

Verifiable Attribute-based Keyword Search over Outsourced Encrypted Data

Qingji Zheng[†] Shouhuai Xu[†] Giuseppe Ateniese[‡]

[†] University of Texas at San Antonio, USA

[‡] Sapienza University of Rome, Italy and Johns Hopkins University, USA

Abstract—It is quite common nowadays for data owners to outsource their data to the cloud. However, since the cloud is not fully trusted, the outsourced data should be encrypted, which brings a range of problems, such as: How can authorized data users search over a data owner’s outsourced encrypted data? How should a data owner grant search capabilities to data users? How can data users be assured that the cloud faithfully executed the search operations? Towards ultimately addressing these problems, in this paper we propose a novel cryptographic scheme, called *verifiable attribute-based keyword search* (VABKS). This scheme allows a data user, whose attributes or credentials satisfy a data owner’s access control policy, to (i) search over the data owner’s outsourced encrypted data, (ii) outsource the tedious search operations to the cloud, and (iii) verify whether the cloud has faithfully executed the search operations. We define VABKS’s security properties, and present concrete constructions that are proven to possess these properties. Performance evaluation shows that the proposed schemes are practical.

I. INTRODUCTION

Cloud computing allows data owners to use massive data storage and vast computation capability at low price. Despite the benefits, data outsourcing deprives data owners of direct control over their outsourced data. To alleviate such concerns, it is natural for data owners to encrypt their outsourced data. However, encrypting outsourced data brings a range of problems, especially when the outsourced data needs to be shared with others. In this paper, we address some of the problems that are encountered when the data owner needs to enforce access control over its outsourced encrypted data, while the data user needs to outsource the search operations to the cloud for performance reasons. To the best of our knowledge, existing solutions are not sufficient for the goals we aim to achieve.

A. Our Contributions

We propose a novel cryptographic scheme, which allows a data user with proper attributes (according to a data owner’s access control policy) to (i) search over the data owner’s outsourced encrypted data, (ii) outsource the actual search operations to the cloud, and (iii) verify whether the cloud has faithfully executed the search operations. This is achieved through a novel cryptographic scheme, called *verifiable attribute-based keyword search* (VABKS). This scheme allows data owners to control the search and use of its outsourced encrypted data according to its access control policy, while the legitimate data users can outsource the often costly search operations to the

cloud and can verify whether the cloud has faithfully executed the search operations.

In addition to formally define the security properties of VABKS, we propose the design of VABKS schemes in a modular fashion, by using attribute-based encryption, bloom filter, digital signature, and a newly introduced building-block called attribute-based keyword search (ABKS), which may be of independent value. The novel application of bloom filter and digital signature allows data users to *efficiently* verify that the cloud honestly conducted search operations over outsourced encrypted data. Experimental evaluation shows that the VABKS (and ABKS) schemes are practical.

B. Related Work

To the best of our knowledge, there is no solution is sufficient for what we want to achieve. Nevertheless, the relevant techniques are briefly reviewed as follows.

Keyword Search over Encrypted Data. Existing solutions for keyword-based search over encrypted data can be classified into two categories: searchable encryption in the symmetric-key setting (e.g., [1]–[9]) and searchable encryption in the public-key setting (e.g., [10]–[14]). In these solutions, the data and the associated keywords are encrypted with symmetric keys or public keys before being outsourced to the server/cloud. The data owner generates some tokens that can be used by data users to search over the encrypted data. Several variants (e.g., [15]–[18]) have been proposed to support complex queries. The work [4], [19] considered searchable encryption in the multi-users setting, where the data owner can enforce access control policies by distributing (stateful) secret keys to authorized users directly and revoke them if necessary. However, these solutions do not solve the problem VABKS aims to solve because (1) part of solutions require interactions between data users and data owners (or trusted authorities, e.g., trapdoor generation entity [14]) to grant search capabilities, and (2) they assume that the server can faithfully execute search operations. Compared with these solutions, VABKS allows data users with proper attributes to issue search tokens so that the cloud can perform keyword search on their behalf, *without* interacting with data owners, and can assure that the cloud honestly executes search operations even if it is untrusted. Note that while ABKS can be cast into the framework of predicate encryption [20], [21], VABKS provides the novel feature of verifiability.

