

CERIAS Tech Report 2012-05
Attribute Based Group Key Management
by Mohamed Nabeel, Elisa Bertino
Center for Education and Research
Information Assurance and Security
Purdue University, West Lafayette, IN 47907-2086

Attribute Based Group Key Management

Mohamed Nabeel, *Student Member, IEEE*, and Elisa Bertino, *Fellow, IEEE*

Abstract—Attribute based systems enable fine-grained access control among a group of users each identified by a set of attributes. Secure collaborative applications need such flexible attribute based systems for managing and distributing group keys. However, current group key management schemes are not well designed to manage group keys based on the attributes of the group members. In this paper, we propose novel key management schemes that allow users whose attributes satisfy a certain access control policy to derive the group key. Our schemes efficiently support rekeying operations when the group changes due to joins or leaves of group members. During a rekey operation, the private information issued to existing members remains unaffected and only the public information is updated to change the group key. Our schemes are expressive; they are able to support any monotonic access control policy over a set of attributes. Our schemes are resistant to collusion attacks; group members are unable to pool their attributes and derive the group key which they cannot derive individually. Experimental results show that our underlying constructs are efficient and practical.

Index Terms—Broadcast group key management, attribute based policies, secret sharing



1 INTRODUCTION

New application domains have pushed novel paradigms and tools for supporting collaboration among (possibly very dynamic) user groups (see for example the notion of group-centric information sharing [1]). An important requirement in collaborative applications is to support operations for user group memberships, like join and leave, based on identity attributes (attributes, for short) of users; we refer to this requirement as *attribute-based group dynamics*. As today enterprises and applications are adopting identity management solutions, it is crucial that these solutions be leveraged on for managing groups. Typically, a user would be automatically assigned (de-assigned) a group membership based on whether his/her attributes satisfy (cease to satisfy) certain group membership conditions. Another critical requirement is to provide mechanisms for group key management (GKM), as very often the goal of a group is to share data. Thus data must be encrypted with keys made available only to the members of the group. The management of these keys, which includes selecting, distributing, storing and updating keys, should directly and effectively support the attribute-based group dynamics and thus requires an *attribute-based group key management* (AB-GKM) scheme, by which group keys are assigned (or de-assigned) to users in a group based on their identity attributes. This scheme recalls the notion of attribute-based encryption (ABE) [2], [3], [4]; however, as we discuss later on, ABE has several shortcomings when applied to GKM. Therefore, a different approach is needed.

A challenging well known problem in GKM is how to efficiently handle group dynamics, i.e., a new user

joining or an existing group member leaving. When the group changes, a new group key must be shared with the existing members, so that a new group member cannot access the data transmitted before she joined (forward secrecy) and a user who left the group cannot access the data transmitted after she left (backward secrecy). The process of issuing a new key is called *rekeying* or *update*. Another challenging problem is to defend against collusion attacks by which a set of colluding fraudulent users are able to obtain group keys which they are not allowed to obtain individually.

In a traditional GKM scheme, when the group changes, the private information given to all or some existing group members must be changed which requires establishing private communication channels. Establishing such channels is a major shortcoming especially for highly dynamic groups. Recently proposed broadcast GKM (BGKM) schemes [5], [6] have addressed such shortcoming. BGKM schemes allow one to perform rekeying operations by only updating some public information without affecting private information existing group members possess. However, BGKM schemes do not support group membership policies over a set of attributes. In their basic form, they can only support 1-out-of- n threshold policies by which a group member possessing 1 attribute out of the possible n attributes is able to derive the group key. In this paper we develop novel expressive AB-GKM schemes which allow one to express any threshold or monotonic¹ conditions over a set of identity attributes.

A possible approach to construct an AB-GKM scheme is to utilize attribute-based encryption (ABE) primitives [2], [3], [4]. Such an approach would work as follows. A key generation server issues each group member a private key (a set of secret values) based on

• M. Nabeel and E. Bertino are with the Department of Computer Science, Purdue University, West Lafayette, IN, 47907.
E-mail: nabeel@cs.purdue.edu

1. Monotone formulas are Boolean formulas that contain only conjunction and disjunction connectives, but no negation.

the attributes and the group membership policies. The group key, typically a symmetric key, is then encrypted under a set of attributes using the ABE encryption algorithm and broadcast to all the group members. The group members whose attributes satisfy the group membership policy can obtain the group key by using the ABE decryption primitive. One can use such an approach to implement an expressive collusion-resistant AB-GKM scheme. However, such an approach suffers from some major drawbacks. Whenever the group dynamic changes, the rekeying operation requires to update the private keys given to existing members in order to provide backward/forward secrecy. This in turn requires establishing private communication channels with each group member which is not desirable in a large group setting. Further, in applications involving stateless members where it is not possible to update the initially given private keys and the only way to revoke a member is to exclude it from the public information, an ABE based approach does not work. Another limitation is that whenever the group membership policy changes, new private keys must be re-issued to members of the group. Our constructions address these shortcomings.

Our AB-GKM schemes are able to support a large variety of conditions over a set of attributes. When the group changes, the rekeying operations do not affect the private information of existing group members and thus our schemes eliminate the need of establishing private communication channels. Our schemes provide the same advantage when the group membership conditions change. Furthermore, the group key derivation is very efficient as it only requires a simple vector inner product and/or polynomial interpolation. Additionally, our schemes are resistant to collusion attacks. Multiple group members are unable to combine their private information in a useful way to derive a group key which they cannot derive individually.

Our AB-GKM constructions are based on an optimized version of the ACV-BGKM (Access Control Vector BGKM) scheme [6], a provably secure BGKM scheme, and Shamir’s threshold scheme [7]. In this paper, we construct three AB-GKM schemes each of which is more suitable over others under different scenarios. The first construction, inline AB-GKM, is based on the ACV-BGKM scheme. Inline AB-GKM supports arbitrary monotonic policies over a set of attributes. In other words, a user whose attributes satisfy the group policies is able to derive the symmetric group key. However, inline AB-GKM does not efficiently support d -out-of- m ($d \leq m$) attribute threshold policies over m attributes. The second construction, threshold AB-GKM, addresses this requirement. The third construction, access tree AB-GKM, is an extension of threshold AB-GKM and is the most expressive scheme. It efficiently supports arbitrary policies. The second and third schemes are constructed by using a modified version of ACV-BGKM, also proposed in this paper.

The remainder of the paper is organized as follows:

TABLE 1
Acronyms

Acronym	Description
GKM	Group Key Management
BGKM	Broadcast GKM
ABE	Attribute Based Encryption
CP-ABE	Ciphertext Policy ABE
KP-ABE	Key Policy ABE
ACV	Access Control Vector
KEV	Key Extraction Vector
ABAC	Attribute Based Access Control
AB-GKM	Attribute Based GKM
PI	Public Information tuple
UA	User-Attribute matrix

Section 2 describes related work. Section 3 summarizes the ACV-BGKM scheme [6]. Sections 4, 5, 6 present the construction of the inline AB-GKM, threshold AB-GKM, and access tree AB-GKM schemes, respectively, and analyze their security and performance. Section 7 shows an example application using the access tree AB-GKM scheme. Section 8 provides experimental results of our underlying optimized ACV-BGKM scheme used in all three schemes against the CP-ABE (ciphertext policy attribute based encryption) scheme. Section 9 concludes the paper. Table 1 lists, for the convenience of the reader, the acronyms used in the paper.

2 RELATED WORK

Group Key Management (GKM): Early approaches to GKM rely on a key server to share a secret with users to distribute decryption keys [8], [9]. Such approaches do not efficiently handle join and leave operations, as in order to achieve forward and backward security, they require sending $O(n)$ private rekey information, where n is the number of users. Hierarchical key management schemes [10], [11] were introduced to reduce this overhead. However, they only reduce the size of the rekey information to $O(\log n)$, and furthermore each user needs to manage at worst $O(\log n)$ hierarchically organized redundant keys.

Broadcast GKM (BGKM) schemes perform the rekey operation with only one broadcast without affecting the secret information issued to existing users. Approaches have also been proposed to make the rekey operation a one-off process [12], [5]. However, these schemes are not formally proven to be secure. Recently Shang et. al. introduced the first provably secure BGKM scheme called ACV-BGKM [6]. Existing BGKM schemes require sending $O(n)$ public information when rekeying. We improve the complexity by utilizing subset-cover techniques [13], [14]. Such improved BGKM schemes efficiently handle group dynamics and lay the foundation for AB-GKM. However such schemes cannot directly handle expressive conditions against attributes.

Attribute-Based Encryption (ABE) and GKM: In an ABE system [2], the plaintext is encrypted with a set of attributes. The key generation server, which possesses

the master key, issues different private keys to users after authenticating the attributes they possess. Thus, these private keys are associated with the set of attributes each user possesses. In its basic form, a user can decrypt a ciphertext if and only if there is a match between the attributes of the ciphertext and the user's key. The initial ABE system was limited to only threshold policies by which there should be at least k out of n attributes common between the attributes used to encrypt the plaintext and the attributes users possess. Since the definition of the initial threshold scheme, a few variants have been introduced to provide more expressive ABE systems. Goyal et al. [3] introduced the idea of key-policy ABE (KP-ABE) systems and Bethencourt et al. [4] introduced the idea of CP-ABE systems. Even though these constructs are expressive and provably secure, they are unable to efficiently support group management, and especially to provide backward secrecy when a user leaves the group (i.e. attribute revocation) and to provide forward secrecy when a new user joins the group. These schemes require sending $O(n)$ private rekey messages in order to handle group management operations. The proposers of some of these schemes have suggested using an expiration attribute along with other attributes for attribute revocation. However, such a solution is not suitable for highly dynamic groups where joins and leaves are frequent. Traynor et al. [15] propose to improve the performance of ABE by grouping users and assigning a unique group attribute to each group. However, their approach only considers one attribute per user and does not support membership policy based group key management.

