

Mediated Ciphertext-Policy Attribute-Based Encryption and Its Application

Luan Ibraimi^{1,2}, Milan Petkovic², Svetla Nikova¹, Pieter Hartel¹,
and Willem Jonker^{1,2}

¹ Faculty of EEMCS, University of Twente, The Netherlands

² Philips Research, The Netherlands

Abstract. In Ciphertext-Policy Attribute-Based Encryption (CP-ABE), a user secret key is associated with a set of attributes, and the ciphertext is associated with an access policy over attributes. The user can decrypt the ciphertext if and only if the attribute set of his secret key satisfies the access policy specified in the ciphertext. Several CP-ABE schemes have been proposed, however, some practical problems, such as attribute revocation, still needs to be addressed. In this paper, we propose a mediated Ciphertext-Policy Attribute-Based Encryption (mCP-ABE) which extends CP-ABE with instantaneous attribute revocation. Furthermore, we demonstrate how to apply the proposed mCP-ABE scheme to securely manage Personal Health Records (PHRs).

1 Introduction

Modern distributed information systems require flexible access control models which go beyond discretionary, mandatory and role-based access control. Recently proposed models, such as attribute-based access control, define access control policies based on different attributes of the requester, environment, or the data object. On the other hand, the current trend of service-based information systems and storage outsourcing require increased protection of data including access control methods that are cryptographically enforced. The concept of Attribute-Based Encryption (ABE) fulfills the aforementioned requirements. It provides an elegant way of encrypting data such that the encryptor defines the attribute set that the decryptor needs to possess in order to decrypt the ciphertext. Since Sahai and Waters [1] proposed the basic ABE scheme, several more advanced schemes have been developed, such as most notably Ciphertext-Policy ABE schemes (CP-ABE) [2,3,4]. In these schemes, a ciphertext is associated with an access policy and the user secret key is associated with a set of attributes. A secret key holder can decrypt the ciphertext if the attributes associated with his secret key satisfy the access policy associated with the ciphertext. For example, consider a situation when two organizations, a Hospital and a University, conduct research in the field of neurological disorders. The Hospital wants to allow access to their research results to all staff from the University who have the role Professor and belong to the Department of Neurology (DN). To enforce the policy, the Hospital encrypts the data according to the access policy

$\tau_{Results} = (\text{University Professor} \wedge \text{Member of DN})$. Only users who have a secret key associated with a set of attributes $\omega = (\text{University Professor}, \text{Member of DN})$ can satisfy the access policy $\tau_{Results}$ and be able to decrypt the ciphertext.

The state-of-the-art CP-ABE schemes provide limited support for revocation of attributes, a feature, which is becoming increasingly important in modern access control systems. In general, attribute revocation may happen due to the following reasons: 1) an attribute is not valid because it has expired, for instance, the attribute "project manager-January 2009" is valid until January 2009, or 2) a user is misusing her secret key associated with a set of attributes, for instance, Alice might give a copy of her secret key to Bob who is not a legitimate user. In particular, attribute revocation is an important requirement in the domain of access control to personal health data, which is our application field for attribute-based encryption.

Contribution. In this paper, we propose a new scheme for attribute revocation in CP-ABE called mediated Ciphertext-Policy Attribute-Based Encryption (mCP-ABE). Previous CP-ABE systems proposed to use a system where attributes are valid within a specific time frame [5]. However, the drawback of this approach is that there is no way to revoke an attribute before the expiration date. In our scheme the secret key is divided into two shares, one share for the mediator and the other for the user. To decrypt the data, the user must contact the mediator to receive a decryption token. The mediator keeps an attribute revocation list (ARL) and refuses to issue the decryption token for revoked attributes. Without the token, the user cannot decrypt the ciphertext, therefore the attribute is implicitly revoked. In our scheme we assume that each user has a unique identifier I_u (in CP-ABE the user is identified only with a set of attributes) and may have many attributes. The identifier is used by the mediator to check if there are revoked attributes related to I_u . Different users having different identifiers, may have the same attribute set. For example, Alice with an identifier I_{Alice} , and Bob with an identifier I_{Bob} , may have the same attribute set $\omega = (att_1, att_2)$. The technique of splitting the attribute components of the secret key into two shares, and the technique of using an identifier I_u for each user, helps us to achieve the following attribute revocations: i) revoking an attribute from a single user without affecting other users, and ii) revoking an attribute from the system where all users are affected.

We also define a security model for the proposed scheme which formalizes the security attacks and provide a security proof under the generic group model. Finally, we demonstrate the applicability of the proposed scheme to securely manage Personal Health Records (PHRs).

1.1 Related Work

Attribute-Based Encryption. Sahai and Waters in their seminal paper [1] introduce the concept of ABE. There are two types of ABE schemes: Key-Policy ABE schemes (KP-ABE) [6] and Ciphertext-Policy ABE schemes (CP-ABE) [2,3,4]. In KP-ABE, a ciphertext is associated with a set of attributes and a user secret

key is associated with an access policy. A secret key holder can decrypt the ciphertext if the attributes associated with the ciphertext satisfy the access policy associated with the secret key. In CP-ABE the idea is reversed. A ciphertext is associated with an access policy and the user secret key is associated with a set of attributes. A secret key holder can decrypt the ciphertext if the attributes associated with the secret key satisfy the access policy associated with the ciphertext.