Attribute-Based Encryption (ABE). ABE allows entities with proper credentials to decrypt a ciphertext in question [22]. It has two variants depending on how access control is enforced. Specifically, key-policy ABE (KP-ABE), where the decryption key is associated with an access control policy [23], and ciphertext-policy ABE (CP-ABE), where the ciphertext is associated with an access control policy [24]. ABE has become the de facto standard for enforcing access control policies via cryptographic means, and has been enriched with various features (e.g., [25]–[28]). In this paper, we propose extending ABE with the novel feature of keyword search, namely *attribute-based keyword search*, by which keywords are encrypted together with some access control policies so that only data users with the right credentials can generate proper tokens for searching on encrypted data. This effectively prevents data owners from knowing keywords a data consumer is searching for and requires no interaction between data users and data owners/trusted authorities. Note that although [14] allows data users to obtain search token from data owners/trusted authorities without revealing keywords, data owners/trusted authorities must always be online and available.

Verifiable Keyword Search. Recently, verifiable keyword search solutions have been proposed in [29]–[31], where each keyword is represented as a root of a polynomial. It is possible to check whether keywords are in the data set by evaluating the polynomial on them and verifying that its output is zero. However, these approaches work only if keywords are sent in the clear to the cloud, thus they are not suitable for the cloud data sharing scenario. Verifiable keyword search in the symmetric-key setting was investigated in [32], which is *not* secure in the public-key setting because the attacker can easily infer keywords directly from the small alphabet space via keyword guessing attack.

Paper Organization: Section II reviews some cryptographic preliminaries. Section III defines ABKS and its security properties, presents the KP-ABKS and CP-ABKS schemes and analyzes their security. Section IV defines VABKS and its security properties, presents a VABKS scheme and analyzes its security. Section V evaluates the performance of the ABKS and VABKS schemes. Section VI concludes the paper.

II. PRELIMINARIES

Given a set S , let $a \leftarrow S$ denote selecting an element a from S uniformly at random. Let $U = \{\text{at}_1, \dots, \text{at}_n\}$ be the set of all relevant attributes and \parallel denote the concatenation.

A. Cryptographic Assumption

Let p be an ℓ -bit prime and G, G_T be cyclic groups of prime order p with generators g, g_T respectively. Let e be a bilinear map: $e : G \times G \rightarrow G_T$ satisfying: (i) $\forall a, b \leftarrow \mathbb{Z}_p$, $e(g^a, g^b) = e(g, g)^{ab}$, (ii) $e(g, g) \neq 1$, and (iii) e can be computed efficiently.

Decisional Linear Assumption (DL). Given $(g, f, h, f^{r_1}, g^{r_2}, Q)$ where $g, f, h, Q \leftarrow G$, $r_1, r_2 \leftarrow \mathbb{Z}_p$, this assumption says that any probabilistic polynomial-time algorithm \mathcal{A} determines $Q \stackrel{?}{=} h^{r_1+r_2}$ with a negligible advantage in security

parameter ℓ , where the advantage is defined as

$$|\Pr[\mathcal{A}(g, f, h, f^{r_1}, g^{r_2}, h^{r_1+r_2}) = 1] - \Pr[\mathcal{A}(g, f, h, f^{r_1}, g^{r_2}, Q) = 1]|.$$

Generic Bilinear Group [33]. The generic bilinear map is described as follows: Let ψ_0, ψ_1 be two random encodings of the additive group \mathbb{Z}_p^+ , such that ψ_0, ψ_1 are injective maps from \mathbb{Z}_p^+ to $\{0, 1\}^m$, where $m > 3 \log(p)$. Let $G = \{\psi_0(x) | x \in \mathbb{Z}_p\}$ and $G_T = \{\psi_1(x) | x \in \mathbb{Z}_p\}$. In addition, there is another oracle to compute $e : G \times G \rightarrow G_T$. G is referred to as a generic bilinear group. Let g denote $\psi_0(1)$, g^x denote $\psi_0(x)$, $e(g, g)$ denote $\psi_1(1)$, and $e(g, g)^y$ denote $\psi_1(y)$.