Despite the limitations of ABE schemes with respect to revocation, flat table based GKM schemes² based on ABE have been proposed [16], [17]. These schemes further suffer from the inherent limited expressibility and scalability of flat table based GKM [18], [19].

GKM Schemes for Selective Dissemination Systems: Selective dissemination or broadcast encryption systems allow one to encrypt a message once and broadcast it to all the users in a group, but only a subset of users who have the correct key can decrypt the message. The database and security communities have carried out extensive research concerning techniques for the selective dissemination of documents based on access control policies with their own GKM schemes [20], [21], [22], [23], [24]. In such approaches, users are able to decrypt the subdocuments, that is, portions of documents, for which they have the keys. However, such approaches require all [21] or some [22] keys be distributed in advance during user registration phase. This requirement makes it difficult to assure forward and backward key secrecy when user groups are dynamic with frequent join and leave operations. Further, the rekey operation is not transparent, thus shifting the burden of acquiring new

keys on existing users when others leave or join. Thus the proposed GKM schemes are not efficient. In contrast, our GKM schemes make rekey transparent to users by not distributing actual keys.

3 BACKGROUND

In this section, we provide an overview of the Broadcast Group Key Management (BGKM) scheme in general and a description of a provably secure BGKM scheme called ACV-BGKM (Access Control Vector BGKM) proposed by Shang et al. [6], [25] in order for readers to better understand our constructions. It should be noted that we use ACV-BGKM in Section 4 and a modified version of ACV-BGKM in our constructions in Sections 5, and 6.

BGKM schemes are a special type of GKM scheme where the rekey operation is performed with a single broadcast without using private communication channels. Unlike conventional GKM schemes, BGKM schemes do not give users the private keys. Instead users are given a secret which is combined with public information to obtain the actual private keys. Such schemes have the advantage to require a private communication only once for the initial secret sharing. The subsequent rekeying operations are performed using one broadcast message. Further, in such schemes achieving forward and backward secrecy requires only to change the public information and does not affect the secret shares given to existing users. In general, a BGKM scheme consists of the following five algorithms:

Setup(ℓ): It initializes the BGKM scheme using a security parameter ℓ . It also initializes the set of used secrets \mathbf{S} , the secret space \mathcal{SS} , and the key space \mathcal{KS} .

SecGen(): It picks a random bit string $s \notin \mathbf{S}$ uniformly at random from \mathcal{SS} , adds s to \mathbf{S} and outputs s .

KeyGen(\mathbf{S}): It picks a group key k uniformly at random from \mathcal{KS} and outputs the public information tuple PI computed from the secrets in \mathbf{S} and the group key k .

KeyDer(s, PI): It takes the user's secret s and the public information PI to output the group key. The derived group key is equal to k if and only if $s \in \mathbf{S}$.

Update(\mathbf{S}): Whenever the set \mathbf{S} changes, a new group key k' is generated. Depending on the construction, it either executes the **KeyGen** algorithm again or incrementally updates the output of the last **KeyGen** algorithm.

Using the above abstract algorithms, we now provide an overview of the construction of the ACV-BGKM scheme under a client-server architecture. The ACV-BGKM scheme satisfies the requirements of *minimal trust*, *key indistinguishability*, *key independence*, *forward secrecy*, *backward secrecy* and *collusion resistance* [26]. The ACV-BGKM algorithms are executed by a trusted key server Svr and a group of users $\text{Usr}_i, i = 1, 2, \dots, n$.

Setup(ℓ): Svr initializes the following parameters: an ℓ -bit prime number q , the maximum group size $N (\geq n$ and N is usually set to $n + 1$), a cryptographic hash function $H(\cdot) : \{0, 1\}^* \rightarrow \mathbb{F}_q$, where \mathbb{F}_q is a finite field with q elements, the keyspace $\mathcal{KS} = \mathbb{F}_q$, the secret space $\mathcal{SS} = \{0, 1\}^\ell$ and the set of issued secrets $\mathbf{S} = \emptyset$.

2. A flat table GKM scheme assigns each member a unique n -bit string. The group key is managed through a set of auxiliary keys which are tied to the unique strings given to the group members.

SecGen(): Svr chooses the secret $s_i \in \mathcal{SS}$ uniformly at random for Usr_i such that $s_i \notin \mathbf{S}$, adds s_i to \mathbf{S} and finally outputs s_i .

KeyGen(S): Svr picks a random $k \in \mathcal{KS}$ as the group key. Svr chooses N random bit strings $z_1, z_2, \dots, z_N \in \{0, 1\}^\ell$. Svr creates an $n \times (N + 1)$ \mathbb{F}_q -matrix

$$A = \begin{pmatrix} 1 & a_{1,1} & a_{1,2} & \dots & a_{1,N} \\ 1 & a_{2,1} & a_{2,2} & \dots & a_{2,N} \\ \vdots & \vdots & \vdots & \vdots & \vdots \\ 1 & a_{n,1} & a_{n,2} & \dots & a_{n,N} \end{pmatrix},$$

where

$$a_{i,j} = H(s_i || z_j), 1 \leq i \leq n, 1 \leq j \leq N, s_i \in \mathcal{S}. \quad (1)$$

Svr then solves for a nonzero $(N + 1)$ -dimensional column \mathbb{F}_q -vector Y such that $AY = 0$. Note that such a nonzero Y always exists as the nullspace of matrix A is nontrivial by construction. Here we require that Svr chooses Y from the nullspace of A uniformly at random. Svr constructs an $(N + 1)$ -dimensional \mathbb{F}_q -vector $ACV = k \cdot e_1^T + Y$, where $e_1 = (1, 0, \dots, 0)$ is a standard basis vector of \mathbb{F}_q^{N+1} , v^T denotes the transpose of vector v , and k is the chosen group key. The vector ACV controls the access to the group key k and is called an *access control vector*. Svr lets $PI = \langle ACV, (z_1, z_2, \dots, z_N) \rangle$, and outputs public PI and private k .

KeyDer(s_i, PI): Using its secret s_i and the public information tuple PI , Usr_i computes $a_{i,j}, 1 \leq j \leq N$, as in formula (1) and sets an $(N + 1)$ -dimensional row \mathbb{F}_q -vector $v_i = (1, a_{i,1}, a_{i,2}, \dots, a_{i,N})$. v_i is called a Key Extraction Vector (KEV) and corresponds to a unique row in the access control matrix A . Usr_i derives the key k' from the inner product of v_i and ACV : $k' = v_i \cdot ACV$.

The derived key k' is equal to the actual group key k if and only if s_i is a valid secret used in the computation of PI , i.e., $s_i \in \mathbf{S}$.

Update(S): It runs the **KeyGen(S)** algorithm and outputs the new public information PI' and the new group key k' .

The above construction becomes impractical with large numbers of users since the complexity of the matrix and the public information is $O(n)$. We propose an improved scheme in Section 4.4 while keeping the underlying scheme unchanged.

4 SCHEME 1: INLINE AB-GKM

Recall that in its basic form, a BGKM scheme can be considered as a 1-out-of- m AB-GKM scheme. If Usr_i possesses the attribute attr_j , Svr shares a unique secret $s_{i,j}$ with Usr_i . Usr_i is thus able to derive the symmetric group key if and only if Usr_i shares at least one secret with Svr and that secret is included in the computation of the public information tuple PI . In order for Svr to revoke Usr_j , it only needs to remove the secrets it shares with Usr_j from the computation of PI ; the secrets issued to other group members are not affected. We extend this scheme to support arbitrary monotonic policies, \mathbf{P} s, over

a set of attributes. A user is able to derive the symmetric group key if and only if the set of attributes the user possesses satisfy \mathbf{P} .

As in the basic BGKM scheme, Usr_i having attr_j is associated with a unique secret value $s_{i,j}$. However, unlike the basic BGKM scheme, PI is generated by using the aggregated secrets that are generated combining the secrets issued to users according to \mathbf{P} . For example, if \mathbf{P} is a conjunction of two attributes, that is $\text{attr}_r \wedge \text{attr}_s$, the corresponding secrets $s_{i,r}$ and $s_{i,s}$ for each Usr_i are combined as one aggregated secret $s_{i,r} || s_{i,s}$ and PI is computed using these aggregated secrets. By construction, the aggregated secrets are unique since the constituent secrets are unique. Any Usr_i is able to derive the symmetric group key if and only if Usr_i has at least one aggregated secret used to compute PI . Notice that multiple users cannot collude to create an aggregated secret which they cannot individually create since $s_{i,j}$'s are unique and each aggregated secret is tied to one specific user. Hence, colluding users cannot derive the group symmetric key. Now we give a detailed description of our first AB-GKM scheme, inline AB-GKM.

4.1 Our construction

Inline AB-GKM consists of the following five algorithms:

Setup(ℓ): The Svr initializes the following parameters: an ℓ -bit prime number q , a cryptographic hash function $H(\cdot) : \{0, 1\}^* \rightarrow \mathbb{F}_q$, where \mathbb{F}_q is a finite field with q elements, the keyspace $\mathcal{KS} = \mathbb{F}_q$, the secret space $\mathcal{SS} = \{0, 1\}^\ell$, and the set of issued secrets $\mathbf{S} = \emptyset$. The user-attribute matrix UA is initialized with empty elements and the maximum group size N is decided in the **KeyGen**. It defines the universe of attributes $\mathcal{A} = \{\text{attr}_1, \text{attr}_2, \dots, \text{attr}_m\}$.

SecGen(γ_i): For each attribute $\text{attr}_j \in \gamma_i$, where $\gamma_i \subset \mathcal{A}$ and γ_i is the attribute set of Usr_i , the Svr chooses the secret $s_{i,j} \in \mathcal{SS}$ uniformly at random for Usr_i such that $s_{i,j} \notin \mathbf{S}$, adds $s_{i,j}$ to \mathbf{S} , sets $UA(i, j) = s_{i,j}$, where $UA(i, j)$ is the $(i, j)^{\text{th}}$ element of the user-attribute matrix UA , and finally outputs $s_{i,j}$.