Mediated Cryptography. Boneh et al. [7,8] introduce a method for fast revocation of public key certificates and security capabilities in a RSA cryptosystem called mediated RSA (mRSA). The method uses an online semi-trusted mediator (SEM) which has a share of each users secret key, while the user has the remaining share of the secret key. To decrypt or sign a message, a user must first contact SEM and receive a message-specific token. Without the token, the user cannot decrypt or sign a message. Instantaneous user revocation is obtained by instructing SEM to stop issuing tokens for future decrypt/sign requests. Thus, in mediated cryptography the Trusted Authority (TA) responsible to generate a user key pair, does not deliver the full decryption key to users, but it delivers only a share of it. This method achieves faster revocation of user's security capabilities compared to previous certification techniques such as Certificate Revocation List (CRL) and Online Certificate Status Protocol (OCSP). Libert and Quisquater [9] show that the architecture for revoking security capabilities can be applied to several existing public key encryption schemes including the Boneh-Franklin scheme, and several signature schemes including the GDH scheme. Nali et al. [10] present a mediated hierarchical identity-based encryption and signature scheme. The hierarchical nature of the schemes and the instant revocation capability offered by the SEM architecture allows to enforce access control cryptographically in hierarchically structured communities of users whose access privileges change dynamically. Nali et al. [11] also show how to extend the Libert and Quisquater mediated identity-based cryptographic scheme to allow the enforcement of role-based access control (RBAC).

Revocation. Credential revocation is a critical issue for access control systems. For ABE systems, Pirretti et al. [5] propose to use user attributes for a limited time period. After a specific time period the attribute would become invalid. However, in such systems an attribute cannot be revoked before the expiration date. This approach also requires the list of keys that correspond to attributes to be updated regularly, which would also require the users secret keys to be updated regularly. Boldyreva et al. [12] proposes a revocable IBE scheme. The proposed idea for the revocation is based on binary tree data structure, proposed previously in the PKI setting [13,14]. Boldyreva approach is an improvement to Boneh and Franklin [15] idea, however, when the number of revoked users increases, then the advantage of the proposed scheme is lost over that proposed by Boneh and Franklin, especially when the number of revoked users r becomes close to the total number of users n in the system. Even if r is less than n , still the key update complexity is bounded by $O(r \log(\frac{n}{r}))$ while in a realistic solution, the key update complexity should depend on the number of revoked

users. Ostrovsky et al.[16] proposes a Key-Policy Attribute-Based Encryption scheme, where the user secret key may be associated with a non-monotonic access policy. The non-monotonic access policy can be represented by a boolean formula such as AND, OR, NOT, and Out Of (threshold) operations. The main drawback of the scheme is that the size of attributes in the ciphertext is fixed, which restricts the expressivity of the scheme. We note that the concept of revoking an attribute is similar to the concept of revoking an identity. Hence, one can revoke an identity of the user instead of revoking an attribute in an access structure. In this paper we propose a mediated CP-ABE scheme which is not limited to the fixed size of attributes which can be revoked.

The rest of this paper is organized as follows. Section 2 provides background information. In Section 3 we give a formal definition of the mCP-ABE scheme. Section 4 describes the construction of mCP-ABE scheme. In Section 5 we apply the mCP-ABE scheme and describe a general architecture for secure management of Personal Health Records (PHRs). The last section concludes the paper.

2 Background

In this section, we briefly review the basics of bilinear pairing and the security proof in the generic group model, and give a formal definition of CP-ABE.

2.1 Bilinear Pairing

Let \mathbb{G}_0 and \mathbb{G}_1 be two multiplicative groups of prime order p , and let g be a generator of \mathbb{G}_0 . A pairing (or bilinear map) $\hat{e} : \mathbb{G}_0 \times \mathbb{G}_0 \rightarrow \mathbb{G}_1$ satisfies the following properties [15]:

1. Bilinear: for all $u, v \in \mathbb{G}_0$ and $a, b \in \mathbb{Z}_p^*$, we have $\hat{e}(u^a, v^b) = \hat{e}(u, v)^{ab}$.
2. Non-degenerate: $\hat{e}(g, g) \neq 1$.

\mathbb{G}_0 is said to be a bilinear group if the group operation in \mathbb{G}_0 and the bilinear map $\hat{e} : \mathbb{G}_0 \times \mathbb{G}_0 \rightarrow \mathbb{G}_1$ can be computed efficiently. Note that the map is symmetric since $\hat{e}(g^a, g^b) = \hat{e}(g, g)^{ab} = \hat{e}(g^b, g^a)$.

2.2 Security in the Generic Group Model

We prove the security of the scheme based on the generic group model, introduced by Shoup [17]. A proof in the generic group model is based on the fact that the discrete logarithm and the Diffie-Hellman problem are hard to solve as long as the order of the group is a large prime number. The same applies to a group with bilinear pairing where finding the discrete logarithm is a hard problem. In the generic group model group elements are encoded as unique random strings, in such a way that the adversary cannot test any property other than equality.