Pseudorandom Generator [34]. A pseudorandom generator $H : \{0, 1\}^\ell \rightarrow \{0, 1\}^m$, $\ell < m$ is a deterministic algorithm that takes as input a short random seed and generates a long string that is pseudorandom.

B. Bloom Filter for Membership Query

A Bloom filter [35] is a data structure suitable for representing a set of items while allowing membership queries. Initially, a m -bit Bloom filter is an array of m bits, which are set to 0, and uses k independent hash functions (i.e. universal hash) H'_1, \dots, H'_k of range $0, \dots, m - 1$. For each element $w \in S = \{w_1, \dots, w_n\}$, the bits $H'_j(w)$ are set to 1 for $1 \leq j \leq k$. The locations in the m -bit Bloom filter can be set to 1 more than once, but only the first change has an effect. To check whether w is an element of S , it is enough to verify that all $H'_j(w)$, for $1 \leq j \leq k$, are set to 1. If not, w is certainly not an element of S ; otherwise, w might be a member of S with high probability (a false-positive rate should be considered). In the m -bit Bloom filter, the false positive rate can be calculated given the assumption that hash functions are perfectly random: After n elements are hashed into the m -bit Bloom filter, the false positive rate is $(1 - (1 - \frac{1}{m})^{kn})^k \approx (1 - e^{-kn/m})^k$. We can calculate the optimal number of hash functions minimizing the false positive rate, which is $k = (\ln 2)m/n$. Thus, the false positive rate in this case is $(0.6185)^{m/n}$. A m -bit Bloom filter has two associated algorithms:

- $\text{BF} \leftarrow \text{BFGen}(\{H'_1, \dots, H'_k\}, \{w_1, \dots, w_n\})$: This algorithm generates a m -bit Bloom filter by hashing the data set $\{w_1, \dots, w_n\}$ with $\{H'_1, \dots, H'_k\}$.
- $\{0, 1\} \leftarrow \text{BFVerify}(\{H'_1, \dots, H'_k\}, \text{BF}, w)$: This algorithm returns 0 if w is not an element of the data set, and 1 otherwise.

C. Access Trees for Representing Access Control Policies

Access trees are convenient for representing access control policies [23]. In an access tree, a leaf is associated with an attribute and an inner node represents a threshold gate (of which AND and OR gates are special cases). Let num_v be the number of children of the node v , and label the ordered children as 1 to num_v . Let k_v be the threshold value associated with an inner node v where $1 \leq k_v \leq \text{num}_v$. Let $\text{parent}(v)$ denote the parent of the node v , $\text{ind}(v)$ denote the label associated with node v , $\text{att}(v)$ denote the attribute associated with leaf node v , $\text{lvs}(T)$ denote the set of leaves of the access

tree T , and T_v denote the subtree of T rooted at the node v (thus, $T_{\text{root}} = T$).

Given an attribute set $\text{Atts} \subseteq U$, we use a function F such that $F(\text{Atts}, T_v) = 1$ means Atts satisfies the access control policy represented by the subtree T_v . The evaluation of $F(\text{Atts}, T_v)$ can be performed iteratively as follows:

- In the case v is a leaf: If $\text{att}(v) \in \text{Atts}$, set $F(\text{Atts}, T_v) = 1$; otherwise, set $F(\text{Atts}, T_v) = 0$.
- In the case v is an inner node with children $v_1, \dots, v_{\text{num}_v}$: If there exists a subset $I \subseteq \{1, \dots, \text{num}_v\}$ such that $|I| \geq k_v$ and $\forall j \in I, F(\text{Atts}, T_{v_j}) = 1$, set $F(\text{Atts}, T_v) = 1$; otherwise, set $F(\text{Atts}, T_v) = 0$.