KeyGen(P): We first give a high-level description of the algorithm and then the details. Svr transforms the policy \mathbf{P} to disjunctive normal form (DNF). For each disjunctive clause of \mathbf{P} in DNF, it creates an aggregated secret (\hat{s}) from the secrets corresponding to each of the attributes in the conjunctive clause. \hat{s} is formed by concatenation only if secrets exist for all the attributes in a given row of the user-attribute matrix UA . The construction creates a unique aggregated secret \hat{s} since the corresponding secrets are unique. For example, if the conjunctive clause is $\text{attr}_p \wedge \text{attr}_q \wedge \text{attr}_r$, for each row i in UA , the aggregated secret \hat{s}_i is formed only if all elements $UA(i, p)$, $UA(i, q)$ and $UA(i, r)$ have secrets assigned. All the aggregated secrets are added to the set \mathbf{AS} . Finally, Svr invokes algorithm **KeyGen(AS)** from the underlying BGKM scheme to output the public information PI and the symmetric group key k .

Now we give the details of the algorithm. Svr converts P to DNF as follows

$$P = \bigvee_{i=1}^{\alpha} \text{conjunct}_i \text{ where there are } \alpha \text{ conjuncts and}$$

$$\text{conjunct}_i = \bigwedge_{j=1}^{\phi_i} \text{cond}_j^{(i)},$$

where each conjunct_i has ϕ_i conditions.

A simple multiplication of clauses ($x \wedge (y \vee z) = (x \wedge y) \vee (x \wedge z)$) and then application of the absorption law ($x \vee (x \wedge y = x)$) are sufficient to convert monotone policies to DNF. Even though there can be an exponential blow up of clauses during multiplication, it has been shown that with the application of the absorption law the number of clauses in the DNF, at the end, is always polynomially bounded. Svr selects N such that

$$N \geq \sum_{i=1}^{\alpha} N_{U_i} = N_U$$

where N_{U_i} is the number of users satisfying conjunct_i ³. Svr creates N_U \hat{s}_i 's and adds them to **AS**. Svr picks a random $k \in \mathcal{KS}$ as the shared group key. Svr chooses N random bit strings $z_1, z_2, \dots, z_N \in \{0, 1\}^\ell$. Svr creates an $m \times (N + 1)$ \mathbb{F}_q -matrix A such that for $1 \leq i \leq N_U$

$$a_{i,j} = \begin{cases} 1 & \text{if } j = 1 \\ H(\hat{s}_i || z_j) & \text{if } 2 \leq j \leq N; \hat{s}_i \in \mathbf{AS} \end{cases} \quad (2)$$

Svr then solves for a nonzero $(N + 1)$ -dimensional column \mathbb{F}_q -vector Y such that $AY = 0$ and sets

$$ACV = k \cdot e_1^T + Y, \text{ and}$$

$$PI = \langle ACV, (z_1, z_2, \dots, z_N) \rangle$$

KeyDer(β_i, PI): Given β_i , the set of secrets for Usr_i , it computes the aggregated secret \hat{s} . Using \hat{s} and the public information PI , it computes $a_{i,j}, 1 \leq j \leq N$, as in formula 2 and sets an $(N + 1)$ -dimensional row \mathbb{F}_q -vector $v_i = (1, a_{i,1}, a_{i,2}, \dots, a_{i,N})$. Usr_i derives the group key k' by the inner product of the vectors v_i and ACV : $k' = v_i \cdot ACV$. The derived group key k' is equal to the actual group key k if and only if the computed aggregated secret $\hat{s} \in \mathbf{AS}$.

Update(S): The composition of the user group changes when one of the following occurs:

- Identity attributes are added or removed resulting in the change in S and UA ⁴.
- The underlying policy P changes.

When such a change occurs, a new symmetric key k' is selected and **KeyGen**(P) is invoked to generate the updated public information PI' . Notice that the secrets shared with existing users are not affected by the group change. It outputs the public PI' and private k' .

3. It should be noted that N_U can be reduced to n , the number of users in the group, by exploiting the relationships between conjuncts and letting the users know the conjunct, out of the many they satisfy, they have to use to derive the key. We leave this optimization to keep the scheme simple.

4. A change in a user attribute is viewed as two events; removing the existing attribute and adding a new attribute.

4.2 Security

We can easily show that if an *unbounded* adversary \mathcal{A} can break the inline AB-GKM scheme in the random oracle model, a simulator \mathcal{S} can be constructed to break the ACV-BGKM scheme.

Definition 1 (Security game for AB-GKM):

Setup The challenger runs the Setup algorithm of AB-GKM and gives the public parameters to the adversary.

Phase 1 The adversary is allowed to request secrets for any set of attributes γ_i and the public information tuples for a policy satisfying these attributes. The public information along with the secrets allows the adversary to derive the private key.

Challenge The adversary declares the set of attributes γ that it wishes to be challenged upon. γ is different from any of the attribute sets γ_i that the adversary queried earlier. The adversary submits two keys k_0 and k_1 . The challenger flips a random coin b and chooses k_b . The challenger generates public information for a policy P satisfying γ , but not any γ_i , using the **KeyGen** algorithm and give it to the adversary. The public information hides the group key k_b .

Phase 2 Phase 1 is repeated as many times provided that the adversary's attribute set does not satisfy P .

Guess The adversary outputs a guess b' of b .

The advantage of an adversary \mathcal{A} in this game is defined as $Pr[b' = b] - 1/2$.

Now we define security under the random oracle model.

Definition 2 (Security under the random oracle model):

An AB-GKM scheme is secure under the random oracle model of security if all adversaries have at most a negligible advantage in the above game.

Shang et al. [6], [25] have shown that the probability of breaking ACV-BGKM is a negligible $1/q$, where q is the ℓ bit large prime number initialized in **Setup**. We capture the hardness of the ACV-BGKM scheme in the following assumption:

Definition 3 (ACV-BGKM Assumption): No adversary without any valid secrets in the random oracle model can break the ACV-BGKM scheme with more than a negligible probability.

Theorem 1: If an adversary can break the inline AB-GKM scheme in the random oracle model, then a simulator can be constructed to break the ACV-BGKM scheme with non-negligible advantage.

Proof: Suppose that there exists an adversary \mathcal{A} that can break our scheme in the random oracle model with advantage ϵ . We build a simulator \mathcal{B} that can break the ACV-BGKM scheme with the advantage at most ϵ . The simulation proceeds as follows:

The challenger runs the setup algorithm of ACV-BGKM and generates secrets for each attributes per user outside of \mathcal{B} 's view. The simulator \mathcal{B} runs \mathcal{A} . \mathcal{B} is given

an instance of ACV-BGKM and gives the public parameters to \mathcal{A} . We assume that all policies are in DNF such that each conjunctive term has only one attribute. The intuition behind the assumption is that inline AB-GKM is an extension of ACV-BGKM to support aggregate secrets and, therefore, in the absent of aggregate secrets, inline AB-GKM is equivalent to ACV-BGKM.

Phase 1 \mathcal{A} submits sets of attributes γ_i to \mathcal{B} and \mathcal{B} sends the secrets using the ACV-BGKM instance.

Challenge \mathcal{A} submits the attribute set $\gamma \neq \gamma_i$ as the challenge and two keys k_0 and k_b . \mathcal{B} flips a random coin b and chooses k_b and then using the ACV-BGKM instance, it generates the public information for a policy P that only γ satisfies hiding k_b .

Phase 2 \mathcal{A} and \mathcal{B} repeats Phase 1 as many times provided \mathcal{A} 's attribute sets do not satisfy P .

Guess Using the public information and the information gathered from the two phases, \mathcal{A} outputs a guess b' of b . Notice that the view of \mathcal{A} when it is run as a subroutine of \mathcal{B} and when it is run directly with the inline AB-GKM scheme is identical. In other words, \mathcal{B} simulates an instance of the inline AB-GKM for \mathcal{A} using an instance of the ACV-BGKM scheme. The simulation is trivial as the aggregate secrets in AB-GKM is the same the secrets in ACV-BGKM. It should be noted that \mathcal{A} does not have an advantage more than ϵ from the information gather from the repeated execution of Phase 1 due to the key indistinguishability and key independence properties of the ACV-BGKM scheme [25].

It can easily be seen that \mathcal{B} has the same advantage of breaking the ACV-BGKM scheme as \mathcal{A} has on the inline AB-GKM scheme. As per the definitions, \mathcal{B} breaks the ACV-BGKM with $Pr[b' = b] = 1/2 + \epsilon$. According to the assumption on the hardness of the ACV-BGKM scheme in Theorem 1, it follows that ϵ must be negligible. \square

4.3 Performance

Now, we discuss the efficiency of inline AB-GKM with respect to computational costs and required bandwidth for rekeying.

For any Usr_i in the group, deriving the shared group key requires N hashing operations (evaluations of $H(\cdot)$) and an inner product computation $v_i \cdot ACV$ of two $(N + 1)$ -dimensional \mathbb{F}_q -vectors, where N is the maximum group size. Therefore the overall computational complexity is $O(n)$.

For every rekeying operation, Svr needs to form a matrix A by performing N^2 hashing operations, and then solve a linear system of size $N \times (N + 1)$. Solving the linear system is the most costly operation as N gets large for computation on Svr . It requires $O(n^3)$ field operations in \mathbb{F}_q when the method of Gauss-Jordan elimination [27] is applied. Experimental results about the ACV-BGKM scheme [6] have shown that this can be performed in a short time when N is small.

When a rekeying process takes place, the new information to be broadcast is $PI = \langle ACV, (z_1, \dots, z_N) \rangle$, where

ACV is a vector consisting of $(N + 1)$ elements in \mathbb{F}_q , and without loss of generality we can pick z_i to be strings of fixed length. This gives an overall communication complexity $O(n)$. An advantage of inline AB-GKM is that no peer-to-peer private channel is needed for any persisting group members when rekeying is executed.