We prove the security of the mCP-ABE scheme based on the argument that no adversary that acts generally on the groups can break the security of our scheme. This means that if there is an efficient adversary who can discover vulnerabilities

in our scheme, then these vulnerabilities can be used to exploit mathematical properties of groups used in the scheme. In the generic model, the adversary has access to the oracles that compute group operations in $\mathbb{G}_0, \mathbb{G}_1$, and to the oracle that performs non-degenerate pairing \hat{e} , while the adversary can test the equality by itself.

2.3 Formal Definition of CP-ABE

The building block of our construction is a CP-ABE scheme. In CP-ABE a message is encrypted under an access policy τ over the set of possible attributes, and a user secret key sk_ω is associated with an attribute set ω . A secret key sk_ω can decrypt the message encrypted under the access policy τ , if and only if the user attribute set ω satisfies the access policy τ . CP-ABE scheme consists of two entities: TA and users. The four algorithms: Setup, Keygen, Encrypt and Decrypt are defined as follows [2]:

- Setup(k): run by the TA, this algorithm takes as input a security parameter k and outputs the public key pk and a master key mk .
- Keygen(ω, mk): run by the TA, this algorithm takes as input the master key mk and a set of attributes ω . The algorithm outputs a secret key sk_ω associated with ω .
- Encrypt(m, τ, pk): run by the encryptor, this algorithm takes as input a public key pk , a message m , and an access policy represented by an access tree τ . The algorithm returns the ciphertext c_τ such that only users who have the secret key shares associated with attributes that satisfy the access tree τ will be able to decrypt the message.
- Decrypt(c_τ, sk_ω): run by the decryptor, this algorithm takes as input a ciphertext c_τ , a secret key sk_ω associated with ω , and it outputs a message m , or an error symbol \perp when the attribute set ω does not satisfy the access tree τ .

3 Mediated Ciphertext-Policy Attribute-Based Encryption (mCP-ABE)

In this section, first, we give a formal definition of our proposed scheme, and later we give the security model in which our scheme is proven to be secure.

3.1 Formal Definition of mCP-ABE

The mCP-ABE scheme consists of three entities: TA, a mediator and users. The TA uses the master key to generate a user secret key, which is then divided into two shares such that the first share of the user secret key is sent to the mediator and the second share of the user secret key is sent to a user. The mediator has to stay online all the time, while the TA can be put off-line once it has generated secret keys for all users. The mCP-ABE scheme consists of five algorithms: Setup, Keygen, Encrypt, m-Decrypt, and Decrypt (Setup and Encrypt algorithms are same as in CP-ABE scheme):

- $\text{Keygen}(mk, \omega, I_u)$: run by the TA, this algorithm takes as input the master key mk , the user attribute set ω , and the user identifier I_u . The algorithm outputs two secret key shares associated with ω and I_u : $sk_{\omega I_u,1}$ and $sk_{\omega I_u,2}$. The first share of the secret key $sk_{\omega I_u,1}$ is delivered to the mediator, and the second share of the secret key $sk_{\omega I_u,2}$ is delivered to the user. The secret key shares are delivered through a secure channel to the mediator and to the user.
- $\text{m-Decrypt}(c_\tau, I_i, sk_{\omega I_i,1})$: run by the mediator, this algorithm takes as input a ciphertext c_τ , the identifier I_i and the secret key $sk_{\omega I_i,1}$, and outputs a message \hat{c}_τ , or an error symbol \perp when the non-revoked attributes from the set ω do not satisfy the access tree τ .
- $\text{Decrypt}(\hat{c}_\tau, sk_{\omega I_u,2})$: run by the message receiver, this algorithm takes as input a ciphertext \hat{c}_τ , and a secret key $sk_{\omega I_u,2}$, and outputs a message m , or an error symbol \perp when the non-revoked attributes from the set ω does not satisfy the access tree τ .

In practice, there might be multiple entities acting as mediators, and a global entity acting as TA. For example, a healthcare organization may choose Proxy₁ as its mediator and a government organization may choose Proxy₂ as its mediator, where each mediator has the first share of the secret key for registered users in the hospital organization, respectively in the government organization. Vanrenen et al. [18] propose the use of peer-to-peer networking (P2P) which would allow users to require a decryption token from every mediator, such as the mediator either tries to compute a decryption token by itself, or forwards the request to its neighbors.

3.2 Security Model

In our scheme the TA is a fully trusted entity which stores securely the master key. We skip discussions about the key escrow problem, since different existing threshold schemes [19,20] can be applied to solve this problem. A mediator is a semi-trusted entity, namely, it should issue decryption tokens to users, but it is not trusted in the sense that it should not obtain information about the plaintext. We define semantic security of mCP-ABE scheme following the security model of Libert and Quisquater [9]. For an encryption scheme to be semantically secure the adversary must not learn anything about the plaintext when the ciphertext and the public key used to create the ciphertext are given. In the security game, the challenger simulates the game and answers adversary \mathcal{A} queries as follows:

1. **Setup.** The challenger runs the **Setup** algorithm to generate (pk, mk) and gives the public key pk to the adversary \mathcal{A} .
2. **Phase1.** \mathcal{A} performs a polynomially bounded number of queries:
 - $\text{Keygen}^1(\omega, I_u)$. \mathcal{A} asks for a secret key for the attribute set ω and identifier I_u , and receives the mediator share of the secret key $sk_{\omega I_u,1}$.
 - $\text{Keygen}^2(\omega, I_u)$. \mathcal{A} asks for a secret key for the attribute set ω and identifier I_u , and receives the user share of the secret key $sk_{\omega I_u,2}$.