Given an access tree T , we denote the algorithm for distributing a secret s according to the access tree T as:

$$\{q_v(0) \mid v \in \text{lvs}(T)\} \leftarrow \text{Share}(T, s).$$

The algorithm operates as follows. Generate a polynomial q_v for each node v in T with the respective threshold value k_v in a top-down fashion according to T (for each leaf node, its threshold value is naturally 1).

- If v is the root of T ($v = \text{root}$), set $q_v(0) = s$ and randomly pick $k_v - 1$ coefficients for polynomial q_v .
- If v is a leaf of T , set $q_v(0) = q_{\text{parent}(v)}(\text{ind}(v))$.
- If v is an inner node, set $q_v(0) = q_{\text{parent}(v)}(\text{ind}(v))$ and randomly select $k_v - 1$ coefficients for polynomial q_v .

When the algorithm stops, each leaf v is associated with a value $q_v(0)$, which is the secret share of s .

Given an access tree T and a set of values $\{E_{u_1}, \dots, E_{u_m}\}$, where u_1, \dots, u_m are leaves of T such that $F(\{\text{att}(u_1), \dots, \text{att}(u_m)\}, T) = 1$, $E_{u_j} = e(g, h)^{q_{u_j}(0)}$ for $1 \leq j \leq m, g, h \in G$, e is a bilinear map and $q_{u_1}(0), \dots, q_{u_m}(0)$ are secret shares of s according to T , we denote the algorithm for reconstructing $e(g, h)^s$ as

$$e(g, h)^s \leftarrow \text{Combine}(T, \{E_{u_1}, \dots, E_{u_m}\}).$$

The algorithm operates as follows. Execute the following steps with respect to node v in a bottom-top fashion according to T .

- If $F(\{\text{att}(u_1), \dots, \text{att}(u_m)\}, T_v) = 0$, then continue.
- If v is a leaf and $F(\{\text{att}(u_1), \dots, \text{att}(u_m)\}, T_v) = 1$, then $E_v = E_{u_j}(0) = e(g, h)^{q_{u_j}(0)}$ where $v = u_j$ for some j .
- Otherwise, v is an inner node and $F(\{\text{att}(u_1), \dots, \text{att}(u_m)\}, T_v) = 1$. Then for v 's children nodes $\{v_1, \dots, v_{\text{num}_v}\}$, there should exist a set of indices S , such that $|S| = k_v$, such that $j \in S, F(\{\text{att}(u_1), \dots, \text{att}(u_m)\}, T_{v_j}) = 1$. Let $\Delta_{v_j} = \prod_{l \in S, l \neq j} \frac{-j}{l-j}$ so that

$$E_v = \prod_{j \in S} E_{v_j}^{\Delta_{v_j}} = \prod_{j \in S} (e(g, h)^{q_{v_j}(0)})^{\Delta_{v_j}} = e(g, h)^{q_v(0)}$$

When the algorithm stops, the root of T is with a value $E_{\text{root}} = e(g, h)^{q_{\text{root}}(0)} = e(g, h)^s$.

III. ATTRIBUTE-BASED KEYWORD SEARCH

In order to effectively grant search capability over encrypted keywords, it is expected that data owners can specify access control policies when they encrypt keywords. After that it requires no further interaction between data owners and data users, which is similar to the intuition of attribute-based encryption. This inspire us to devise a new cryptographic scheme, called Attribute-Based Keyword Search (ABKS). It allows data owners to bind keyword ciphertexts with access control policies, so that authorized data users, whose attributes (or credentials) satisfy access control policies, can conduct keyword search (i.e. generating proper tokens) over keyword ciphertexts.

