

Attribute-based encryption (ABE): more fine-grained control to encrypted data

IBE: can encrypt to an identity (policy is basically checking equality)

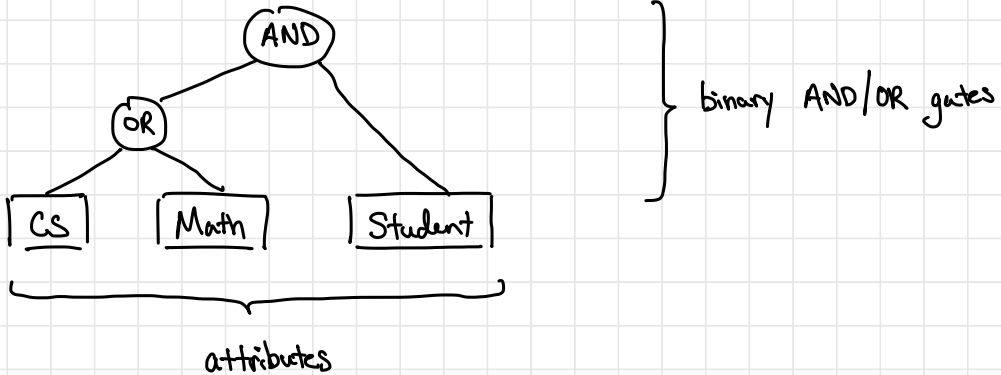
ABE: can encrypt to general policies

Key-policy ABE: ciphertexts are associated with attributes
secret keys are associated with policy

Example: attribute could be a classification level (unclassified, secret, top secret)
policy could be (unclassified or secret)

attribute could describe a role (e.g., CS, math, physics, etc.)
policy could be (CS or math)

We will describe policies as a Boolean formula:



In a formula, every gate has fan-out 1 (each output of a gate can only be used once)

Will assume a polynomial number of attributes today

We will label the attributes by integers $1, 2, \dots, n$

Goyal-Pandey-Sahai-Waters ABE scheme:

Setup: For each $i \in [n]$, sample $t_i \xleftarrow{R} \mathbb{Z}_p$
Sample $\gamma \xleftarrow{R} \mathbb{Z}_p$

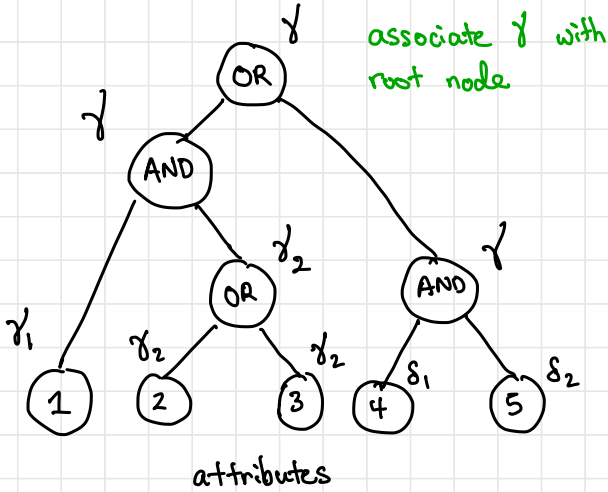
Output attribute keys $T_1 = g^{t_1}, \dots, T_n = g^{t_n}, h = e(g, g)^\gamma$

$\text{mpk} = (T_1, \dots, T_n, h)$ $\text{msk} = (t_1, \dots, t_n, \gamma)$


Encrypt (mpk, S, m): On input $\text{mpk} = (T_1, \dots, T_n, h)$ and a set of attributes $S \subseteq [n]$, and the message $m \in \mathbb{G}_r$, sample $r \xleftarrow{R} \mathbb{Z}_p$ and output

$$\text{ct} = (S, T_i^r, h^r \cdot m)$$

KeyGen (msk, P): We will start with an example, and formalize later.
Idea is we will secret share the "secret key" γ according to the policy P .



for AND gate




sample $\gamma_1, \gamma_2 \stackrel{R}{\leftarrow} \mathbb{Z}_p$ such that $\gamma = \gamma_1 + \gamma_2$

need both γ_1 and γ_2 to obtain γ

$$\gamma_1, \gamma_2 \stackrel{R}{\leftarrow} \mathbb{Z}_p \text{ s.t. } \gamma_1 + \gamma_2 = \gamma$$

$$\delta_1, \delta_2 \stackrel{R}{\leftarrow} \mathbb{Z}_p \text{ s.t. } \delta_1 + \delta_2 = \gamma$$

for OR gate



$$\text{set } \gamma_1 = \gamma_2 = \gamma$$

either γ_1 or γ_2 needed to obtain γ

We refer to the exponents associated with the leaf nodes as the shares of the master secret key associated with the attribute.