Nowadays we generally care less about storage costs on both Svr and Usrs . Nevertheless, for a group of maximum N users, in the worst case, inline AB-GKM only requires each Usr to store $(O(|\mathcal{A}|))$ secrets, one secret per attribute that Usr possesses, and Svr to keep track of all $O(n|\mathcal{A}|)$ secrets.

4.4 Improving the Complexity using Subset-Cover

The above approach becomes inefficient if each N_{U_i} values are large as the computational and communication complexities are still proportional to N_{U_i} values. We utilize the result from previous research on broadcast encryption [13], [14] to improve the complexity. Based on that, one can make the complexity sub-linear in the number of users by giving more than one secret during **SecGen** for each attribute users possess. The secrets given to each user overlaps with different subsets of users. During the **KeyGen**, Svr identifies the minimum number of subsets to which all the users belong and uses one secret per the identified subset. During **KeyDer**, a user identifies the subset it belongs to and uses the corresponding secret to derive the group key. Group dynamics are handled by making some of the secrets given to users invalid.

We give a high-level description of the basic *subset-cover* approach. In the basic scheme, n users are organized as the leaves of a balanced binary tree of height $\log n$. A unique secret is assigned to each vertex in the tree. Each user is given $\log n$ secrets that correspond to the vertices along the path from its leaf node to the root node. In order to provide backward secrecy when a single user is revoked, the updated tree is described by $\log n$ subtrees formed after removing all the vertices along the path from the user leaf node to the root node. To rekey, Svr executes **Update** using the $\log n$ secrets corresponding to the roots of these subtrees. Naor et. al. [13] improve this technique to simultaneously revoke r users and describe the exiting users using $r \log(n/r)$ subtrees. Since then, there have been many improvements to the basic scheme. In the remainder of the paper, in order to maintain the simplicity the **SecGen** algorithm generates only one secret per attribute but the schemes can be trivially modified to use any subset-cover approach.

5 SCHEME 2: THRESHOLD AB-GKM

Consider now the case of policies by which a user can derive the symmetric group key k , if it possesses at least d attributes out of the m attributes associated with the group. We refer to such policies as *threshold policies*. Under the inline AB-GKM scheme presented in

Section 4, with such threshold policies the size of the access control matrix (A) increases exponentially if users are not informed which attributes to use. Specifically, to support d -out-of- m , the inline AB-GKM scheme may require creating a matrix of dimension up to $O(nm^d)$ where n is the number of users in the group. Thus, the inline AB-GKM scheme is not suitable for threshold policies. In this section, we construct a new scheme, threshold AB-GKM, which overcomes this shortcoming.

An initial construction to enforce threshold policies is to associate each user with a random $d - 1$ degree polynomial, $q(x)$, with the restriction that each polynomial has the same value at $x = 0$ and $q(0) = k$, where k is the symmetric group key. For each attribute users have, they are given a secret value. The secret values given to a user are tied to its random polynomial $q(x)$. A user having d or more secrets can perform a Lagrange interpolation to obtain $q(x)$ and thus the symmetric group key $k = q(0)$. Since the secrets are tied to random polynomials, multiple users are unable to combine their secrets in any way that makes possible collusion attacks. However, revocation is difficult in this simple approach and requires re-issuing all the secrets again.

Our approach to address the revocation problem is to use a layer of indirection between the secrets given to users and the random polynomials such that revocations do not require re-issuing all the secrets again. We use a modified ACV-BGKM construction as the indirection layer. We cannot directly use the ACV-BGKM construction since, multiple instances of ACV-BGKM allow collusion attacks in which colluding users can recover the group key which they cannot obtain individually. We first show the details of the modified ACV-BGKM scheme and then present the threshold AB-GKM which uses the modified ACV-BGKM scheme and Shamir's secret sharing scheme.

5.1 Modified ACV-BGKM Scheme

The modified ACV-BGKM works under similar conditions as ACV-BGKM, but instead of giving the same key k to all the users, the **KeyDer** algorithm gives each Usr_i a different key k_i when the public information tuple PI is combined with their unique secret s_i .

The algorithms are executed with a trusted key server Svr and a group of users Usr_i , $i = 1, 2, \dots, n$ with the attribute universe $\mathcal{A} = \{\text{attr}_1, \text{attr}_2, \dots, \text{attr}_m\}$. The construction is as follows:

Setup(ℓ): Svr initializes the following parameters: an ℓ -bit prime number q , the maximum group size $N (\geq n)$, a cryptographic hash function $H(\cdot) : \{0, 1\}^* \rightarrow \mathbb{F}_q$, where \mathbb{F}_q is a finite field with q elements, the key space $\mathcal{KS} = \mathbb{F}_q$, the secret space $\mathcal{SS} = \{0, 1\}^\ell$ and the set of issued secret tuples $\mathbf{S} = \emptyset$. Each Usr_i is given a unique secret index $1 \leq i \leq N$.

SecGen(\emptyset): The Svr chooses the secret $s_i \in \mathcal{SS}$ uniformly at random for Usr_i such that s_i is unique among all the users, adds the secret tuple $\langle i, s_i \rangle$ to \mathbf{S} , and outputs $\langle i, s_i \rangle$.

KeyGen(\mathbf{S}, \mathbf{K}): Given the set of secret tuples $\mathbf{S} = \{\langle i, s_i \rangle | 1 \leq i \leq N\}$ and a random set of keys $\mathbf{K} = \{k_i | 1 \leq i \leq N\}$, it outputs the public information tuple PI which allows each Usr_i to derive the key k_i using its secret s_i . The details follow.

Svr chooses N random bit strings $z_1, z_2, \dots, z_N \in \{0, 1\}^\ell$ and creates an $N \times 2N$ \mathbb{F}_q -matrix A where for a given row i , $1 \leq i \leq N$

$$a_{i,j} = \begin{cases} 1 & \text{if } i = j \\ 0 & \text{if } 1 \leq j \leq N \text{ and } i \neq j \\ H(s_i || z_j) & \text{if } N < j \leq 2N \end{cases}$$

Like in the ACV-BGKM scheme, Svr computes the null space of A with a set of its N basis vectors, and selects a vector Y as one of the basis vectors. Svr constructs an $2N$ -dimensional \mathbb{F}_q -vector

$$ACV = \left(\sum_{i=1}^N k_i \cdot e_i^T \right) + Y,$$

where e_i is the i^{th} standard basis vector of \mathbb{F}_q^{2N} . Notice that, unlike ACV-BGKM, a unique key corresponding to Usr_i , $k_i \in \mathbf{K}$ is embedded into each location corresponding to a valid index i . Like, ACV-BGKM, Svr sets $PI = \langle ACV, (z_1, z_2, \dots, z_N) \rangle$, and outputs PI via the broadcast channel.

KeyDer(s_i, PI): Usr_i , using its secret s_i and public PI , derives the $2N$ -dimensional row \mathbb{F}_q -vector v_i which corresponds to a row in A . Then Usr_i derives the specific key as $k_i = v_i \cdot ACV$.

Update(\mathbf{S}, \mathbf{K}'): If a user leaves or join the group, a new set of keys \mathbf{K}' is selected. **KeyGen**(\mathbf{S}, \mathbf{K}') is invoked to generate the updated public information PI' . Notice that the secrets shared with existing users are not affected by the group change. It outputs the public PI' .

We refer the reader to our technical report [28] for the proof of security of the above modified ACV-BGKM scheme.

5.2 Our Construction

Now we provide our construction of the threshold AB-GKM scheme which utilizes the modified ACV-BGKM scheme.

Recall that in this scheme, we wish to allow a user to derive the symmetric group key k if the user possesses at least d attributes out of m . For each user Usr_i we associate a random $d - 1$ degree polynomial $q_i(x)$ with the restriction that each polynomial has the same value k , the symmetric group key, at $x = 0$, that is, $q_i(0) = k$. We associate a random secret value with each user attribute. For each attribute attr_i , we generate a public information tuple (PI_i) using the modified ACV-BGKM scheme with the restriction that the temporary key that each Usr_j derives is tied to its random polynomial $q_j(x)$, that is $q_j(i) = k_i$. Notice that each user obtains different temporary keys from the same PI . If a user can derive d temporary keys corresponding to d attributes, it can compute its random function $q(x)$ and obtain the group

symmetric key k . Notice that, since the temporary keys are tied to a unique polynomial, multiple users are unable to collude and combine their temporary keys in order to obtain the symmetric group key which they are not allowed to obtain individually. Thus, our construction prevents collusion attacks.

A detailed description of our threshold AB-GKM scheme follows.

Setup(ℓ) Svr initializes the parameters of the underlying modified ACV-BGKM scheme: the ℓ -bit prime number q , the maximum group size N ($\geq n$), the cryptographic hash function H , the key space \mathcal{KS} , the secret space \mathcal{SS} , the set of issued secrets \mathbf{S} , the user-attribute matrix UA and the universe of attributes $\mathcal{A} = \{\text{attr}_1, \text{attr}_2, \dots, \text{attr}_m\}$.

Svr defines the Lagrange coefficient $\Delta_{i,\mathbf{Q}}$ for $i \in \mathbb{F}_q$ and a set, \mathbf{Q} of elements in \mathbb{F}_q :

$$\Delta_{i,\mathbf{Q}}(x) = \prod_{j \in \mathbf{Q}, j \neq i} \frac{x - j}{i - j}.$$

SecGen(γ_i) For each attribute $\text{attr}_j \in \gamma_i$, where $\gamma_i \subset \mathcal{A}$ and γ_i is the attribute set of Usr_i , Svr invokes **SecGen**() of the modified ACV-BGKM scheme in order to obtain the random secret $s_{i,j}$. It returns β_i , the set of secrets for all the attributes in γ_i .