To ease the presentation, we use I_{Enc} to denote the input to the encryption function Enc , and I_{KeyGen} to denote the input to the key generation function KeyGen . In the case of ciphertext-policy ABKS (CP-ABKS), I_{Enc} and I_{KeyGen} are an access structure and an attribute set, respectively. In the case of key-policy ABKS (KP-ABKS), I_{Enc} and I_{KeyGen} are an attribute set and an access structure, respectively. Let $F(I_{\text{KeyGen}}, I_{\text{Enc}}) = 1$ denote that I_{KeyGen} satisfies I_{Enc} in CP-ABKS or I_{Enc} satisfies I_{KeyGen} in KP-ABKS.

A. ABKS Definition and Security

Definition 1: An ABKS scheme consists of the following algorithms:

- $(\text{mk}, \text{pm}) \leftarrow \text{Setup}(1^\ell)$: This algorithm initializes the public parameter pm and generates a master key mk .
- $\text{sk} \leftarrow \text{KeyGen}(\text{mk}, I_{\text{KeyGen}})$: Taking as input the master key mk and I_{KeyGen} , this algorithm outputs credential sk .
- $\text{cph} \leftarrow \text{Enc}(w, I_{\text{Enc}})$: Taking as input keyword w and I_{Enc} , this algorithm outputs a ciphertext cph .
- $\text{tk} \leftarrow \text{TokenGen}(\text{sk}, w)$: Taking as input credential sk and keyword w , this algorithm outputs a search token tk .
- $\{0, 1\} \leftarrow \text{Search}(\text{cph}, \text{tk})$: This algorithm determines whether cph contains the keyword w specified by tk . It returns 1 if so, and 0 otherwise.

The correctness of the ABKS scheme requires that, given $(\text{mk}, \text{pm}) \leftarrow \text{Setup}(1^\ell)$, $\text{sk} \leftarrow \text{KeyGen}(\text{mk}, I_{\text{KeyGen}})$ and $F(I_{\text{KeyGen}}, I_{\text{Enc}}) = 1$, for any w , $\text{cph} \leftarrow \text{Enc}(w, I_{\text{Enc}})$ and $\text{tk} \leftarrow \text{TokenGen}(\text{sk}, w)$, then $1 \leftarrow \text{Search}(\text{cph}, \text{tk})$ always holds.

Similar to the searchable encryption in the literature, the intuition for ABKS security is that the adversary should learn nothing but search results when executing function Search . Concretely, an ABKS scheme should satisfy security requirements as follows even in the presence of a probabilistic polynomial-time adversary \mathcal{A} :

- Selective security against chosen-keyword attack: Given only keyword ciphertexts (i.e. without seeing any matched search tokens), no partial information about keywords is leaked in the selective model (meaning that I_{Enc} for encrypting challenge keywords has to be selected before the public parameter is generated). To be specific,

\mathcal{A} cannot distinguish the encryption of two challenge keywords of its choice. We formalize this security property via the selectively chosen-keyword attack game.

- **Keyword secrecy:** In the public-key setting, it is inherently impossible to protect any information encoded by search tokens (aka. predicate privacy [36]) due to the *keyword guessing attack*: one can encrypt any keyword of his choice and check whether the resulting ciphertext matches the target token. Therefore, a weaker notion, called *keyword secrecy*, is introduced to assure that the probability of \mathcal{A} learning the keyword from the ciphertext and its search token is no more than that of a random keyword guess. We formalize this security property via the keyword secrecy game.

Selectively Chosen-Keyword Attack (SCKA) Game:

Setup: \mathcal{A} selects a non-trivial challenge I_{Enc}^* (a trivial challenge I_{Enc}^* is one that can be satisfied by any data user who does not have any credential), and gives it to the challenger. Then the challenger runs $\text{Setup}(1^\ell)$ to generate the public parameter pm and the master key mk .

