based on access policies with **AND** and **NOT** gates. With the help of a Subset-Cover framework [20], this scheme can also implement the **OR** of the two gates above hence making it a DNF-type scheme. The authors also exhibit a solution on the way to perform direct user revocation with the Subset-Cover framework by adding $2n$ attributes and at most $\log_2(n) + 1$ new attributes to each user, n being the total number of users in the system. Below we provide two tables for comparing our scheme versus several others with direct user revocation in mind.

Lubicz and Sirvent scheme [18].

In this scheme (see §A for a quick review of it), given n attributes in the system, the encryption key contains $3n + 2$ group elements and n elements of $(\mathbb{Z}/p\mathbb{Z})^*$. It should be noted that if we would like to revoke individual receivers, 2ℓ new attributes are added to the system and each user will belong to $\log_2(\ell) + 1$ additional groups. The revocation is achieved using the SD-method [20]. The decryption key of a user u with $\kappa(u)$ attributes contains $\kappa(u) + 2$ group elements and $\kappa(u)$ elements of $(\mathbb{Z}/p\mathbb{Z})^*$. For a DNF access policy with N clauses (this scheme can only handle this type of access policies by concatenation) having R revoked attributes, the size of the header is $N \cdot (R + 2)$ group elements and $2N$ elements of $(\mathbb{Z}/p\mathbb{Z})^*$. It should also be noted that in case of direct user revocation, R will be a function of r , the number of revoked users. The decryption time is mainly given by the time to perform the $N \cdot \kappa$ group exponentiations.

Attrapadung and Imai schemes [1].

There are two CP-ABE schemes, both having an explicit capability of conjunctively revoking individual receivers; the schemes rely on the LSSS technique, hence implying DNF formulas. The individual revocation is achieved by mathematically joining a BE scheme to a CP-ABE. This translates into adding an **AND** gate with a disjunctive list of authorized receivers to the DNF expression. In these schemes, there is a maximum allowed number m for the attribute set within an individual user and κ_{max} - the maximum number of attributes in the access policy. Since these parameters should be fixed at the system deployment, it can already be seen as a limiting factor. In the BCP-ABE1 scheme authors are combining mathematically the Waters [26] scheme (for attribute-based broadcast) with the Boneh-Gentry-Waters [6] broadcast encryption scheme for direct revocation. It should be noted that in the case of BGW scheme, the receiver must store the public key to be able to decrypt the ciphertext. It means that the private key size is $O(n)$ - comparable to our scheme. The advantage of our scheme is that there is no limiting factors on the attribute set and the number of attributes per clause that should be fixed prior to system deployment.

Our scheme.

With n attributes and ℓ users, the encryption and the decryption keys contain $O(n + \ell)$ group elements. For a CNF access policy with N clauses, the size of the header is $2N + 1$ group elements and this, independently of the

	DNF (with N clauses)				CNF (with N clauses)			
	this paper	[18]	[1] ^{CP-ABE1}	[1] ^{CP-ABE2}	this paper	[18]	[1] ^{CP-ABE1}	[1] ^{CP-ABE2}
ciphertext	$O(N \cdot M)$	$O(N \cdot R)^b$	$O(t)^c$	$O(\kappa_{max} + r)^c$	$O(N)$	–	–	–
decryption	$O(M + n)$	$O(N \cdot \kappa)$	$O(t)^c$	$O(\kappa_{max} + r)^c$	$O(N)$	–	–	–

Table 1: Bandwidth and decryption complexity comparison. N - number of clauses in a policy, i.e. $\mathbb{A} = \beta_1 \vee \dots \vee \beta_N$, M - maximum number of attributes in a given clause, i.e. $\beta_i = A_1 \wedge \dots \wedge A_M$, n - total number of attributes in the system, ℓ - total number of users, m - maximum number of attributes within individual user, R - number of revoked (negated attributes) in a clause, r - number of revoked users, κ_{max} - maximum number of attributes in the access policy, t - number of attributes in the access policy, κ - number of attributes for a given user (positive and negated). We assume that $\ell \gg n$.

	this paper	[18]	[1] ^{CP-ABE1}	[1] ^{CP-ABE2}
encryption key size	$O(n + \ell)$	$O(n + \ell)$	$O(n + \ell)$	$O(m + \kappa_{max})$
private key size	$O(n + \ell)$	$O(\kappa + \log_2(\ell))$	$O(n + \ell)$	$O(n + \kappa)$