KeyGen(α, d) Taking α , a subset of attributes from the attribute universe \mathcal{A} and d , the threshold value, for each user Usr_i , Svr assigns a random degree $d-1$ polynomial $q_i(x)$ with $q_i(0)$ set to the group symmetric key k . For each attribute attr_j in the set of attributes α ($\alpha \subset \mathcal{A}$ and $|\alpha| \geq d$), it selects the set of secrets corresponding to attr_j , \mathbf{S}_j and invokes **KeyGen**($S_j, \{q_1(j), q_2(j), \dots, q_N(j)\}$) of the modified ACV-BGKM scheme to obtain PI_j , the public information tuple for attr_j . It outputs the private group key k and the set of public information tuples $\mathbf{PI} = \{PI_j \mid \text{attr}_j \in \alpha\}$.

KeyDer(β_i, \mathbf{PI}) Using the set of d secrets $\beta_i = \{s_{i,j} \mid 1 \leq j \leq N\}$ for the d attributes attr_j , $1 \leq j \leq N$, and the corresponding d public information tuples $PI_j \in \mathbf{PI}$, $1 \leq j \leq N$, it derives the group symmetric key k as follows.

First, it derives the temporary key k_j for each attribute attr_j using the underlying modified ACV-BGKM scheme as **KeyDer**($s_{i,j}, PI_j$). Then, using the set of d points $\mathbf{Q}_i = \{(j, k_j) \mid 1 \leq j \leq N\}$, it computes $q_i(x)$ as follows:

$$\Delta_{j,\mathbf{Q}_i}(x) = \prod_{j \in \mathbf{Q}_i, j \neq i} \frac{x - j}{i - j}$$

$$q_i(x) = \sum_{j \in \mathbf{Q}_i} k_j \Delta_{j,\mathbf{Q}_i}(x).$$

It outputs the group key $k = q_i(0)$.

Update(α, d) The Update algorithm is invoked whenever α , the attribute set considered, or d , the threshold value, or the group members satisfying the threshold policy change. The group membership changes due to similar reasons mentioned under the **Update** algorithm in Section 4.1. In such a situation, a new symmetric group key

k' is selected and **KeyGen**(α, d) is invoked to generate the set of new public information tuples \mathbf{PI}' . Notice that the secrets shared with existing users are not affected by the group change.

5.3 Security

If an unbounded adversary can break our threshold AB-GKM scheme, a simulator can be constructed to break the modified ACV-BGKM scheme. We only give a high-level detail of the reduction based proof due to the space constraint.

Proof: Suppose that an unbounded adversary \mathcal{A} having a set of $d-1$ attributes α can break our scheme in the random oracle model with advantage ϵ . Note that this is the most powerful adversary as it possesses $d-1$ attributes out of the d attributes required to derive the group key. We build a simulator \mathcal{B} that can derive the key k_d from PI_d corresponding to $\text{attr}_d \notin \alpha$ with the same advantage ϵ using \mathcal{A} as subroutine. In other words, we build a simulator to break the modified ACV-BGKM scheme.

The intuition behind our proof is that, by construction, the modified ACV-BGKM instances corresponding to the attributes are independent. In other words, a user who can access the key for one attribute only has a negligible advantage in obtaining the key for another attribute using the known attributes due to the key indistinguishability and independence properties of the ACV-BGKM scheme [25].

The challenger creates an instance of the modified ACV-BGKM scheme for each of the n attributes. \mathcal{A} obtains secrets $\{s_i \mid i = 1, 2, \dots, d-1\}$ for the attributes α it has from \mathcal{B} . The challenger constructs the public information tuples $\{PI_i \mid i = 1, 2, \dots, d\}$, each having a random key k_i and gives them to \mathcal{B} . \mathcal{B} in turn gives them to \mathcal{A} . Notice that the view of \mathcal{A} is identical to that of \mathcal{A} interacting directly with an instance of the threshold AB-GKM scheme, even though it is simulated. The random keys correspond to a random degree $d-1$ polynomial $q(x)$. Notice that \mathcal{A} possesses secrets to obtain the random keys k_i , $1 \leq i \leq d-1$ and can derive the secret key k_d with an advantage ϵ from the public information tuples.

We omit the details of the security game defined in the previous section. As mentioned in the game, \mathcal{A} may execute the threshold AB-GKM scheme for different sets of attributes that do not satisfy the challenge threshold policy and do not include attr_d . As mentioned earlier, \mathcal{A} does not gain any additional advantage by such executions.

After executing the phase 1 of the security game as many times, \mathcal{A} outputs k , which is equal to $q(0)$. This allows \mathcal{B} to fully determine $q(x)$ as it now has d points and derive the key $k_d = q(d)$. In other words, it allows \mathcal{B} to break the modified ACV-BGKM scheme to recover the intermediate key k_d from the public information tuple PI_d without the knowledge of the secret s_d . In

our technical report [28], we show that the probability of breaking the modified ACV-BGKM scheme is a negligible $1/q^N$ where q is the ℓ bit prime number and N is the maximum number of users. Therefore, it follows that ϵ must be negligible. \square

5.4 Performance

We now discuss the efficiency of the threshold AB-GKM with respect to computational costs and required bandwidth for rekeying.

For any Usr_i in the group deriving the shared group key requires: $\sum_{i=1}^d N_i$ hashing operations (evaluations of $H(\cdot)$), where N_i is the maximum number of users having attr_i ; and d inner product computations $v_i \cdot \text{ACV}_i$ of two $(2N_i)$ -dimensional \mathbb{F}_q -vectors and the Lagrange interpolation $O(m \log^2 m)$, where $m = |\mathcal{A}|$. Therefore, the overall computational complexity is $O(dn + m \log^2 m)$. Notice that the inner product computations are independent and can be parallelized to improve performance.

For every rekeying phase, for each attr_i , Svr needs to form a matrix A_i by performing N_i^2 hashing operations, and then solve a linear system of size $N_i \times (2N_i)$. Solving the linear system is the most costly operation as N_i gets large for computation on Svr ; it requires $O(\sum_{i=1}^m n^3)$ field operations in \mathbb{F}_q .

When a rekeying process takes place, the new information to be broadcast is $PI_i = \langle \text{ACV}_i, (z_1, \dots, z_{N_i}) \rangle$, $i = 1, 2, \dots, m$, where ACV_i is a vector consisting of $(2N_i)$ elements in \mathbb{F}_q , and without loss of generality we can pick z_i to be strings with a fixed length. This gives an overall communication complexity $O(\sum_{i=1}^m n)$.

For a group of maximum N users, in the worst case, the threshold AB-GKM only requires each Usr to store $(O(m))$ secrets, one secret per attribute that Usr possesses and Svr to keep track of all $O(nm)$ secrets.

6 SCHEME 3: ACCESS TREE AB-GKM

In the inline AB-GKM scheme, the policy \mathbf{P} is embedded into the BGKM scheme itself. As discussed in Section 5, while this approach works for many different types of policies, such an approach is not able to efficiently support threshold access control policies. Scheme 2, threshold AB-GKM, on the other hand, is able to efficiently support threshold policies, but it is unable to support other policies. In order to support more expressive policies, we extend the threshold AB-GKM scheme. Like threshold AB-GKM, instead of embedding \mathbf{P} in the BGKM scheme, we construct a separate BGKM instance for each attribute. Then, we embed \mathbf{P} in an access structure \mathcal{T} . \mathcal{T} is a tree with the internal nodes representing threshold gates and the leaves representing attributes. The construction of \mathcal{T} is similar to that of the approach by Goyal et al. [3]. However, unlike Goyal et al.'s approach, the goal of our construction is to derive the group key for the users whose attributes satisfy the access structure \mathcal{T} .

TABLE 2
Access tree functions

Function	Description
$\text{index}(x)$	Returns the index of node x
$\text{parent}(x)$	Returns the parent node of node x
$\text{attr}(x)$	Returns the index of the attribute associated with a leaf node x
q_x	The polynomial assigned to node x
$\text{sat}(\mathcal{T}_x, \alpha)$	Returns 1 if the set of attributes α satisfies \mathcal{T}_x , the subtree rooted at node x , and 0 otherwise

6.1 Access Tree

Let \mathcal{T} be a tree representing an access structure. Each internal node of the tree represents a threshold gate. A threshold gate is described by its child nodes and a threshold value. If n_x is the number of children of a node x and t_x is its threshold value, then $0 < t_x \leq n_x$. Notice that when $t_x = 1$, the threshold gate is an OR gate and when $t_x = n_x$, it is an AND gate. Each leaf node x of the tree is described by an attribute, a corresponding BGKM instance and a threshold value $t_x = 1$. The children of each node x are indexed from 1 to n_x .

We define the functions in Table 2 in order to construct our scheme. All the functions except sat are straightforward to implement. A brief description of sat follows:

The function $\text{sat}(\mathcal{T}_x, \alpha)$ works as a recursive function. If x is a leaf node, it returns 1, provided that the attribute associated with x is in the set of attributes α and 0 otherwise. If x is an internal node, if at least t_x child nodes of x return 1, then $\text{sat}(\mathcal{T}_x, \alpha)$ returns 1 and 0 otherwise.

6.2 Our Construction

The access tree AB-GKM scheme consists of five algorithms:

Setup(ℓ): Svr initializes the parameters of the underlying modified ACV-BGKM scheme: the prime number q , the maximum group size N ($\geq n$), the cryptographic hash function H , the key space \mathcal{KS} , the secret space \mathcal{SS} , the set of issued secrets \mathbf{S} , the user-attribute matrix UA and the universe of attributes $\mathcal{A} = \{\text{attr}_1, \text{attr}_2, \dots, \text{attr}_m\}$.

Svr defines the Lagrange coefficient $\Delta_{i, \mathbf{Q}}$ for $i \in \mathbb{F}_q$ and a set, \mathbf{Q} of elements in \mathbb{F}_q :

$$\Delta_{i, \mathbf{Q}}(x) = \prod_{j \in \mathbf{Q}, j \neq i} \frac{x - j}{i - j}.$$

SecGen(γ_i): Taking γ_i , the attribute set of Usr_i , as input, for each attribute $\text{attr}_j \in \gamma_i$, where $\gamma_i \subset \mathcal{A}$, Svr invokes **SecGen**($\mathbf{0}$) of the modified ACV-BGKM scheme to obtain the random secret $s_{i,j}$. It returns β_i , the set of secrets for all the attributes in γ_i .