Phase 1: \mathcal{A} can query the following oracles for polynomially many times, and the challenger keeps a keyword list L_{kw} , which is initially empty.

- $\mathcal{O}_{\text{KeyGen}}(I_{\text{KeyGen}})$: If $F(I_{\text{KeyGen}}, I_{\text{Enc}}^*) = 1$, then abort; otherwise, the challenger returns to \mathcal{A} credential sk corresponding to I_{KeyGen} .
- $\mathcal{O}_{\text{TokenGen}}(I_{\text{KeyGen}}, w)$: The challenger generates credential sk with I_{KeyGen} , and returns to \mathcal{A} a search token tk by running algorithm TokenGen with inputs sk and w . If $F(I_{\text{KeyGen}}, I_{\text{Enc}}^*) = 1$, the challenger adds w to L_{kw} .

Challenge phase: \mathcal{A} chooses two keywords w_0 and w_1 , where $w_0, w_1 \notin L_{kw}$. The challenger selects $\lambda \leftarrow \{0, 1\}$, computes $\text{cph}^* \leftarrow \text{Enc}(w_\lambda, I_{\text{Enc}}^*)$, and delivers cph^* to \mathcal{A} . Note that the requirement of $w_0, w_1 \notin L_{kw}$ is to prevent \mathcal{A} from trivially guessing λ with tokens from $\mathcal{O}_{\text{TokenGen}}$.

Phase 2: \mathcal{A} continues to query the oracles as in Phase 1. The restriction is that (I_{KeyGen}, w_0) and (I_{KeyGen}, w_1) cannot be the input to $\mathcal{O}_{\text{TokenGen}}$ if $F(I_{\text{KeyGen}}, I_{\text{Enc}}^*) = 1$.

Guess: \mathcal{A} outputs a bit λ' , and wins the game if $\lambda' = \lambda$.

Let $|\Pr[\lambda = \lambda'] - \frac{1}{2}|$ be the advantage of \mathcal{A} winning the above SCKA game. Thus, we have

Definition 2: An ABKS scheme is *selectively secure against chosen-keyword attack* if the advantage of any \mathcal{A} winning the SCKA game is negligible in security parameter ℓ .

Keyword Secrecy Game:

Setup: The challenger runs $\text{Setup}(1^\ell)$ to generate the public parameter pm and the master key mk .

Phase 1: \mathcal{A} can query the following oracles for polynomially many times:

- $\mathcal{O}_{\text{KeyGen}}(I_{\text{KeyGen}})$: The challenger returns to \mathcal{A} credential sk corresponding to I_{KeyGen} . It adds I_{KeyGen} to the list L_{KeyGen} , which is initially empty.
- $\mathcal{O}_{\text{TokenGen}}(I_{\text{KeyGen}}, w)$: The challenger generates credential sk with I_{KeyGen} , and returns to \mathcal{A} a search token tk by running algorithm TokenGen with input sk and w .

Challenge phase: \mathcal{A} chooses a non-trivial I_{Enc}^* and gives it to the challenger. The challenger selects w^* from the message space uniformly at random and selects I_{KeyGen}^* such that $F(I_{\text{KeyGen}}^*, I_{\text{Enc}}^*) = 1$. The challenger runs $\text{cph} \leftarrow \text{Enc}(w^*, I_{\text{Enc}}^*)$ and $\text{tk} \leftarrow \text{TokenGen}(\text{sk}, w^*)$ and delivers (cph, tk) to \mathcal{A} . We require that $\forall I_{\text{KeyGen}} \in L_{\text{KeyGen}}, F(I_{\text{KeyGen}}, I_{\text{Enc}}^*) = 0$.

Guess: After guessing q distinct keywords, \mathcal{A} outputs a keyword w' , and wins the game if $w' = w$.