Table 2: Key storage complexity comparison.

number of attributes inside each clause. The direct user revocation is achieved using only 1 additional clause. The decryption of one clause of M attributes will be dominated by two pairing operations. The final session key computation is given by one additional pairing operation. For a DNF access policy, the size of the header is $O(N \cdot M)$, where M is the average number of attributes per clause. Hence, as it has already been pointed out, our scheme is hence naturally suitable for CNF type of expressions. We also emphasize that our scheme accepts CNF or DNF expressions, which are the general description of any possible formula. Hence a logical formula needs to be transformed into CNF or DNF form first.

6. CONCLUSION AND OPEN PROBLEMS

We have proposed a new ABBE scheme which allows performing encryptions based on different access policies expressed either in CNF or DNF form along with efficient individual receiver revocation ability. Since these two forms is the most general way of expressing Boolean access policies, we are relying on it to achieve the generality that other ABBE scheme do not necessarily provide. The security of our scheme is proven in the generic model of groups with pairing. While we understand that a security proof in a more tight assumption might be seen as a plus, we leave the proposal for an efficient and flexible ABBE with a security proof in the standard assumption (i.e., q -BDHE) as an open problem.

7. REFERENCES

- [1] N. Attrapadung and H. Imai. Conjunctive broadcast and attribute-based encryption. In H. Shacham and B. Waters, editors, *Pairing-Based Cryptography - Pairing 2009, Third International Conference, Palo Alto, CA, USA, August 12-14, 2009. Proceedings*,

^b R is a function of r in this case.

^cPlus the transmission of a $\kappa_{max} \times t$ access structure matrix Π with elements in $\mathbb{Z}/p\mathbb{Z}$.

volume 5671 of *Lecture Notes in Computer Science*, pages 248–265. Springer-Verlag, 2009.

- [2] J. Benaloh. General linear secret sharing. Unpublished manuscript available at <http://research.microsoft.com/pubs/68477/glss.ps>, 1996.
- [3] S. Berkovits. How to broadcast a secret. In D. Davies, editor, *Advances in Cryptology - EUROCRYPT '91: Workshop on the Theory and Application of Cryptographic Techniques, Brighton, UK, April 1991. Proceedings*, volume 547 of *Lecture Notes in Computer Science*, pages 535–541. Springer-Verlag, 1991.
- [4] J. Bethencourt, A. Sahai, and B. Waters. Ciphertext-policy attribute-based encryption. In *2007 IEEE Symposium on Security and Privacy (S&P 2007), 20-23 May 2007, Oakland, California, USA*, pages 321–334. IEEE Computer Society, 2007.
- [5] D. Boneh, X. Boyen, and J.-E. Goh. Hierarchical identity based encryption with constant size ciphertext. In R. Cramer, editor, *Advances in Cryptology - EUROCRYPT 2005: 24th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Aarhus, Denmark, May 22-26, 2005. Proceedings*, volume 3494 of *Lecture Notes in Computer Science*, pages 440–456. Springer-Verlag, 2005.
- [6] D. Boneh, C. Gentry, and B. Waters. Collusion resistant broadcast encryption with short ciphertexts and private keys. In V. Shoup, editor, *Advances in Cryptology - CRYPTO 2005, 25th Annual International Cryptology Conference, Santa Barbara, California, USA, August 14-18, 2005. Proceedings*, volume 3621 of *Lecture Notes in Computer Science*, pages 258–275. Springer-Verlag, 2005.
- [7] D. Boneh and B. Waters. A fully collusion resistant broadcast, trace and revoke system. In A. Juels, R. Wright, and S. De Capitani di Vimercati, editors, *Proceedings of the 13th ACM Conference on Computer and Communications Security, CCS 2006, Alexandria,*