3. **Challenge.** \mathcal{A} sends to the challenger two messages m_0, m_1 , and the challenge access policy τ^* , such that none of the full secret keys $sk_{\omega_{I_u}}$ (both $sk_{\omega_{I_u},1}$ and $sk_{\omega_{I_u},2}$) generated from the interaction with Keygen^1 and Keygen^2 oracles satisfies τ^* . The challenger picks a random bit $b \in (0, 1)$ and returns $c_{\tau^*} = \text{Encrypt}(m_b, \tau^*, pk)$.
4. **Phase2.** \mathcal{A} can continue querying with the restriction that none of the full secret keys $sk_{\omega_{I_u}}$ generated from the interaction with Keygen^1 and Keygen^2 oracles satisfies τ^* .
5. **Guess.** \mathcal{A} outputs a guess $b' \in (0, 1)$.

Definition 1. *The mCP-ABE scheme is said to be semantically secure if any polynomial-time adversary has only a negligible advantage in the security game, where the advantage is defined to be $|\Pr[b' = b] - \frac{1}{2}|$.*

Note that the security game formally captures the following security requirements:

- Resistance against secret key collusion, where different users cannot combine their attribute sets to extend their decryption power. For example, suppose there is a message encrypted under the access tree $\tau = (a_1 \wedge a_2 \wedge a_3)$. Suppose Alice has a secret key $sk_{\omega_A I_A}$ associated with an attribute set $\omega_A = (a_1, a_2)$, and Bob has a secret key $sk_{\omega_B I_B}$ associated with an attribute set $\omega_B = (a_3, a_4)$. Neither Alice's secret key, nor Bob's secret key satisfies the access tree τ . But, if Alice and Bob combine their attribute sets $\omega_A \cup \omega_B = (a_1, a_2, a_3, a_4)$, then the combined attribute sets satisfies the access tree τ . Therefore in the security game we allow the adversary to make secret key queries associated with different attribute sets, say ω_1 and ω_2 , such that neither ω_1 , nor ω_2 alone can satisfy the challenge access policy τ^* , but $\omega_1 \cup \omega_2$ can satisfy τ^* .
- Resistance against malicious cooperation between the mediator and some users to decrypt the ciphertext associated with an access policy, when the users secret key does not satisfy the access policy. For example, even if a user with attribute set $\omega = (a_1, a_2)$ collude with the mediator, the user should not be capable to decrypt a ciphertext encrypted under a challenge access policy $\tau^* = (a_1 \wedge a_2 \wedge a_3)$, since ω does not satisfy τ^* . Therefore in the security game the adversary is allowed to ask the mediator share (first share of the secret key $sk_{\omega_{I_u},1}$) and the user share of a secret key (second share of the secret key $sk_{\omega_{I_u},2}$) for any set of attributes which does not satisfy the challenge access policy τ^* .

4 mCP-ABE Scheme

In this section, we give a description of the access policy associated with the ciphertext, and then we give the construction of the scheme. Finally, we analyze the security of the scheme and describe how to revoke user attributes.

4.1 Access Policy

In mCP-ABE scheme, an access policy is represented by an access tree τ , in which inner nodes are either \wedge (and) or \vee (or) boolean operators, and leaf nodes are attributes. The access tree τ specifies which combination of attributes the decryptor needs to possess in order to decrypt the ciphertext. Figure 1 presents an example of an access tree τ representing an access policy: $(a_1 \wedge a_4) \vee (a_3 \vee a_5)$. To decrypt an encrypted message under the access tree τ , the decryptor must possess a secret key which is associated with the attribute set which satisfies τ . Attributes are interpreted as logic variables, and possessing a secret key associated with an attribute makes the corresponding logical variable *true*. There are several different sets of attributes that can satisfy the access tree τ presented in Figure 1, such as the attribute set (a_1, a_4) , the attribute (a_3) , or the attribute (a_5) . In our scheme, we assume that attributes are ordered in the access tree e.g. $\text{index}(a_1)=1$, $\text{index}(a_4)=2$, $\text{index}(a_3)=3$ and $\text{index}(a_5)=4$.




Fig. 1. Access tree $\tau = (a_1 \wedge a_4) \vee (a_3 \vee a_5)$

4.2 Main Construction

1. **Setup**(k) : On input of the security parameter k , the algorithm generates a group \mathbb{G}_0 of prime order p with a generator g and a bilinear map $\hat{e} : \mathbb{G}_0 \times \mathbb{G}_0 \rightarrow \mathbb{G}_1$. The algorithm generates the system attribute set $\Omega = (a_1, a_2, \dots, a_n)$, for some integer n , and for each $a_j \in \Omega$ chooses a random elements $t_j \in \mathbb{Z}_p$. Let $y = \hat{e}(g, g)^\alpha$, where α is chosen at random from \mathbb{Z}_p , and $\{T_j = g^{t_j}\}_{j=1}^n$. The public key is published as:

$$pk = (g, y, \{T_j\}_{j=1}^n)$$

The master secret key consists of the following components:

$$mk = (\alpha, \{t_j\}_{j=1}^n)$$

2. **Keygen**(mk, ω, I_u) : To generate a secret key for the user with an attribute set ω and an identifier I_u , the **Keygen** algorithm performs as follows:
 - (a) Compute the base component of the secret key: $d_0 = g^{\alpha - u_{id}}$ where $u_{id} \in_R \mathbb{Z}_p$ (for each user with an identifier I_u a unique random value u_{id} is generated).