Definition 3: An ABKS scheme achieves *keyword secrecy* if the probability that \mathcal{A} wins the keyword secrecy game is at most $\frac{1}{|\mathcal{M}| - q} + \epsilon$, where \mathcal{M} is the keyword space, q is the number of distinct keywords that the adversary has attempted, and ϵ is a negligible in security parameter ℓ .

B. ABKS Construction

Assume that $H_1 : \{0, 1\}^* \rightarrow G$ is a secure hash function that modeled as random oracle, and $H_2 : \{0, 1\}^* \rightarrow \mathbb{Z}_p$ is a secure one-way hash function, where G is the cyclic group of bilinear map $e : G \times G \rightarrow G_T$.

In ABKS it should consider (1) how to encrypt a keyword, and (2) how to generate a search token from credentials and a keyword, so that if data users' attributes (or credentials) satisfy the specified access control policy, then it can determine whether keywords are equal with respect to the ciphertext and search token.

By taking KP-ABKS as an example, where credentials are generated by letting $t \leftarrow \mathbb{Z}_p$, $A_v = g^{q_v(0)} H_1(\text{att}(v))^t$, $B_v = g^t$ for each leaf v , where g is the generator of G , $q_v(0)$ is the share of secret ac for leaf v according to the access tree T , the ciphertext and search token are generated as follows:

- Keyword w is encrypted into two parts: one is to blind keyword with randomness by letting $W' = g^{cr_1}$, $W = g^{a(r_1+r_2)} g^{bH_2(w)r_1}$ and $W_0 = g^{r_2}$ where $g^a, g^b, g^c \in G$ are public keys and $r_1, r_2 \leftarrow \mathbb{Z}_p$, and the other is associated with the attribute set Atts by letting $W_j = H_1(\text{at}_j)^{r_2}$ for each $\text{at}_j \in \text{Atts}$, both of which are binding together with r_2 .
- Given credentials, a search token for keyword w can be generated with two parts: one is associated with keywords, $\text{tok}_1 = (g^a g^{bH_2(w)})^s$ and $\text{tok}_2 = g^{cs}$ by selecting $s \leftarrow \mathbb{Z}_p$, and the other is associated with credentials by letting $A'_v = A_v^s, B'_v = B_v^s$ for each $v \in \text{lvs}(T)$, both of which are binding together with the random number s .

As long as the attribute set satisfies the access tree T , with A'_v, B'_v and W_0, W_j , it can recover $e(g, g)^{acr_2s}$, which can be used to test the keyword equality for the keyword ciphertext and search token, the detail of which is described in the following KP-ABKS construction.

KP-ABKS Construction. The KP-ABKS can be constructed as follows:

$\text{Setup}(1^\ell)$: This algorithm selects a bilinear map $e : G \times G \rightarrow G_T$, such that G and G_T are cyclic groups of order p , an ℓ -bit prime. It selects $a, b, c \leftarrow \mathbb{Z}_p$ and $g \leftarrow G$. Let

$$\text{pm} = (e, g, p, g^a, g^b, g^c, G, G_T), \text{mk} = (a, b, c).$$

KeyGen(mk, T): This algorithm executes $\{q_v(0)|v \in \text{lvs}(\text{T})\} \leftarrow \text{Share}(\text{T}, ac)$ to obtain secret shares of ac for all leaves in T. For each leaf v in T, it picks $t \leftarrow \mathbb{Z}_p$, and computes $A_v = g^{q_v(0)} H_1(\text{att}(v))^t$ and $B_v = g^t$. Let

$$\text{sk} = (\text{T}, \{(A_v, B_v)|v \in \text{lvs}(\text{T})\}).$$

Enc(w , Atts): This algorithm selects $r_1, r_2 \leftarrow \mathbb{Z}_p$, and computes $W' = g^{cr_1}$, $W = g^{a(r_1+r_2)} g^{bH_2(w)r_1}$ and $W_0 = g^{r_2}$. For each $\text{at}_j \in \text{Atts}$, it computes $W_j = H_1(\text{at}_j)^{r_2}$. Let