- VA, USA, October 30 - November 3, 2006., pages 211–220. ACM Press, 2006.
- [8] R. Canetti, J. Garay, G. Itkis, D. Micciancio, M. Naor, and B. Pinkas. Multicast security: a taxonomy and some efficient constructions. In *Proceedings of IEEE INFOCOM'99, The Conference on Computer Communications, Eighteenth Annual Joint Conference of the IEEE Computer and Communications Societies, March 21–25 1999, New-York, NY, USA*, volume 2, pages 708–716. IEEE, 1999.
- [9] R. Canetti, T. Malkin, and K. Nissim. Efficient communication-storage tradeoffs for multicast encryption. In J. Stern, editor, *Advances in Cryptology – EUROCRYPT'99: International Conference on the Theory and Application of Cryptographic Techniques, Prague, Czech Republic, May 1999. Proceedings*, volume 1592 of *Lecture Notes in Computer Science*, pages 459–474. Springer-Verlag, 1999.
- [10] M. Chase. Multi-authority attribute based encryption. In S. Vadhan, editor, *Theory of Cryptography, 4th Theory of Cryptography Conference, TCC 2007, Amsterdam, The Netherlands, February 21–24, 2007. Proceedings*, volume 4392 of *Lecture Notes in Computer Science*, pages 515–534. Springer-Verlag, 2007.
- [11] C. Delerablée, P. Paillier, and D. Pointcheval. Fully collusion secure dynamic broadcast encryption with constant-size ciphertexts or decryption keys. In T. Takagi, T. Okamoto, E. Okamoto, and T. Okamoto, editors, *Pairing-Based Cryptography - Pairing 2007, First International Conference, Tokyo, Japan, July 2–4, 2007. Proceedings*, volume 4575 of *Lecture Notes in Computer Science*, pages 39–59. Springer-Verlag, 2007.
- [12] Y. Dodis and N. Fazio. Public-key broadcast encryption for stateless receivers. In J. Feigenbaum, editor, *Security and Privacy in Digital Rights Management, ACM CCS-9 Workshop, DRM 2002, Washington, DC, USA, November 18, 2002. Revised Papers*, volume 2696 of *Lecture Notes in Computer Science*, pages 61–80. Springer-Verlag, 2002.
- [13] A. Fiat and M. Naor. Broadcast encryption. In D. Stinson, editor, *Advances in Cryptology – CRYPTO'93: 13th Annual International Cryptology Conference, Santa Barbara, California, USA, August 22–26, 1993. Proceedings*, volume 773 of *Lecture Notes in Computer Science*, pages 480–491. Springer-Verlag, 1994.
- [14] M. Goodrich, J. Sun, and R. Tamassia. Efficient tree-based revocation in groups of low-state devices. In M. Franklin, editor, *Advances in Cryptology – CRYPTO 2004, 24th Annual International Cryptology Conference, Santa Barbara, California, USA, August 15–19, 2004. Proceedings*, volume 3152 of *Lecture Notes in Computer Science*, pages 511–527. Springer-Verlag, 2004.
- [15] V. Goyal, A. Jain, O. Pandey, and A. Sahai. Bounded ciphertext-policy attribute based encryption. In L. Aceto, I. Damgård, L. A. Goldberg, M. Halldórsson, A. Ingólfssdóttir, and I. Walukiewicz, editors, *Automata, Languages and Programming, 35th International Colloquium, ICALP 2008, Reykjavik, Iceland, July 7–11, 2008, Proceedings, Part II - Track B: Logic, Semantics, and Theory of Programming & Track C: Security and Cryptography Foundations*, volume 5126 of *Lecture Notes in Computer Science*, pages 579–591. Springer-Verlag, 2008.
- [16] V. Goyal, O. Pandey, A. Sahai, and B. Waters. Attribute-based encryption for fine-grained access control of encrypted data. In A. Juels, R. Wright, and S. De Capitani di Vimercati, editors, *Proceedings of the 13th ACM Conference on Computer and Communications Security, CCS 2006, Alexandria, VA, USA, October 30 - November 3, 2006.*, pages 72–81. ACM Press, 2006.
- [17] D. Halevy and A. Shamir. The LSD broadcast encryption scheme. In M. Yung, editor, *Advances in Cryptology – CRYPTO 2002: 22nd Annual International Cryptology Conference Santa Barbara, California, USA, August 18–22, 2002. Proceedings*, volume 2442 of *Lecture Notes in Computer Science*, pages 47–60. Springer-Verlag, 2002.
- [18] D. Lubicz and T. Sirvent. Attribute-based broadcast encryption scheme made efficient. In S. Vaudenay, editor, *Progress in Cryptology - AFRICACRYPT 2008, First International Conference on Cryptology in Africa, Casablanca, Morocco, June 11–14, 2008. Proceedings*, volume 5023 of *Lecture Notes in Computer Science*, pages 325–342. Springer-Verlag, 2008.
- [19] S. Müller, S. Katzenbeisser, and C. Eckert. Distributed attributed-based encryption. In P. J. Lee and J. H. Cheon, editors, *Information Security and Cryptology – ICISC 2008: 11th International Conference, Seoul, Korea, December 3–5, 2008. Revised Selected Papers*, volume 5461 of *Lecture Notes in Computer Science*, pages 20–36. Springer-Verlag, 2008.
- [20] D. Naor, M. Naor, and J. Lotspiech. Revocation and tracing schemes for stateless receivers. In J. Kilian, editor, *Advances in Cryptology – CRYPTO 2001: 21st Annual International Cryptology Conference, Santa Barbara, California, USA, August 19–23, 2001. Proceedings*, volume 2139 of *Lecture Notes in Computer Science*, pages 41–62. Springer-Verlag, 2001.
- [21] R. Ostrovsky, A. Sahai, and B. Waters. Attribute-based encryption with non-monotonic access structures. In P. Ning, S. De Capitani di Vimercati, and P. Syverson, editors, *Proceedings of the 2007 ACM Conference on Computer and Communications Security, CCS 2007, Alexandria, VA, USA, October 28–31, 2007.*, pages 195–203. ACM Press, 2007.
- [22] A. Perrig, D. Song, and D. Tygar. ELK, a new protocol for efficient large-group key distribution. In *Proceedings of the IEEE Symposium on Security and Privacy, 14–16 May, 2001, Oakland, California, USA*, pages 247–262. IEEE Computer Society, 2001.
- [23] A. Sahai and B. Waters. Fuzzy identity-based encryption. In R. Cramer, editor, *Advances in Cryptology – EUROCRYPT 2005: 24th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Aarhus, Denmark, May 22–26, 2005. Proceedings*, volume 3494 of *Lecture Notes in Computer Science*, pages 457–473. Springer-Verlag, 2005.