KeyGen(\mathbf{P}): Svr transforms the policy \mathbf{P} into an access tree \mathcal{T} . The algorithm outputs the public information which a user can use to derive the group key if and only if the user's attributes satisfy the access tree \mathcal{T} built

for the policy \mathbf{P} . The algorithm constructs the public information as follows.

For each user Usr_i having the intermediate set of keys $\mathbf{K}_i = \{k_{i,j} | 1 \leq j \leq m\}$, where $k_{i,j}$ represents the intermediate key for Usr_i and attr_j , the following construction is performed. For each attribute attr_i , there is a leaf node in \mathcal{T} . The construction of the tree is performed top-down. Each node x in the tree is assigned a polynomial q_x . The degree d_x of the polynomial q_x is set to $t_x - 1$, that is, one less than the threshold value of the node. For the root node r , $q_r(0)$ is set to the group key k and d_r other points are chosen uniformly at random so that q_r is a unique polynomial of degree d_r fully defined through Lagrange interpolation. For any other node x , $q_x(0)$ is set to $q_{\text{parent}(x)}(\text{index}(x))$ and d_x other points are chosen uniformly at random to uniquely define q_x . For each leaf node x corresponding to a unique attribute attr_j , $q_x(0)$ is set to $q_{\text{parent}(x)}(1)$ and $k_{i,j} = q_x(0)$.

At the end of the above computation, we have all the sets of intermediate keys $\mathbf{K} = \{\mathbf{K}_i | \text{Usr}_i, 1 \leq i \leq N\}$. For each leaf node x , the modified BGKM algorithm $\text{KeyGen}(\mathbf{S}_x, \mathbf{K}_x)$, where \mathbf{S}_x is the set of secrets corresponding to the attribute associated with the node x and $\mathbf{K}_x = \{k_{i,j} | 1 \leq i \leq N, \text{attr}_j\}$, $j = \text{attr}(x)$, is invoked to generate public information tuple PI_x . We denote the set of all the public information tuples $\mathbf{PI} = \{PI_j | \text{attr}_j, 1 \leq j \leq m\}$.

KeyDer(β_i, \mathbf{PI}): Given β_i , a set of secret values corresponding to the attributes of Usr_i , and the set of public information tuples \mathbf{PI} , it outputs the group key k .

The key derivation is a recursive procedure that takes β_i and \mathbf{PI} to derive k bottom-up. Note that a user can obtain the key if and only if its attributes satisfy the access tree \mathcal{T} , i.e., $\text{sat}(\mathcal{T}_r, \beta_i) = 1$. The high-level description of the key derivation is as follows.

For each leaf node x corresponding to the attribute with the user's secret value $s_x \in \beta_i$, the user derives the intermediate key k_x using the underlying modified BGKM scheme $\text{KeyDer}(s_x, PI_x)$. Using Lagrange interpolation, the user recursively derives the intermediate key k_x for each internal ancestor node x until the root node r is reached and $k_r = k$. Notice that since intermediate keys are tied to unique polynomials, users cannot collude to derive the group key k if they are unable to derive it individually. A detailed description follows.

If x is a leaf node, it returns an empty value \perp if $\text{attr}(x) \notin \beta_i$, otherwise it returns the key $k_x = v_x \cdot ACV_x$, where v_x is the key derivation vector corresponding to the attribute $\text{attr}_{\text{attr}(x)}$ and ACV_x the access control vector in PI_x .

If x is an internal node, it returns an empty value \perp if the number of children nodes having a non-empty key is less than t_x , otherwise it returns k_x as follows:

Let the set \mathbf{Q}_x contain the indices of t_x children nodes having non-empty keys $\{k_i | i \in \mathbf{Q}_x\}$.

$$\begin{aligned} \Delta_{i, \mathbf{Q}_x}(y) &= \prod_{i \in \mathbf{Q}_x, i \neq j} \frac{y - i}{j - i} \\ q_x(y) &= \sum_{i \in \mathbf{Q}_x} k_i \Delta_{i, \mathbf{Q}_x}(y) \\ k_x &= q_x(0). \end{aligned}$$

The above computation is performed recursively until the root node is reached. If Usr_i satisfies \mathcal{T} , Usr_i gets $k = q_r(0)$, where r is the root node. Otherwise, Usr_i gets an empty value \perp .

Update(P) The group members change due to the similar reasons mentioned for the **Update** algorithm in Section 4.1. In such a situation, a new symmetric group key k' is selected and **KeyGen(P)** is invoked to generate the set of new public information tuples \mathbf{PI}' . Like the previous two schemes, the secrets shared with existing users are not affected by the group change.

6.3 Security

If an unbounded adversary can break our access tree AB-GKM scheme, a simulator can be constructed to break the modified ACV-BGKM scheme. Like the previous scheme, we only give a high-level detail of the reduction based proof.

Proof: Suppose that an unbounded adversary \mathcal{A} using a set of attributes α as the challenge set that does not satisfy the access tree \mathcal{T} breaks our scheme in the random oracle model with advantage at most ϵ . Let the root node of \mathcal{T} be r and the group key $k = q_r(0)$. Notice that since \mathcal{A} does not satisfy \mathcal{T} and $q_r(x)$ a t_r -out-of- n_r threshold scheme, which represents any type of threshold node, \mathcal{A} satisfies no more than $t_r - 1$ subtrees rooted at children of r out of the n_r subtrees. By inference, it is easy to see that \mathcal{A} does not satisfy at least one leaf node.

The challenger constructs modified ACV-BGKM instances for each of the attributes and gives them to \mathcal{B} . \mathcal{A} obtains secrets for each of the attributes in α . \mathcal{B} sends the public information tuples and the access tree \mathcal{T} to \mathcal{A} . Notice that \mathcal{A} can easily derive the keys for any attribute in α , but it can derive the keys for any other attribute only with an advantage of ϵ . According to the assumption, \mathcal{A} does not satisfy at least one attribute required to satisfy \mathcal{T} . Let that attribute be attr_x . \mathcal{A} derives k_x from PI_x corresponding to one such unsatisfied leaf node with advantage ϵ . Therefore, \mathcal{A} derives the group key k with an advantage of at most ϵ .

Like the proof in Section 5, \mathcal{A} derives the group key k , after executing the phase 1 of the security game as many times and give k to \mathcal{B} . Now, \mathcal{B} works downwards \mathcal{T} to recover the keys for nodes originally unsatisfied by \mathcal{A} using Lagrange interpolation. For example, using k and $t_r - 1$, \mathcal{B} obtains the key k_{i_r} for the t_r^{th} child node of r . Finally, \mathcal{B} obtains the key k_x for an unsatisfied leaf node x corresponding to attr_x . In other words, it

allows \mathcal{B} to break the modified ACV-BGKM scheme to recover the key k_x from the public information tuple PI_x without the knowledge of the secret s_x . As mentioned earlier, the probability of breaking the modified ACV-BGKM scheme by applying the **KeyDer** algorithm is a negligible $1/q^N$ where q is the ℓ bit prime number and N is the maximum number of users. Therefore, it follows that ϵ must be negligible. \square

6.4 Performance

We now discuss the efficiency of access tree AB-GKM with respect to computational costs and required bandwidth for rekeying.

For any U_{sr_i} in the group, deriving the shared group key requires: $\sum_{i=1}^d N_i$ hashing operations (evaluations of $H(\cdot)$), where $d = |\beta_i|$, N_i is the maximum number of users having $attr_i$, and d inner product computations $v_i \cdot ACV_i$ of two $(2N_i)$ -dimensional \mathbb{F}_q -vectors and M Lagrange interpolations $O(Mm \log^2 m)$, where M is equal to the number of internal nodes in \mathcal{T} and $m = |\mathcal{A}|$. Therefore, the overall computational complexity is $O(dn + Mm \log^2 m)$. Notice that the inner product computations are independent and can be parallelized to improve performance.

The cost of rekeying, communication and storage are comparable to those of the threshold scheme presented in Section 5.

7 EXAMPLE APPLICATION

Among other applications, fine-grained access control in a group setting using broadcast encryption is an important application of the AB-GKM schemes. We illustrate the access-tree AB-GKM scheme using a healthcare scenario [29], [6]. We refer the reader to our technical report [28] for more examples. A hospital (Svr) supports fine-grained access control on electronic health records (EHRs) [30], [31] by encrypting and making the encrypted records available to hospital employees (Usrs). Typical hospital users include employees playing different roles such as receptionist, cashier, doctor, nurse, pharmacist, system administrator and non-employees such as patients. An EHR document is divided into subdocuments including BillingInfo, ContactInfo, Medication, PhysicalExam, LabReports and so on. In accordance with regulations such as health insurance portability and accountability act (HIPAA), the hospital policies specify which users can access which subdocument(s). A cashier, for example, need not have access to data in EHRs except for the BillingInfo, while a doctor or a nurse need not have access to BillingInfo. These policies can be based on the content of EHRs itself. An example of such policies is that “information about a patient with cancer can only be accessed by the primary doctor of the patient”. In addition, patients define their own privacy policies to protect their EHRs. For example, a patient’s policy may specify that “only the doctors and nurses who support her insurance plan can view her EHR”.

In order to support content-based access control, the hospital maintains some associations among users and data. Table 3 shows the insurance plans supported by each doctor and nurse, identified by the pseudonym “Employee ID”.

TABLE 3
Insurance Plans Supported by Doctors/Nurses

EmployeeID	Role/level	Insurance Plan(s)
emp ₁	doctor	MedB, ACME
emp ₂	doctor	ACME
emp ₃	nurse/junior	ACME
emp ₄	nurse/senior	MedA
emp ₅	nurse/senior	MedC
emp ₆	doctor	MedA
emp ₇	doctor	MedB, ACME
emp ₈	nurse/senior	MedA
emp ₉	nurse/senior	MedA, MedB, ACME

The hospital runs **Setup** algorithm to initialize system parameters and issues secrets to employees by running the **SecGen** algorithm. Table 4 shows the content of the user attribute matrix UA that the hospital maintains. (Small numbers are used for illustrative purposes.)