- (b) Compute the attribute component of the secret key. For each attribute $a_j \in \omega$, choose $u_j \in_R \mathbb{Z}_p$ and compute $d_{j,1} = g^{\frac{u_j}{t_j}}$ and $d_{j,2} = g^{\frac{u_{id} - u_j}{t_j}}$.

The secret key of the form: $sk_{\omega I_u,1} = \{d_{j,1}\}_{a_j \in \omega}$ is delivered to the mediator, and the secret key of the form: $sk_{\omega I_u,2} = (d_0, \{d_{j,2}\}_{a_j \in \omega})$ is delivered to the user.

3. **Encrypt**(m, τ, pk) : To encrypt a message $m \in \mathbb{G}_1$ the algorithm proceeds as follows:

- Select a random element $s \in \mathbb{Z}_p$ and compute:

$$\begin{aligned} c_0 &= g^s \\ c_1 &= m \cdot y^s = m \cdot \hat{e}(g, g)^{\alpha s} \end{aligned}$$

- Set the value of the root node of τ to be s , mark all nodes as un-assigned, and mark the root node assigned. Recursively, for each assigned non-leaf node, suppose its value is s , do the following.
 - If the symbol is \wedge and its child nodes are marked un-assigned, let n be the number of child nodes, set the value of each child node, except the last one, to be $s_i \in_R \mathbb{Z}_p$, and the value of the last node to be $s_n = s - \sum_{i=1}^{n-1} s_i \bmod p$ (i represents the index of an attribute in the access tree). Mark this node assigned.
 - If the symbol is \vee , set the values of its child nodes to be s . Mark this node assigned.
- For each leaf attribute $a_{j,i} \in \tau$, compute $c_{j,i} = T_j^{s_i}$.

Return the ciphertext $c_\tau = (\tau, c_0, c_1, \{c_{j,i}\}_{a_{j,i} \in \tau})$.

4. **m-Decrypt**($c_\tau, sk_{\omega I_i,1}, I_i$) : when receiving the ciphertext c_τ , the recipient I_i firstly chooses the smallest set $\omega' \subseteq \omega$ that satisfies τ and forwards to the mediator (c_τ, ω', I_i) . The mediator checks the Attribute Revocation List (ARL) if any $a_j \in \omega'$ is revoked either from system attribute set Ω or from the user attribute set ω .

- (a) If an attribute is revoked, the mediator returns an error symbol \perp and does not perform further computations.
- (b) If no attribute is revoked, the mediator computes \hat{c}_τ as follows:

$$\begin{aligned} \hat{c}_\tau &= \prod_{a_j \in \omega'} \hat{e}(T_j^{s_i}, g^{\frac{u_j}{t_j}}) \\ &= \hat{e}(g, g)^{\sum_{a_j \in \omega'} u_j s_i} \end{aligned}$$

Sends \hat{c}_τ to the recipient.

5. **Decrypt**($\hat{c}_\tau, sk_{\omega I_i,2}$) : To decrypt the ciphertext the recipient proceeds as follows:

(a) compute:

$$\begin{aligned}
 c''_{\tau} &= \prod_{a_j \in \omega'} \hat{e}(T_j^{s_i}, g^{\frac{u_{id}-u_j}{t_j}}) \\
 &= \prod_{a_j \in \omega'} \hat{e}(g^{t_j s_i}, g^{\frac{u_{id}-u_j}{t_j}}) \\
 &= \hat{e}(g, g)^{\sum_{a_j \in \omega'} (u_{id}-u_j) s_i}
 \end{aligned}$$

(b) compute:

$$\begin{aligned}
 \hat{e}(c_0, d_0) \cdot \hat{c}_{\tau} \cdot c''_{\tau} &= \hat{e}(g^s, g^{\alpha-u_{id}}) \cdot \hat{e}(g, g)^{\sum_{a_j \in \omega'} u_j s_i} \cdot \hat{e}(g, g)^{\sum_{a_j \in \omega'} (u_{id}-u_j) s_i} \\
 &= \hat{e}(g^s, g^{\alpha-u_{id}}) \cdot \hat{e}(g, g)^{u_{id} s} \\
 &= \hat{e}(g^s, g^{\alpha})
 \end{aligned}$$

(c) return m , where

$$\begin{aligned}
 m &= \frac{c_1}{\hat{e}(g^s, g^{\alpha})} \\
 &= \frac{m \cdot \hat{e}(g, g)^{\alpha s}}{\hat{e}(g^s, g^{\alpha})}
 \end{aligned}$$