- [24] V. Shoup. Lower bounds for discrete logarithms and related problems. In W. Fumy, editor, *Advances in Cryptology – EUROCRYPT ’97: International Conference on the Theory and Application of Cryptographic Techniques, Konstanz, Germany, May 1997. Proceedings*, volume 1233 of *Lecture Notes in Computer Science*, pages 256–266. Springer-Verlag, 1997.
- [25] D. Wallner, E. Harder, and R. Agee. Key management for multicast: issues and architectures. RFC 2627, 1999. Available on <http://www.ietf.org>.
- [26] B. Waters. Ciphertext-policy attribute-based encryption: an expressive, efficient, and provably secure realization. Available on <http://eprint.iacr.org/2008/290/>, 2008.
- [27] C. Wong, M. Gouda, and S. Lam. Secure group communications using key graphs. In *Proceedings of the ACM SIGCOMM’98 Conference on Applications, Technologies, Architectures, and Protocols for Computer Communication, August 31 - September 4, 1998, Vancouver, British Columbia, Canada*, pages 68–79. ACM Press, 1998.

APPENDIX

A. THE LUBICZ-SIRVENT SCHEME

We briefly review the public-key ABBE scheme disclosed by Lubicz and Sirvent [18]. It is important to note that they consider Boolean access policies in DNF consisting of a *single* clause, such as $\mathbb{A} = A_1 \wedge A_2 \wedge \dots \wedge A_r \wedge \dots \wedge A_{r+s}$. As before, we will denote by \mathcal{B} and $\overline{\mathcal{B}}$ the sets of positive attributes $\mathcal{B} = \{A_1, \dots, A_r\}$ and of negative attributes $\overline{\mathcal{B}} = \{A_{r+1}, \dots, A_{r+s}\}$, respectively. The **Setup**(.) algorithm is specified as follows: four elements $\alpha, \beta, \gamma, \delta$ are chosen uniformly at random in $(\mathbb{Z}/p\mathbb{Z})^*$. Each group of users identified to a Boolean attribute A_i is then associated with an element $\mu_i \in_R (\mathbb{Z}/p\mathbb{Z})^*$, simply called “attribute”, such that all these elements are pairwise different and different from α . Another attribute μ_0 chosen under the same constraints will correspondant to an attribute assigned to *no* user. The encryption key ek is defined as

$$\text{ek} = \left(g, g^{\beta\gamma\delta}, (\mu_j)_{0 \leq j \leq n}, \left(g^{\alpha^j} \right)_{0 \leq j \leq n}, \left(g^{\gamma\alpha^j} \right)_{0 \leq j \leq n}, \left(g^{\delta\alpha^j} \right)_{0 \leq j \leq n} \right).$$