TABLE 4
User Attribute Matrix

Emp ID	doctor	nurse	senior	junior	MedA	MedB	MedC	ACME
emp ₁	100	⊥	⊥	⊥	⊥	111	⊥	102
emp ₂	120	⊥	⊥	⊥	⊥	⊥	⊥	105
emp ₃	⊥	106	⊥	120	⊥	⊥	⊥	121
emp ₄	⊥	103	150	⊥	175	⊥	⊥	⊥
emp ₅	⊥	133	151	⊥	⊥	⊥	161	⊥
emp ₆	129	⊥	⊥	⊥	141	⊥	⊥	⊥
emp ₇	119	⊥	⊥	⊥	⊥	133	⊥	137
emp ₈	⊥	143	152	⊥	115	⊥	⊥	⊥
emp ₉	⊥	109	156	⊥	117	119	⊥	124

Now we illustrate the use of the access tree AB-GKM scheme. Consider the following policy specification on the Medication subdocument of the EHR. “A senior nurse supporting at least two insurance plans can access Medication of any patient”. In order to implement this access control policy, we need to consider attributes role, level and insurance plan. The access control policy looks as follows:

$$P = (\text{“role = nurse”} \wedge \text{“level = senior”} \wedge \text{“2-out-of-}\{\text{MedA, MedB, MedC, ACME}\}\text{”})$$

TABLE 5
List of employees satisfying each insurance plan

Attribute	Employee IDs
MedA	emp ₄ , emp ₆ , emp ₈ , emp ₉
MedB	emp ₁ , emp ₇ , emp ₉
MedC	emp ₅
ACME	emp ₁ , emp ₂ , emp ₃ , emp ₇ , emp ₉

In addition to Table 5 containing the list of employees satisfying insurance plans, the hospital maintains the list of employees satisfying the attributes nurse and senior as shown in Table 6.

TABLE 6
List of employees satisfying attributes

Attribute	Employee IDs
nurse	emp ₃ , emp ₄ , emp ₅ , emp ₈ , emp ₉
senior	emp ₄ , emp ₅ , emp ₈ , emp ₉

The above policy can be represented using an access tree with two internal nodes and six leaf nodes. The root node is an AND gate and has three children. The first and second children of the root node represent the attributes nurse and senior, respectively, and the third child of the root node is a 2-out-of-4 threshold gate which has four children representing the four insurance plans.

The hospital executes the **KeyGen** algorithm to generate six PI tuples and encrypts the Medication subdocuments with the group symmetric key k :

$$PI_{MedA} = \langle ACV_{MedA}, (z_1, z_2, z_3, z_4) \rangle$$

$$PI_{MedB} = \langle ACV_{MedB}, (z_5, z_6, z_7) \rangle$$

$$PI_{MedC} = \langle ACV_{MedC}, (z_8) \rangle$$

$$PI_{ACME} = \langle ACV_{ACME}, (z_9, z_{10}, z_{11}, z_{12}, z_{13}) \rangle$$

$$PI_{nurse} = \langle ACV_{nurse}, (z_{14}, z_{15}, z_{16}, z_{17}, z_{18}) \rangle$$

$$PI_{senior} = \langle ACV_{senior}, (z_{19}, z_{20}, z_{21}, z_{22}) \rangle$$

Expressive access control. Notice that only one employee, emp₉, can derive the group key k using **KeyDer** algorithm to decrypt Medication subdocuments.

Collusion resistance. Notice that emp₄ supports MedA and emp₅ supports MedC and both of them are senior nurses. It may appear that these two employees can collude to derive the group key k . Since, in this particular example, the access tree AB-GKM scheme associates each user with two unique polynomials, one for the AND gate and another for the threshold gate, none of them individually satisfies the access tree and **KeyDer** results in an incorrect key.

Handling user dynamics. Assume that emp₄ starts to support the insurance plan ACME in addition to MedA. The hospital re-generates the public information by adding emp₄ to the calculation of PI_{ACME} and associating a new group key k' . Now emp₄ is able to derive k' using **KeyDer** as its attributes satisfy the access tree. Notice that the change in the user attributes does not affect the secret information each existing employees have. A similar approach is taken when one or more of these attributes are revoked from an existing employee. It should be noted that, like the first two schemes, this scheme has the added flexibility to support changes to the access tree by requiring only changes to the public information.

8 EXPERIMENTAL RESULTS

In this section we provide experimental results for the underlying optimized ACV-BGKM scheme used with all

TABLE 7
Average Time for CP-ABE algorithms

Algorithm	Time (ms)
Setup	34.395
Key generation	26.725
Encryption	24.453
Decryption	13.415

three AB-GKM schemes presented earlier. We compare our results with CP-ABE scheme with comparable security parameters.

The experiments were performed on a machine running GNU/Linux kernel version 2.6.32 with an Intel® Core™ 2 Duo CPU E8400 3.00GHz and 3.2 Gbytes memory. Only one processor was used for computation. Our prototype system is implemented in C/C++. We use V. Shoup's NTL library [32] version 5.4.2 for finite field arithmetic, and SHA-1 and AES-128 implementations of OpenSSL [33] version 1.0.0d for cryptographic hashing and symmetric key encryption. We use Bethencourt et. al.'s cpabe [34] library to gather experimental results for CP-ABE. The cpabe library uses PBC library [35] for pairing based cryptography.

We implemented the ACV-BGKM scheme with subset cover optimization. We utilized the complete subset algorithm introduced by Naor et al. [13] as the subset cover. All finite field arithmetic operations in ACV-BGKM scheme are performed in an 512-bit prime field. We used comparable and efficient pairing parameters for CP-ABE. The size of the base finite field is set to the 512-bit prime number

87807107996633125224377819847540498158068831994
14208211028653399266475630880222957078625179422
66222142315585876958231745927771336731748132492
5129998224791

and the group order to the 160-bit number

730750818665451621361119245571504901405976559617.

Following the well-known security practice, we generate symmetric keys and use them for encrypting documents. Then we encrypt such encryption keys with either the ACV-BGKM generated symmetric keys or the CP-ABE generated public keys. Therefore, in the experiments we measure the time to encrypt and decrypt the document encryption keys only. For all the ACV-BGKM experiments, we assume that 5% of users have left the group after executing the setup.

First we give experimental results for the most simplest case where a single attribute condition is considered. Then we provide, experimental results for multiple attribute conditions.

Table 7 shows the average time required to execute setup, key generation, encryption and decryption algorithms of CP-ABE scheme for one attribute condition.

Figure 1 reports the average time required to execute the key generation algorithm of ACV-BGKM and CP-ABE with different group sizes. In both ACV-BGKM and CP-ABE the time increases linearly with the group

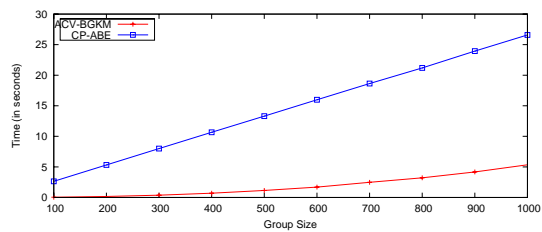


Fig. 1. Average Key Generation Time for Different Group Sizes

size. However, ACV-BGKM is much more efficient as it does not involve any expensive pairing operations. It only uses efficient hashing and binary operations over a finite field. Further, the subset cover technique applied to ACV-BGKM reduces the computational complexity of the underlying scheme. Without the subset cover optimization, ACV-BGKM has a non-linear computational complexity and becomes inefficient for large groups. We omit the comparison experimental result due to lack of space.

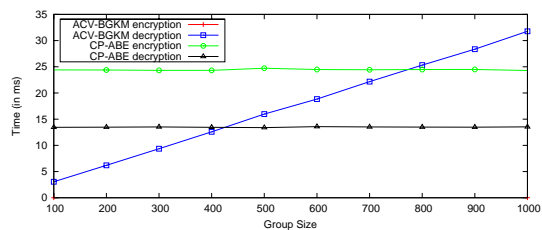


Fig. 2. Average Encryption/Decryption Time for Different Group Sizes

Figure 2 reports the average time required to perform encryption and decryption in ACV-BGKM and CP-ABE schemes for one attribute condition with different group sizes. The decryption time of ACV-BGKM is taken as the time to derive the key as well as to decrypt the encryption key. The encryption and decryption times of CP-ABE remain constant whereas the decryption time of ACV-BGKM increases linearly with the group size. As the group size increases, the key derivation algorithm of ACV-BGKM requires to spend more time to build larger KEVs. The encryption time of ACV-BGKM is negligible and remains constant as it involves an efficient symmetric encryption only. The average encryption time of ACV-BGKM is 8.8 microseconds (as these times are very small, the line plotting them is very close to zero in the graph in Figure 2 and thus overlaps with the x-axis). It should be noted that if one caches the KEVs, the decryption time of ACV-BGKM also becomes negligible as it involves only modular multiplications.

Figure 3 reports the average time required to execute the key generation algorithm with varying number of attribute conditions with the group size set to 1000. The time of both techniques increases linearly with the number of attribute conditions. However, similar to Figure 1,

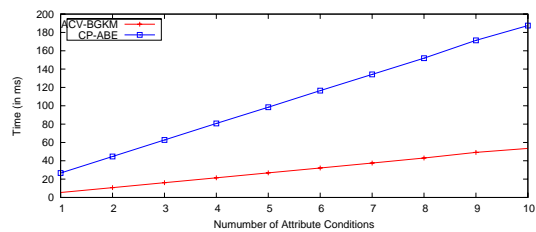


Fig. 3. Average Key Generation Time for Varying Attribute Counts

the ACV-BGKM key generation is much more efficient than the CP-ABE key generation.

As can be seen from the experiments, our constructs are more efficient in handling scenarios where the key generation algorithm has to be executed frequently due to changes in user dynamics.