Efficiency. Our scheme is similar to the work of Cheung and Newport [3] on ciphertext-policy attribute-based encryption, however we make major changes in the Key Generation phase, Encryption phase and Decryption phase in order to improve the expressivity of the scheme (the scheme in [3] supports only access policies with logical conjunction), and we improve the efficiency of the scheme (in [3] the size of the ciphertext and secret key increases linearly with the total number of attributes in the system). In our proposed scheme, the size of the shares of the secret key $sk_{\omega_{I_u,1}}$ and $sk_{\omega_{I_u,2}}$ depend on the number of attributes the user has and consists of $|\omega| + 1$ group elements in \mathbb{G}_0 ($|\omega|$ is the cardinality of a set ω). The size of the ciphertext c_{τ} depends on the size of the access policy τ and has $|\tau| + 1$ group elements in \mathbb{G}_0 , and one group element in \mathbb{G}_1 . In the m-Decrypt phase, the mediator has to compute ω' pairing operations, where $\omega' \subseteq \omega$ is the attribute set which satisfies the access policy τ . In the decryption phase, to reveal the message, the user has to compute $\omega' + 1$ pairing operations.

For the sake of simplicity, we mentioned only access policies which consist of \wedge (and) and \vee (or) nodes. Note that our scheme, in addition to \wedge (and) and \vee (or) nodes, can support threshold nodes or Out Of nodes. For example, the encryptor may specify the access policy 2 Out Of (a_1, a_2, a_3) , which implies that the user must have at least two attributes in order to satisfy the access policy and be able to decrypt. If the access policy contains threshold nodes, then the attribute shares s_i can be generated using threshold secret sharing techniques e.g. using Shamir's secret sharing technique. This would require, to use Lagrange basis polynomials in the decryption phase in order to reconstruct the value s .

4.3 Security Analysis

We give a brief discussion about the security of the proposed scheme. A full formal security proof using the generic group model is given in the full version of this paper [21]. To decrypt a ciphertext without satisfying the access policy, the adversary has to construct $\hat{e}(g^s, g^\alpha)$, and then divide c_1 with $\hat{e}(g^s, g^\alpha)$ to obtain m . To obtain $\hat{e}(g^s, g^\alpha)$, the adversary must first obtain $\hat{e}(g, g)^{su_{id}}$, which can be calculated by pairing the components of the secret key $g^{\frac{u_{id}-u_j}{t_j}}$ with the components of the ciphertext $g^{t_j s_i}$, and then multiply the result with the decryption token $\hat{e}(g, g)^{\sum_{a_j \in \omega'} u_j s_i}$ received from the mediator. However, $\hat{e}(g, g)^{su_{id}}$ can be computed only if the adversary has enough attributes which satisfy the access policy, otherwise this would not be possible. Also note that, if a user acting as an adversary is revoked, then the user will not get the decryption token $\hat{e}(g, g)^{\sum_{a_j \in \omega'} u_j s_i}$ from the mediator, and as a result of this the user cannot reconstruct $\hat{e}(g, g)^{su_{id}}$ even if the user has a secret key with attributes which satisfy the access policy. If we assume that the adversary is able to compromise the mediator, then the adversary will be able to learn the mediator share of user secret key $sk_{\omega I_u, 1}$, and be able to compute the decryption token $\hat{e}(g, g)^{\sum_{a_j \in \omega'} u_j s_i}$. However, the decryption token will not help the adversary to decrypt the ciphertext since the adversary does not know the second share of the user secret key $sk_{\omega I_u, 2}$.

The very important security of mCP-ABE scheme is a collusion resistance of user secret keys - it should not be possible for different users to combine their secret keys in order to extend their decryption power. Therefore, to prevent collusion, the Keygen algorithm of our scheme generates a random value u_{id} for each user, which is embedded in each component of the user secret key. Users cannot combine components of the secret key since different users have different random value in their secret keys.

4.4 Attribute Revocation

As already mentioned in section 1, there can be many reasons why an attribute can be revoked. We assume that the mediator maintains an Attribute Revocation List (ARL) which simply has information about attributes revoked from the system attribute set Ω , and attributes revoked from user attribute set ω . The basic idea is that, when an attribute a_j is revoked from the system attribute set Ω , the TA removes a_j from the system attribute set, and notifies the mediator to stop performing decryption tasks for all users whose attribute secret key involves a_j . When an attribute is revoked from a specific user I_u , TA notifies the mediator to stop helping the user I_u to perform decryption tasks for the attribute a_j . Therefore, the attribute revocation is achieved immediately after the revocation decision is made. We assume that there is a policy of revocation authorization maintained by the mediator that describes who is responsible to revoke system or user attributes. At least, the TA should be able to revoke the system and user attributes, and the owner of the attribute should be able to revoke its attribute because the owner may be the first to notice the compromise of her secret key.

5 Applying mCP-ABE in Practice

In this section we describe an application of mCP-ABE. We propose to use mCP-ABE to securely manage Personal Health Records (PHRs). This application demonstrates the practicality and usefulness of our scheme.