Each user $u_i \in \mathcal{U}$ with $1 \leq i \leq \ell$ is assigned a value $s_{u_i} \in_R (\mathbb{Z}/p\mathbb{Z})^*$. Let $\Omega(u_i)$ be the set of attributes corresponding to the groups he belongs to: $\Omega(u_i) = \{\mu_j : j \in \mathfrak{B}(u_i)\}$ and let us denote $\kappa(u_i) = \#\Omega(u_i)$ its cardinality. Finally, let $\Pi(u_i) = \prod_{\mu \in \Omega(u_i)} (\alpha - \mu)$. Then, the decryption key of user u_i is defined as

$$\text{dk}_i = \left(\Omega(u_i), g^{\delta(\beta+s_{u_i})}, g^{s_{u_i}\Pi(u_i)\gamma}, \left(g^{\gamma\delta s_{u_i}\alpha^j} \right)_{0 \leq j \leq \kappa(u_i)} \right).$$

We now describe the **Encrypt**(.) algorithm. A trivial case occurs when $\mathcal{B} \cap \overline{\mathcal{B}} \neq \emptyset$: in that case, **Encrypt**(.) returns \perp , since a user cannot simultaneously be inside and outside a given group of users. Let $\Omega = \{\mu_j : i \in \mathcal{B}\}$ and $\overline{\Omega} = \{\mu_j : i \in \overline{\mathcal{B}}\}$, as well as their respective cardinalities $\kappa = \#\Omega$ and $\overline{\kappa} = \#\overline{\Omega}$. Let $\Pi = \prod_{\mu \in \Omega} (\alpha - \mu)$, $\overline{\Pi} = \prod_{\mu \in \overline{\Omega}} (\alpha - \mu)$ and $\overline{\overline{\Pi}} = \prod_{\mu \in \overline{\Omega}} (\alpha - \mu)$. Let $z \in_R (\mathbb{Z}/p\mathbb{Z})^*$. The result of the encryption

operation is given by

$$h = \left(\Pi, \overline{\Pi}, g^{z\overline{\Pi}}, g^{\gamma z \Pi}, \left(g^{\delta z \alpha^j} \right)_{0 \leq j \leq \overline{\kappa}} \right) \text{ and } k = h^{\beta\gamma\delta z \Pi}.$$

When $\overline{\mathcal{B}} = \emptyset$, the encryption algorithm considers that the virtual group containing no user is revoked; hence, then $\overline{\Omega} = \{\mu_0\}$ and $\overline{\kappa} = 1$. Finally, the **Decrypt**(.) algorithm works as follows for the user u_i : if $\Omega \subseteq \Omega(u_i)$ and if $\overline{\Omega} \cap \Omega(u_i) = \emptyset$, then the user u_i is able to decrypt the header h . For this, he uses the extended Euclidean algorithm over the polynomials

$$\prod_{\mu \in (\Omega \cup \overline{\Omega})} (X - \mu) \text{ and } \prod_{\mu \in \Omega(u_i)} (X - \mu).$$

He obtains then two unitary polynomials

$$V(X) = \sum_{0 \leq j \leq \kappa} v_j X^j \text{ and } W(X) = \sum_{0 \leq j \leq \overline{\kappa}} w_j X^j$$

in $(\mathbb{Z}/p\mathbb{Z})[X]$ such that

$$V(X) \prod_{\mu \in (\Omega \cup \overline{\Omega})} (X - \mu) + W(X) \prod_{\mu \in \Omega(u_i)} (X - \mu) = \prod_{\mu \in \Omega} (X - \mu).$$

Given the header $h = \left(\Pi, \overline{\Pi}, g^{z\overline{\Pi}}, g^{\gamma z \Pi}, \left(g^{\delta z \alpha^j} \right)_{0 \leq j \leq \overline{\kappa}} \right)$ and his decryption key dk_i , the user u_i can recover the session key k by computing

$$k = \frac{e \left(g^{\delta(\beta+s_{u_i})}, g^{\gamma z \Pi} \right)}{e \left(\prod_{j=0}^{\kappa-1} g^{v_j \gamma \delta s_{u_i} \alpha^j}, g^{z\overline{\Pi}} \right) e \left(g^{s_{u_i} \Pi(u_i) \gamma}, \prod_{j=0}^{\overline{\kappa}-1} g^{w_j \delta z \alpha^j} \right)}.$$