Key for policy P:

1. Secret share γ according to policy P
2. Let $\gamma_i \in \mathbb{Z}_p$ be the share associated with attribute i
3. The secret key for P is the set

$g^{\gamma_i/t}$ for each attribute i that appears in the policy

along with the policy P

Decryption (we will defer formal description to later)



$$\gamma_1 + \gamma_2 = \gamma$$

$$\delta_1 + \delta_2 = \gamma$$

write each as K_i for each attribute appearing in the policy
[attribute keys]

ciphertext : $S, \{T_i^r\}_{res}, h \cdot m$

$$\parallel \quad \parallel$$

$$g^{rt_i} \quad e(g, g)^{r\gamma} \cdot m$$

Suppose ciphertext has attributes (2, 4, 5)

ciphertext contains $g^{rt_2}, g^{rt_4}, g^{rt_5}$

user can pair $e(K_2, g^{rt_2}) = e(g^{\gamma_2/t_2}, g^{rt_2}) = e(g, g)^{\gamma_2 r}$

can similarly obtain $e(g, g)^{\delta_1 r}$ and $e(g, g)^{\delta_2 r}$

↳ Observe that exponents $\gamma_2 r, \delta_1 r, \delta_2 r$ are secret shares of γr . If the attributes satisfy the policy, can use shares to reconstruct $e(g, g)^{\gamma r}$ and decrypt. This relies on fact that secret sharing scheme has a linear reconstruction algorithm.

Summary of approach: ElGamal-style encryption

secret key: γ
public key: $h = e(g, g)^\gamma$

ciphertext includes $h^r \cdot m$

Normally in ElGamal, ciphertext contains g^r and γ is used to decrypt.
With pairings, we can also use g^γ as the decryption key.
↳ But cannot just give out g^γ .

Instead: secret key contain secret share of g^γ according to policy P , but need a way to combine with attributes.

Approach: set $K_i = g^{d_i/t_i}$ where d_i is a share of γ and t_i is the attribute key.


Pair with $g^{t_i r}$ from ciphertext to obtain share of $e(g, g)^\gamma$.

Secret shares for different keys are independent (no mixing and matching)


To describe the scheme more generally, we describe a linear secret sharing scheme

Suppose we have a secret s . For each attribute t , we can associate with it a share s_t . Given a policy P and shares $\{s_t\}_{t \in T}$, if $P(T) = 1$, then should be able to recover s .

Idea: secret share gate-by-gate:




recover s if you have either s_1 or s_2
 \hookrightarrow set $s_1 = s_2 = s$



recover s if you have s_1 and s_2
 \hookrightarrow sample $s_1 \xleftarrow{R} \mathbb{F}_p$ and set $s_2 = s - s_1$
 \hookrightarrow reconstruct by computing $s_1 + s_2 = s$
 \hookrightarrow no information given out with just 1 share

Compose to support general policies



$$s_1, s_3 \xleftarrow{R} \mathbb{Z}_p$$

$$\begin{aligned} s_A &= s_1 \\ s_B &= s_C = s - s_1 \\ s_D &= s_3 \\ s_E &= s - s_3 \end{aligned}$$

For the policy, we can associate with it a "share generating" matrix

For an AND gate, we can secret share s by sampling $\alpha \leftarrow^R \mathbb{F}_p$ and setting $s_1 = s + \alpha$ and $s_2 = -\alpha$:

$$\begin{bmatrix} s_1 \\ s_2 \end{bmatrix} = \begin{bmatrix} 1 & 1 \\ 0 & -1 \end{bmatrix} \begin{bmatrix} s \\ \alpha \end{bmatrix}$$

For an OR gate, we secret share s by setting $s_1 = s_2 = s$

$$\begin{bmatrix} s_1 \\ s_2 \end{bmatrix} = \begin{bmatrix} 1 \\ 1 \end{bmatrix} [s]$$

General procedure:

1. Associate vector $[1]$ with root node and initialize counter $c \leftarrow 1$

2. Proceed down the tree.

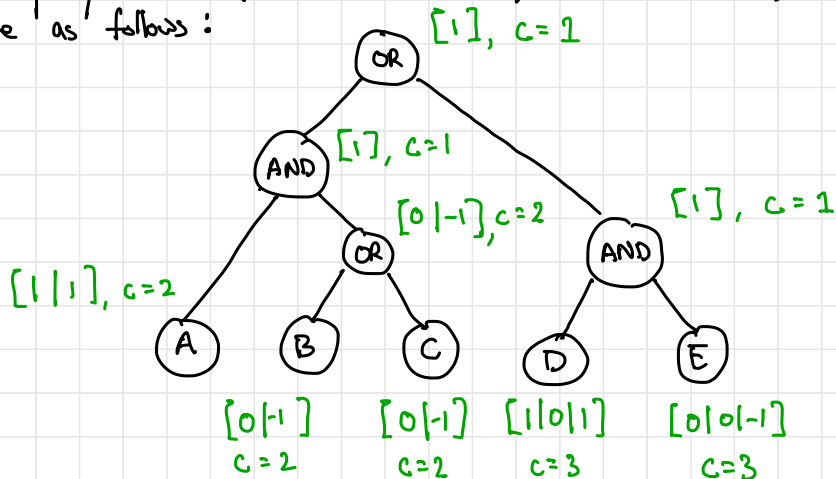
- If a node is an OR with label v , the children are associated with label v (i.e., they are identical to the parent).

- If a node is an AND with label v , we label the children with label $v \parallel 0^{\lfloor \log(v)-c \rfloor} \parallel 1$ and the other label with $0^c \parallel -1$.

Observe that the sum of the shares is $v \parallel 0^{\lfloor \log(v)-c \rfloor}$. Finally, increment $c \leftarrow c + 1$.

3. Pad all vectors with 0s to dimension c

Example: For policy above $(A \text{ AND } (B \text{ OR } C)) \text{ OR } (D \text{ AND } E)$, the vectors are as follows:



Share-generation matrix:

$$\begin{array}{l} SA \\ SB \\ SC \\ SD \\ SE \end{array} \underbrace{\begin{bmatrix} 1 & 1 & 0 \\ 0 & -1 & 0 \\ 0 & -1 & 0 \\ 1 & 0 & 1 \\ 0 & 0 & -1 \end{bmatrix}}_{M_p} \begin{bmatrix} S \\ \alpha_1 \\ \alpha_2 \end{bmatrix} = \begin{bmatrix} S + \alpha_1 \\ -\alpha_1 \\ -\alpha_1 \\ S + \alpha_2 \\ -\alpha_2 \end{bmatrix}$$

same shares we had before!

Main observation: if attributes satisfy policy, then vector $e_i^T = [1, 0, \dots, 0]$ will be in the row span of M_p

can show inductively: for every node v in the tree, if attributes satisfy the node v , then vector associated with v is in the row span of M_p

true for every leaf node and inductive step follows by construction

Converse also holds: if attributes do not satisfy policy, then e_i^T is not in the row span of M_p

Can express Boolean formula access structure as a linear secret sharing scheme.

Boolean formula policy P on k attributes

→ share generation matrix M with k rows

If a set of attributes $S \subseteq [k]$ satisfies the policy, then there exists a vector $v \in \mathbb{Z}_p^k$ where $v_i = 0$ for all $i \notin S$ such that

$$v^T S = [1 \ 0 \ \dots \ 0]$$

(i.e., the first basis vector is in spanned by the rows associated with the attributes in S).

If S does not satisfy the policy, then $[1 \ 0 \ \dots \ 0]$ is not spanned by the rows of S .

↳ Suppose v_1, \dots, v_n are the vectors associated with S

Since $e_1 \notin \text{span}(v_1, \dots, v_n)$, it must be the case that e_1 has a component in the orthogonal complement of $\{v_1, \dots, v_n\}$

Thus, there exists a vector w such that $w^T v_i = 0$ and $w^T e_1 \neq 0$. Without loss of generality, we assume the first component of w is 1.

attribute universe n

Setup (n): Sample $t_1, \dots, t_n \xleftarrow{R} \mathbb{Z}_p$
 $\gamma \xleftarrow{R} \mathbb{Z}_p$

Set $T_i = g^{t_i}$
 $h = e(g, g)^\gamma$

Set $mpk = (T_1, \dots, T_n, h)$

Set $msk = (t_1, \dots, t_n, \gamma)$

Encrypt (mpk, S, m): sample $r \xleftarrow{R} \mathbb{Z}_p$ and set $ct = (\{c_i, T_i^r\}_{i \in S}, h^{r \cdot m})$

KeyGen (msk, M): here $M \in \mathbb{Z}_p^{k \times l}$ is the share-generating matrix associated with the policy

sample $\alpha_1, \dots, \alpha_{l-1} \xleftarrow{R} \mathbb{Z}_p$ and compute the shares of γ :

$$\begin{bmatrix} s_1 \\ \vdots \\ s_k \end{bmatrix} = M \cdot \begin{bmatrix} \gamma \\ \alpha_1 \\ \alpha_2 \\ \vdots \\ \alpha_{l-1} \end{bmatrix}$$

$p(1), \dots, p(l)$ are the attributes the policy is looking at

let $p: [k] \rightarrow [n]$ be the mapping from the row index of M to the attribute index in $[n]$.