5.1 Using mCP-ABE to Securely Manage PHRs

Issues around the confidentiality of health records are considered as one of the primary reason for the lack of the deployment of open interoperable health record systems. Health data is sensitive: inappropriate disclosure of a record can change a patient's life, and there may be no way to repair such harm financially or technically. Although, access to health data in the professional medical domain is tightly controlled by existing legislations, such as the U.S. Health Insurance Portability and Accountability Act (HIPAA) [22], private web PHR systems stay outside the scope of this legislation. Therefore, a number of patients might hesitate to upload their sensitive health records to web PHR systems such as the Microsoft Health Vault, Google Health or WebMD. The scheme presented in this paper allows patients to store their sensitive health records on web PHR systems in an encrypted form, while still giving them control to share their data with healthcare providers and/or with their family members. The reason is because the data is encrypted according to an access policy, and the policy moves with the encrypted data. Thus, even if the server which stores health records gets compromised, the confidentiality of the data is preserved since the data is encrypted, and the attacker cannot decrypt the encrypted data without having a secret key. Figure 2 illustrates a general architecture of a PHR system that uses mCP-ABE. The architecture consists of a publishing server, a data repository that includes a security mediator (Proxy), TA and several data users. The publishing server can be implemented on a home PC of the data source (a patient) or as a trusted service. Its role is to protect and publish health records. The data repository stores encrypted health records, while Proxy is used in the data consumption phase for revocation. The TA is used to set up the keys. Note that the TA and the publishing server do not have to be always online (the TA is needed only in the set-up phase while the publishing server can upload the protected data in an ad-hoc way). There are three basic processes in the management of PHRs:

1. Setup: The steps of this phase are depicted with number 1 in Figure 2. In this phase, the TA distributes the keys to the patients, users and Proxy.
2. Data protection (upload of data to the PHR): When a patient wants to upload protected data to the repository, she contacts the TA to check which attributes are allowed to be used as a policy. Then she creates her access control policy and encrypts the data with the keys corresponding to that policy. Then the data is uploaded to the repository. If she wants to change the policy she can re-encrypt the data and update the repository. All this can be done by a publishing server on behalf of the patient who specifies the access control policy.

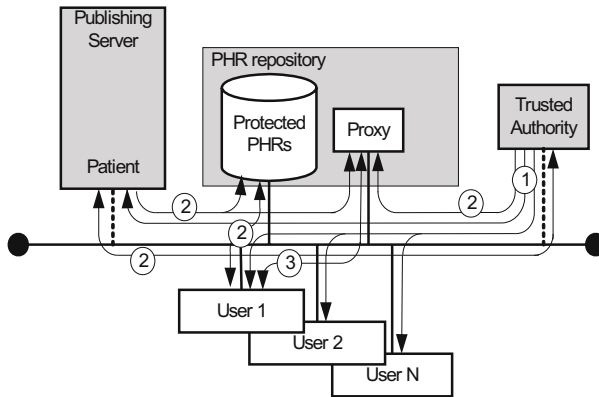


Fig. 2. Secure Management of PHRs

3. Data consumption (doctor's request-response) and revocation: When a user wants to use patient data he contacts the PHR repository and downloads encrypted data. The user makes a request to Proxy for a decryption token. The request contains the encrypted data and a set of user attributes which satisfy the access policy associated with the encrypted data. Proxy checks if any attribute from the user request is not revoked, and, if so, Proxy generates the decryption token and send it to the user. After receiving the decryption token the user decrypts the patient data using the keys corresponding to the appropriate attributes which satisfy the access tree. The steps of this phase are depicted with number 3 in Figure 2.

An addition advantage of an online semi-trusted mediator (Proxy) is that the mCP-ABE scheme can be used to enforce context attributes such as: system date and time or the location from where the request comes from. This is useful for healthcare applications which require context-aware access control where access to patients data depends not only on user roles, but also on the context information. Suppose there is an access policy $\tau = (Location = Hospital \wedge (A \wedge B))$ which says that a doctor (we assume that a doctor is identified with attributes A and B) can view patients health records only if doctor's request comes from inside the hospital. Outside the hospital, no user should be able to decrypt the ciphertext encrypted under the access tree τ , even if the user may own a secret key associated with the attributes (A, B) . Using mCP-ABE scheme, the enforcement of τ , which contain context attributes, is done as follows:

- a patient encrypts her health record according to the access policy $(A \wedge B)$, and then uploads his data to a PHR repository. Hence, part of the access policy is enforced in the **Encryption** phase by the patient.
- a doctor download encrypted data and request from Proxy a decryption token.
- Proxy checks the context attribute inside τ and issues decryption token only if the request comes from inside the hospital (e.g. only if the request comes

from a specific IP address), therefore the context attribute is enforced by Proxy in the m-Decrypt phase.

Note that the involvement of an online semi-trusted mediator (Proxy) plays a crucial role in the enforcement of context attributes, as it is very hard or rather impossible to enforce these attributes without the involvement of an online component. The mCP-ABE scheme can also support the off-line use of data. Then the architecture is slightly changed in a way that Proxy is distributed to the users or their domains within which the data will be used. As a consequence there will be a number of proxies which will be coordinated by the central Proxy. The above defined process will not fundamentally change, except that the central Proxy will update the local ones and that in the data consumption phase, the user will contact only the local Proxy.