Lubicz and Sirvent prove the security of their scheme relatively to an ad-hoc assumption which is an extension of the decisional version of the General Diffie-Hellman Exponent (GDHE) problem studied in [5]; furthermore, they assess the security of this assumption within the framework of the generic model of groups with pairings. We refer the reader to [18] for the details.

B. PROOFS

B.1 Soundness

Provided a policy with N clauses $\beta_1 \wedge \beta_2 \wedge \dots \wedge \beta_N$ and given that the receiver is associated with at least one attribute from every clause β_i by the mean of a private key, the authorized receiver would be able to compute, for every

i , the session subkey SK_i^{su} :

$$\begin{aligned}
SK_i^{su} &= \frac{e(g_k^{su}, C_1)}{e\left(d_k \cdot \prod_{j \in \beta_i, j \neq k} g_{n+1-j+k}^{su}, C_0\right)} \\
&= \frac{e\left(g_k^{su}, (v^r \prod_{j \in \beta_i} g_{n+1-j}^r)^{t_i}\right)}{e\left(g_k^{\gamma \cdot su} \cdot \prod_{j \in \beta_i, j \neq k} g_{n+1-j+k}^{su}, g^{rt_i}\right)} \\
&= \frac{e\left(g_k^{su}, g^{\gamma t_i r}\right) \cdot e\left(g_k^{su}, \prod_{j \in \beta_i} g_{n+1-j}^r\right)^{t_i}}{e\left(g_k^{\gamma \cdot su} \cdot \prod_{j \in \beta_i, j \neq k} g_{n+1-j+k}^{su}, g^{rt_i}\right)} \\
&= \frac{e(g_k, g)^{su \gamma t_i r} \cdot e\left(g_k^{su}, \prod_{j \in \beta_i} g_{n+1-j}^r\right)^{t_i}}{e\left(g_k^{\gamma \cdot su}, g^{rt_i}\right) \cdot e\left(\prod_{j \in \beta_i, j \neq k} g_{n+1-j+k}^{su}, g^{rt_i}\right)} \\
&= \frac{e\left(g_k^{su}, g_{n+1-k}^r \cdot \prod_{j \in \beta_i, j \neq k} g_{n+1-j}^r\right)^{t_i}}{e\left(\prod_{j \in \beta_i, j \neq k} g_{n+1-j+k}^{su}, g\right)^{rt_i}} \\
&= \frac{e\left(g_k^{su}, g_{n+1-k}^r\right)^{t_i} \cdot e\left(g_k^{su}, \prod_{j \in \beta_i, j \neq k} g_{n+1-j}^r\right)^{rt_i}}{e\left(\prod_{j \in \beta_i, j \neq k} g_{n+1-j+k}^{su}, g\right)^{rt_i}} \\
&= \frac{e\left(g_k^{su}, g_{n+1-k}^r\right)^{t_i} \cdot e\left(g^{su}, \prod_{j \in \beta_i, j \neq k} g_{n+1-j+k}\right)^{rt_i}}{e\left(\prod_{j \in \beta_i, j \neq k} g_{n+1-j+k}^{su}, g\right)^{rt_i}} \\
&= e(g, g)^{\alpha^{n+1} r s_u t_i}.
\end{aligned}$$

The final session key SK is then computed as

$$\begin{aligned}
SK &= \frac{e(\text{hdr}_0, g_1^{r(\beta+s_u)})}{\prod_{i=1}^N SK_i^{su}} = \frac{e\left(g_1^t, g_1^{r(\beta+s_u)}\right)}{\prod_{i=1}^N e(g, g)^{\alpha^{n+1} r s_u t_i}} \\
&= \frac{e(g, g)^{r \alpha^{n+1} t (\beta+s_u)}}{e(g, g)^{\alpha^{n+1} r s_u \sum_{i=1}^n t_i}} \\
&= \frac{e(g, g)^{r \alpha^{n+1} t \beta} \cdot e(g, g)^{r \alpha^{n+1} s_u t}}{e(g, g)^{r \alpha^{n+1} s_u t}} \\
&= e(g, g)^{\beta r \alpha^{n+1} t}
\end{aligned}$$