output $sk = M, \{c_i, g^{s_i/t_{p(i)}}\}_{i \in [k]}$

Decrypt (sk, ct): Parse $sk = (M, \{c_i, D_i\}_{i \in [k]})$

$ct = (\{c_i, C_i\}_{i \in S}, Z)$

If S satisfies the policy, then there exists a vector $v \in \mathbb{Z}_p^k$ such that $v^T M = [1, 0, \dots, 0]$ and $v_i = 0$ for all $p(i) \notin S$

Compute $Z / \prod_{i \in [k]} e(C_{p(i)}, D_i)^{v_i}$

[Note: treat $C_{p(i)} = g^0$ if $p(i) \notin S$]

Reduction strategy:

Algorithm A chooses a set of attributes $S \subseteq [n]$

We need to simulate the public key mpk and secret keys for policies that S does not satisfy

Approach: Challenge ciphertext will use T
 $\leftarrow e(g, g)^{\gamma r}$

If $T = e(g, g)^{abc}$, we will need to map abc to γr
and also publish $e(g, g)^{\gamma r}$ in public key

Reduction will use c for the exponent r and ab for the exponent γ . Then,

$$\begin{aligned} e(g, g)^{\gamma} &= e(g, g)^{ab} = e(g^a, g^b) \\ &= e(A, B) \end{aligned}$$

For the attribute-specific components, if $i \in S$, pick $t_i \xleftarrow{R} \mathbb{Z}_p$
and set $T_i = g^{t_i}$ \leftarrow must know t_i

If $i \notin S$, will need to pick carefully

to simulate ciphertext components

In the ciphertext, reduction can set $C_i = C^{t_i} = g^{t_i c}$

Suffices to construct secret keys. Problem: do not (and cannot know secret key ab)

Let $M \in \mathbb{Z}_p^{k \times l}$ be share-generation matrix associated with a policy.
Let $\rho: [k] \rightarrow [n]$ be labeling function.

Normally, we would compute

$$\begin{bmatrix} s_1 \\ \vdots \\ s_k \end{bmatrix} = M \cdot \begin{bmatrix} ab \\ \alpha_1 \\ \vdots \\ \alpha_l \end{bmatrix} = \begin{bmatrix} \text{---} m_1^T \text{---} \\ \vdots \\ \text{---} m_k^T \text{---} \end{bmatrix} \begin{bmatrix} ab \\ \alpha_1 \\ \vdots \\ \alpha_l \end{bmatrix}$$

and give out $D_i = g^{s_i/t_{p(i)}}$ for each $i \in [k]$

Problem: Each s_i is a linear function of ab and we do not have g^{ab}

Let $u = \begin{bmatrix} ab \\ \alpha_1 \\ \vdots \\ \alpha_l \end{bmatrix}$. Normally we compute $s = Mu$.

We will instead write $u = b \underbrace{\begin{bmatrix} 0 \\ \alpha_1 \\ \vdots \\ \alpha_l \end{bmatrix}}_z + ab \cdot w$

z : share of 0

↑ recall first component of w is 1 and w is orthogonal to rows of M corresponding to attributes in S
[$m_i^T w = 0$ if $p(i) \in S$]

Same distribution as before!

Consider two possibilities for D_i :

1) if $p(i) \in S$: $s_i = m_i^T u = b m_i^T z$

$$D_i = g^{s_i/t_{p(i)}} = g^{b m_i^T z / t_{p(i)}} = \underline{B^{m_i^T z / t_{p(i)}}}$$

reduction can compute since it knows M , z and $t_{p(i)}$

2) if $p(i) \notin S$: $s_i = m_i^T u = b m_i^T z + ab m_i^T w$
 $= b(m_i^T z + a m_i^T w)$

$$D_i = g^{s_i/t_{p(i)}} = g^{b(m_i^T z + a m_i^T w) / t_{p(i)}}$$

Seems to rely on knowledge of ab .

But... we have one degree of freedom. We can pick $t_{p(i)}$ strategically here!

Reduction does not need to know $t_{p(i)}$ since these are not present in challenge ciphertext. It just needs to give out $T_{p(i)} = g^{t_{p(i)}}$ in public parameters.

Approach: choose T_i for $i \notin S$ to force cancellation.

Set $T_i := B^{d_i}$ where $d_i \xleftarrow{R} \mathbb{Z}_p$ in public parameters.
||
 g^{bd_i}

Then, in the secret key for $g(i) \notin S$

$$\begin{aligned} D_i &= g^{s_i / t_{p(i)}} = g^{b(m_i^T z + a m_i^T w) / t_{p(i)}} \\ &= g^{m_i^T z + a m_i^T w / d_i} \\ &= g^{m_i^T z / d_i} A^{m_i^T w / d_i} \end{aligned}$$

reduction can compute since it knows M, z, w , and d_i