6 Conclusion

We propose a mediated Ciphertext-Policy Attribute-Based Encryption scheme that supports revocation of user attributes. If an attribute is revoked, the user cannot use it in the decryption phase. The scheme allows the encryptor to encrypt a message according to an access policy over a set of attributes, and only users who satisfy the access policy and whose attributes are not revoked can decrypt the ciphertext. Furthermore, we demonstrate how to use the proposed scheme to solve very important problems in managing Personal Health Records (PHRs). A possible extension to this work would be to provide a scheme which would have a security proof under standard complexity assumptions.

References

1. Sahai, A., Waters, B.: Fuzzy identity-based encryption. In: Cramer, R. (ed.) EUROCRYPT 2005. LNCS, vol. 3494, pp. 457–473. Springer, Heidelberg (2005)
2. Bethencourt, J., Sahai, A., Waters, B.: Ciphertext-policy attribute-based encryption. In: Proceedings of the 2007 IEEE Symposium on Security and Privacy, pp. 321–334. IEEE Computer Society Press, Washington (2007)
3. Cheung, L., Newport, C.: Provably secure ciphertext policy ABE, pp. 456–465. ACM, New York (2007)
4. Ibraimi, L., Tang, Q., Hartel, P., Jonker, W.: Efficient and provable secure ciphertext-policy attribute-based encryption schemes. In: Bao, F., Li, H., Wang, G. (eds.) ISPEC 2009. LNCS, vol. 5451, pp. 1–12. Springer, Heidelberg (2009)
5. Piretti, M., Traynor, P., McDaniel, P., Waters, B.: Secure attribute-based systems. In: Proceedings of the 13th ACM Conference on Computer and Communications Security, pp. 99–112 (2006)
6. Goyal, V., Pandey, O., Sahai, A., Waters, B.: Attribute-based encryption for fine-grained access control of encrypted data. In: Proceedings of the 13th ACM Conference on Computer and Communications Security, pp. 89–98. ACM Press, New York (2006)

7. Boneh, D., Ding, X., Tsudik, G., Wong, C.M.: A method for fast revocation of public key certificates and security capabilities. In: SSYM 2001: Proceedings of the 10th conference on USENIX Security Symposium, Berkeley, CA, USA, p. 22. USENIX Association (2001)
8. Boneh, D., Ding, X., Tsudik, G.: Fine-grained control of security capabilities. *ACM Transactions on Internet Technology* 4, 60–82 (2004)
9. Libert, B., Quisquater, J.J.: Efficient revocation and threshold pairing based cryptosystems. In: Proceedings of the twenty-second annual symposium on Principles of distributed computing, pp. 163–171. ACM, New York (2003)
10. Nali, D., Miri, A., Adams, C.: Efficient revocation of dynamic security privileges in hierarchically structured communities. In: Proceedings of the 2nd Annual Conference on Privacy, Security and Trust (PST 2004), Fredericton, New Brunswick, Canada, pp. 219–223 (2004)
11. Nali, D., Adams, C., Miri, A.: Using mediated identity-based cryptography to support role-based access control. In: Zhang, K., Zheng, Y. (eds.) ISC 2004. LNCS, vol. 3225, pp. 245–256. Springer, Heidelberg (2004)
12. Boldyreva, A., Goyal, V., Kumar, V.: Identity-based encryption with efficient revocation. In: Proceedings of the 15th ACM conference on Computer and communications security, pp. 417–426. ACM, New York (2008)
13. Aiello, W., Lodha, S., Ostrovsky, R.: Fast digital identity revocation (extended abstract). In: Krawczyk, H. (ed.) CRYPTO 1998. LNCS, vol. 1462, pp. 137–152. Springer, Heidelberg (1998)
14. Naor, M., Nissim, K.: Certificate revocation and certificate update. In: USENIX Security Symposium, pp. 561–570. IEEE, Los Alamitos (1998)
15. Boneh, D., Franklin, M.: Identity-based encryption from the weil pairing. In: Kilian, J. (ed.) CRYPTO 2001. LNCS, vol. 2139, p. 213. Springer, Heidelberg (2001)
16. Ostrovsky, A., Sahai, A., Waters, B.: Attribute based encryption with non-monotonic access structures. In: ACM conference on Computer and Communications Security, pp. 195–203. ACM, New York (2007)
17. 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)
18. Vanrenen, G., Smith, S.: Distributing security-mediated PKI. In: Katsikas, S.K., Gritzalis, S., López, J. (eds.) EuroPKI 2004. LNCS, vol. 3093, pp. 218–231. Springer, Heidelberg (2004)
19. Shamir, A.: How to share a secret. *Communications of the ACM* 22(11), 612–613 (1979)
20. Desmedt, Y., Frankel, Y.: Threshold cryptosystems. In: Brassard, G. (ed.) CRYPTO 1989. LNCS, vol. 435, pp. 307–315. Springer, Heidelberg (1990)
21. Ibraimi, L., Petkovic, M., Nikova, S.I., Hartel, P.H., Jonker, W.: Mediated ciphertext-policy attribute-based encryption and its application (extended version). Technical Report TR-CTIT-09-12, Enschede (April 2009)
22. The US Department of Health and Human Services. Summary of the HIPAA privacy rule (2003)