B.2 Proof of Lemma 1

We start by writing all terms with β :

$$\alpha r(\beta + s_u), \alpha^n \beta$$

Multiplying the two values by any other term gives us:

$$\left(\begin{array}{l} \alpha r(\beta + s_u), \alpha r^2(\beta + s_u), \alpha^{n+1} r \beta(\beta + s_u), \gamma \alpha r^2(\beta + s_u), \\ \gamma \alpha^{k+1} r s_u(\beta + s_u), \alpha^2 r^2(\beta + s_u), \alpha^3 r^2(\beta + s_u), \dots, \\ \alpha^{n+1} r^2(\beta + s_u), \alpha^{n+3} r^2(\beta + s_u), \dots, \\ \alpha^{2n+1} r^2(\beta + s_u), \\ \alpha^2 r s_u(\beta + s_u), \alpha^3 r s_u(\beta + s_u), \dots, \\ \alpha^{n+1} r s_u(\beta + s_u), \alpha^{n+3} r s_u(\beta + s_u), \dots, \\ \alpha^{2n+1} r s_u(\beta + s_u), \alpha^{n+1} r t(\beta + s_u), \\ \alpha r^2 t_1(\beta + s_u), \dots, \alpha r^2 t_R(\beta + s_u), \alpha^{n+1} r(\beta + s_u), \\ \alpha r(\beta + s_u) t_{i_1} (\gamma r + r \sum_{j \in \beta_{i_1}} \alpha^{n+1-j}), \dots, \\ \alpha r(\beta + s_u) t_{i_R} (\gamma r + r \sum_{j \in \beta_{i_R}} \alpha^{n+1-j}) \end{array} \right)$$

and

$$\left(\begin{array}{l} \alpha^n \beta, \alpha^n \beta r, \alpha^{n+1} r \beta(\beta + s_u), \alpha^{2n} \beta^2, \alpha^n \beta \gamma r, \alpha^{n+k} \gamma s_u \beta, \\ \alpha^{n+1} \beta r, \alpha^{n+2} \beta r, \dots, \alpha^{2n} \beta r, \dots, \alpha^{2n+2} \beta r, \dots, \alpha^{3n} \beta r, \\ \alpha^{n+1} s_u \beta, \alpha^{n+2} s_u \beta, \dots, \alpha^{2n} s_u \beta, \\ \alpha^{2n+2} s_u \beta, \dots, \alpha^{3n} s_u \beta, \alpha^{2n} \beta t, \alpha^n \beta r t_1, \dots, \alpha^n \beta r t_1, \alpha^{2n} \beta, \\ \alpha^n \beta t_{i_1} (\gamma r + r \sum_{j \in \beta_{i_1}} \alpha^{n+1-j}), \dots, \\ \alpha^n \beta t_{i_R} (\gamma r + r \sum_{j \in \beta_{i_R}} \alpha^{n+1-j}) \end{array} \right)$$

Terms that have $\alpha^{n+1} \beta$ include

$$\left(\begin{array}{l} \alpha^{n+1} r \beta(\beta + s_u), \alpha^{k+1} r s_u(\beta + s_u), \alpha^{n+1} r^2(\beta + s_u), \\ \alpha^{n+1} r s_u(\beta + s_u), \alpha^{n+1} r t(\beta + s_u), \alpha^{n+1} r(\beta + s_u), \\ \alpha^{n+1} r \beta(\beta + s_u), \alpha^{n+k} \gamma s_u \beta, \alpha^{n+1} \beta r, \alpha^{n+1} s_u \beta \end{array} \right).$$

We can now notice that the only term having a t is

$$\alpha^{n+1} r t(\beta + s_u) = \alpha^{n+1} r t \beta + \alpha^{n+1} r t s_u.$$

The only way to obtain the correct session key consists in removing the term $\alpha^{n+1} r t s_u$. The terms containing t in P are

$$\left(\begin{array}{l} \alpha^n t, r t, t_{i_1} (\gamma r + r \sum_{j \in \beta_{i_1}} \alpha^{n+1-j}), \dots, \\ t_{i_R} (\gamma r + r \sum_{j \in \beta_{i_R}} \alpha^{n+1-j}) \end{array} \right).$$