9 CONCLUSION

In this paper, we have presented three attribute based group key management (AB-GKM) schemes: inline AB-GKM, threshold AB-GKM, and access tree AB-GKM. In all our schemes, when the group changes, the rekeying operations do not affect the private information of existing group members and thus our schemes eliminate the need of establishing private communication channels. Our schemes provide the same advantage when the group membership policies change. We have also shown that our schemes are resistant to collusion attacks. Our constructions are based on a provably secure ACV-BGKM scheme and Shamir's threshold scheme. Our experimental results show that our underlying construction is more efficient than the popular CP-ABE scheme.

ACKNOWLEDGMENTS

This material is based upon work supported by the Air Force Office of Scientific Research under Awards No. FA9550-08-1-0260, FA9550-09-0468 and FA9550-08-1-0265.

REFERENCES

- [1] R. Krishnan, R. Sandhu, J. Niu, and W. H. Winsborough, "Foundations for group-centric secure information sharing models," in *Proceedings of the 14th ACM symposium on Access control models and technologies*, ser. SACMAT '09. New York, NY, USA: ACM, 2009, pp. 115–124.
- [2] A. Sahai and B. Waters, "Fuzzy identity-based encryption," in *Eurocrypt 2005, LNCS 3494*. Springer-Verlag, 2005, pp. 457–473.
- [3] V. Goyal, O. Pandey, A. Sahai, and B. Waters, "Attribute-based encryption for fine-grained access control of encrypted data," in *CCS '06: Proceedings of the 13th ACM conference on Computer and communications security*. New York, NY, USA: ACM, 2006, pp. 89–98.
- [4] J. Bethencourt, A. Sahai, and B. Waters, "Ciphertext-policy attribute-based encryption," in *SP '07: Proceedings of the 2007 IEEE Symposium on Security and Privacy*. Washington, DC, USA: IEEE Computer Society, 2007, pp. 321–334.
- [5] X. Zou, Y. Dai, and E. Bertino, "A practical and flexible key management mechanism for trusted collaborative computing," *INFOCOM 2008. The 27th Conference on Computer Communications*. IEEE, pp. 538–546, April 2008.

- [6] N. Shang, M. Nabeel, F. Paci, and E. Bertino, "A privacy-preserving approach to policy-based content dissemination," in *ICDE '10: Proceedings of the 2010 IEEE 26th International Conference on Data Engineering*, 2010.
- [7] A. Shamir, "How to share a secret," *Commun. ACM*, vol. 22, no. 11, pp. 612–613, 1979.
- [8] H. Harney and C. Muckenhirn, "Group key management protocol specification," Network Working Group, United States, Tech. Rep., 1997.
- [9] H. Chu, L. Qiao, K. Nahrstedt, H. Wang, and R. Jain, "A secure multicast protocol with copyright protection," *SIGCOMM Comput. Commun. Rev.*, vol. 32, no. 2, pp. 42–60, 2002.
- [10] C. Wong and S. Lam, "Keystone: a group key management service," in *International Conference on Telecommunications, ICT*, 2000.
- [11] A. Sherman and D. McGrew, "Key establishment in large dynamic groups using one-way function trees," *Software Engineering, IEEE Transactions on*, vol. 29, no. 5, pp. 444–458, May 2003.
- [12] S. Berkovits, "How to broadcast a secret," in *EUROCRYPT '91: Proceedings of the 10th annual international conference on Advances in Cryptology*. Berlin, Heidelberg: Springer-Verlag, 1991, pp. 535–541.
- [13] D. Naor, M. Naor, and J. B. Latspiech, "Revocation and tracing schemes for stateless receivers," in *Proceedings of the 21st Annual International Cryptology Conference on Advances in Cryptology*, ser. CRYPTO '01. London, UK: Springer-Verlag, 2001, pp. 41–62.
- [14] D. Halevy and A. Shamir, "The lsd broadcast encryption scheme," in *Proceedings of the 22nd Annual International Cryptology Conference on Advances in Cryptology*, ser. CRYPTO '02. London, UK: Springer-Verlag, 2002, pp. 47–60.
- [15] P. Traynor, K. R. B. Butler, W. Enck, and P. McDaniel, "Realizing massive-scale conditional access systems through attribute-based cryptosystems," in *Proceedings of the Network and Distributed System Security Symposium, 2008*, ser. NDSS 2008, 2008.
- [16] L. Cheung, J. A. Cooley, R. Khazan, and C. Newport, "Collusion-resistant group key management using attribute-based encryption. cryptology eprint archive report 2007/161," 2007.
- [17] S. Yu, K. Ren, and W. Lou, "Attribute-based on-demand multicast group setup with membership anonymity," in *Proceedings of the 4th international conference on Security and privacy in communication networks*, ser. SecureComm '08. New York, NY, USA: ACM, 2008, pp. 18:1–18:6.
- [18] I. Chang, R. Engel, D. Kandlur, D. Pendarakis, and D. Saha, "Key management for secure internet multicast using boolean function minimization techniques," in *INFOCOM 1999. The 18th Conference on Computer Communications*. IEEE, 1999, pp. 689–698.
- [19] M. Waldvogel, G. Caronni, D. Sun, N. Weiler, and B. Plattner, "The versakey framework: versatile group key management," *Selected Areas in Communications, IEEE Journal on*, vol. 17, no. 9, pp. 1614–1631, Sep. 1999.
- [20] A. Fiat and M. Naor, "Broadcast encryption," in *Advances in Cryptology - CRYPTO 93*, ser. Lecture Notes in Computer Science, D. Stinson, Ed., vol. 773. Springer Berlin / Heidelberg, 1994, pp. 480–491.
- [21] E. Bertino and E. Ferrari, "Secure and selective dissemination of XML documents," *ACM Trans. Inf. Syst. Secur.*, vol. 5, no. 3, pp. 290–331, 2002.
- [22] G. Miklau and D. Suci, "Controlling access to published data using cryptography," in *VLDB '2003: Proceedings of the 29th international conference on Very large data bases*. VLDB Endowment, 2003, pp. 898–909.
- [23] D. Halevy and A. Shamir, "The lsd broadcast encryption scheme," in *Advances in Cryptology CRYPTO 2002*, ser. Lecture Notes in Computer Science, M. Yung, Ed., vol. 2442. Springer Berlin / Heidelberg, 2002, pp. 145–161.
- [24] D. Boneh, C. Gentry, and B. Waters, "Collusion resistant broadcast encryption with short ciphertexts and private keys," in *Advances in Cryptology CRYPTO 2005*, ser. Lecture Notes in Computer Science, V. Shoup, Ed., vol. 3621. Springer Berlin / Heidelberg, 2005, pp. 258–275.
- [25] N. Shang, M. Nabeel, E. Bertino, and X. Zou, "Broadcast group key management with access control vectors," Department of Computer Science, Tech. Rep., 4 2010.
- [26] Y. Challal and H. Seba, "Group key management protocols: A novel taxonomy," *International Journal of Information Technology*, vol. 2, no. 2, pp. 105–118, 2006.
- [27] D. Dummit and R. Foote, "Gaussian-Jordan elimination," in *Abstract Algebra*, 2nd ed. Wiley, 1999, p. 404.
- [28] M. Nabeel and E. Bertino, "Attribute based group key management," Purdue University, Tech. Rep. CERIAS TR 2010, 2010.
- [29] M. Pirretti, P. Traynor, P. McDaniel, and B. Waters, "Secure attribute-based systems," in *CCS '06: Proceedings of the 13th ACM conference on Computer and communications security*. New York, NY, USA: ACM, 2006, pp. 99–112.
- [30] "XML in clinical research and healthcare industries," <http://xml.coverpages.org/healthcare.html>.
- [31] M. Eichelberg, T. Aden, J. Riesmeier, A. Dogac, and G. B. Laleci, "A survey and analysis of electronic healthcare record standards," *ACM Comput. Surv.*, vol. 37, no. 4, pp. 277–315, 2005.
- [32] V. Shoup, "NTL library for doing number theory," <http://www.shoup.net/ntl/>.
- [33] "OpenSSL the open source toolkit for SSL/TLS," <http://www.openssl.org/>.
- [34] J. Bethencourt, A. Sahai, and B. Waters, "Ciphertext policy attribute based encryption library," <http://http://acsc.cs.utexas.edu/cpabe/>.
- [35] B. Lynn, "Pairing based cryptography library," <http://crypto.stanford.edu/pcb/>.

Mohamed Nabeel is a PhD candidate at the department of computer science, Purdue university. He is also a member of the Center for Education and Research in Information Assurance and Security (CERIAS), IEEE and ACM. His research interests are in data privacy, distributed system security and applied cryptography. His PhD thesis topic is "Privacy preserving content dissemination using attribute based group key management". His research adviser is prof. Elisa Bertino. He has published in the areas of privacy preserving content dissemination and group key management. He received the Fulbright fellowship in 2006, Purdue Cyber Center research grant in 2010 and Purdue research foundation grant in 2011.

Elisa Bertino is Professor of Computer Science at Purdue University, and serves as research director of the Center for Education and Research in Information Assurance and Security (CERIAS) and Interim Director of Cyber Center (Discovery Park). Previously, she was a faculty member and department head at the Department of Computer Science and Communication of the University of Milan. Her main research interests include security, privacy, digital identity management systems, database systems, distributed systems, and multimedia systems. She is currently serving as chair of the ACM SIGSAC and as a member of the editorial board of the following international journals: IEEE Security & Privacy, IEEE Transactions on Service Computing, ACM Transactions on Web. She also served as editor in chief of the VLDB Journal and editorial board member of ACM TISSEC and IEEE TDSC. She co-authored the book "Identity Management - Concepts, Technologies, and Systems". She is a fellow of the IEEE and a fellow of the ACM. She received the 2002 IEEE Computer Society Technical Achievement Award for outstanding contributions to database systems and database security and advanced data management systems and the 2005 IEEE Computer Society Tsutomu Kanai Award for pioneering and innovative research contributions to secure distributed systems.