There is no way to construct $\alpha^{n+1} r t s_u$ from $\alpha^n t, r t$ and any other term. The only possibility consists in computing for all t_{i_j} :

$$\begin{aligned}
&\alpha^k s_u (t_{i_1} \gamma r + t_{i_1} r \sum_{j \in \beta_{i_1}} \alpha^{n+1-j}) - \\
&r t (\alpha^k \gamma s_u + s_u \sum_{j \in \beta_{i_1}, j \neq k} \alpha^{n+1-j+k}) = \\
&\alpha^k s_u t \gamma + t \alpha^k s_u r \sum_{j \in \beta_{i_1}} \alpha^{n+1-j} - \\
&r t \alpha^k \gamma s_u - r t s_u \sum_{j \in \beta_{i_1}, j \neq k} \alpha^{n+1-j+k} = \\
&t \alpha^k s_u r \sum_{j \in \beta_{i_1}} \alpha^{n+1-j} - r t s_u \sum_{j \in \beta_{i_1}, j \neq k} \alpha^{n+1-j+k}
\end{aligned}$$

But since $d_{i_j} \in \mathcal{G} \setminus (\beta_1 \cap \beta_2 \dots \cap \beta_n)$, there is no such $j \in \beta_{i_1}$ that we could compute

$$\begin{aligned}
&t \alpha^k s_u r \sum_{j \in \beta_{i_1}} \alpha^{n+1-j} - r t s_u \sum_{j \in \beta_{i_1}, j \neq k} \alpha^{n+1-j+k} = \\
&t \alpha^{k+n+1-k} s_u r + t s_u r \sum_{j \in \beta_{i_1}, j \neq k} \alpha^{n+1-j+k} - \\
&r t s_u \sum_{j \in \beta_{i_1}, j \neq k} \alpha^{n+1-j+k} = t \alpha^{n+1} s_u r
\end{aligned}$$

B.3 Proof of Lemma 2

As in the previous proof, the only two terms having a t are

$$\alpha^{n+1} r t(\beta + s_{u_1}) = \alpha^{n+1} r t \beta + \alpha^{n+1} r t s_{u_1}$$

and

$$\alpha^{n+1} r t(\beta + s_{u_2}) = \alpha^{n+1} r t \beta + \alpha^{n+1} r t s_{u_2}.$$

Hence, to obtain a valid session key, the adversary needs to remove (de-blind) either

$$\alpha^{n+1} r t s_{u_1} = \alpha^{n+1} r t_1 s_{u_1} + \alpha^{n+1} r t_2 s_{u_1}$$

or

$$\alpha^{n+1} r t s_{u_2} = \alpha^{n+1} r t_1 s_{u_2} + \alpha^{n+1} r t_2 s_{u_2}.$$

But since $k_1 \notin \beta_2$ and $k_2 \notin \beta_1$, neither $\alpha^{n+1} r t_1 s_{u_2}$ nor $\alpha^{n+1} r t_2 s_{u_1}$ can be removed even if the two users u_1 and u_2 collude.

B.4 Proof of Lemma 3

Following the same intuition as in the two previous cases, there will be ℓ terms having a t , namely

$$\alpha^{n+1}rt(\beta + s_{u_1}), \alpha^{n+1}rt(\beta + s_{u_2}), \dots, \alpha^{n+1}rt(\beta + s_{u_\ell})$$

For an i , we would have $\alpha^{n+1}rt(\beta + s_{u_i}) = \alpha^{n+1}rt\beta + \alpha^{n+1}rts_{u_i}$. The second term is $\alpha^{n+1}rts_{u_i} = \sum_{j=1}^N \alpha^{n+1}rt_j s_{u_i}$. The only way to construct $\sum_{j=1}^N \alpha^{n+1}rt_j s_{u_i}$ consists in computing

$$\sum_{v=1}^N \left(\alpha^{k_i} s_{u_i} (t_v \gamma r + t_v r \sum_{j \in \beta_v} \alpha^{n+1-j}) - \right. \\ \left. rt(\alpha^k \gamma s_{u_i} + s_{u_i} \sum_{j \in \beta_v, j \neq k_i} \alpha^{n+1-j+k}) \right)$$

But since there is at least one k_i s.t. $k_i \notin \beta_v$, it is not possible.