$$\text{cph} = (\text{Atts}, W', W, W_0, \{W_j|\text{at}_j \in \text{Atts}\}).$$

TokenGen(sk, w): This algorithm selects $s \leftarrow \mathbb{Z}_p$, and computes $A'_v = A_v^s, B'_v = B_v^s$ for each $v \in \text{lvs}(\text{T})$. It computes $\text{tok}_1 = (g^a g^{bH_2(w)})^s$ and $\text{tok}_2 = g^{cs}$. Let

$$\text{tk} = (\text{tok}_1, \text{tok}_2, \text{T}, \{(A'_v, B'_v)|v \in \text{lvs}(\text{T})\})$$

Search(tk, cph): Given the attribute set Atts included in cph, this algorithm selects a minimum attribute set S satisfying the access tree T included in tk. If S does not exist, return 0; otherwise, for each $\text{at}_j \in S$, it computes

$$E_v = e(A'_v, W_0)/e(B'_v, W_j) = e(g, g)^{sr_2 q_v(0)},$$

where $\text{att}(v) = \text{at}_j$ for $v \in \text{lvs}(\text{T})$. Then it computes

$$e(g, g)^{sr_2 q_{\text{root}}(0)} \leftarrow \text{Combine}(\text{T}, \{E_v|\text{att}(v) \in S\})$$

so that $E_{\text{root}} = e(g, g)^{acsr_2}$. It return 1 if $e(W', \text{tok}_1)E_{\text{root}} = e(W, \text{tok}_2)$; otherwise return 0.

This completes the description of KP-ABKS scheme. The correctness can be verified as follows.

$$\begin{aligned} e(W', \text{tok}_1)E_{\text{root}} &= e(g^{cr_1}, (g^a g^{bH_2(w)})^s)E_{\text{root}} \\ &= ee(g, g)^{acsr_1} e(g, g)^{bcsH_2(w)r_1} E_{\text{root}} \\ &= e(g, g)^{acs(r_1+r_2)} e(g, g)^{bcsH_2(w)r_1} \\ \text{and } e(W, \text{tok}_2) &= ee(g^{a(r_1+r_2)} g^{bH_2(w)r_1}, g^{cs}) \\ &= ee(g, g)^{acs(r_1+r_2)} e(g, g)^{bcsH_2(w)r_1} \end{aligned}$$

To show the security of KP-ABKS, we have the following theorems, the proofs of which are deferred to Appendix A and Appendix B, respectively.

Theorem 1: Assume that the DL assumption holds, the KP-ABKS scheme is selectively secure against chosen-keyword attack in the random oracle model.

Theorem 2: Given the one-way hash function H_2 , the KP-ABKS scheme achieves keyword secrecy.

CP-ABKS Construction. The CP-ABKS scheme can be constructed as below:

Setup(1^ℓ): This algorithm selects a bilinear group $e : G \times G \rightarrow G_T$, such that G and G_T are cyclic groups of order p , an ℓ -bit prime. It selects $a, b, c \leftarrow \mathbb{Z}_p$ and $g \leftarrow G$. Let

$$\text{pm} = (e, g, p, g^a, g^b, g^c, G, G_T), \text{mk} = (a, b, c).$$

KeyGen(mk, Atts): This algorithm selects $r \leftarrow \mathbb{Z}_p$, computes $A = g^{(ac-r)/b}$. Then for each $\text{at}_j \in \text{Atts}$, it selects $r_j \leftarrow \mathbb{Z}_p$ and computes $A_j = g^r H_1(\text{at}_j)^{r_j}$ and $B_j = g^{r_j}$. Let

$$\text{sk} = (\text{Atts}, A, \{(A_j, B_j)|\text{at}_j \in \text{Atts}\}).$$

Enc(w , T): This algorithm selects $r_1, r_2 \leftarrow \mathbb{Z}_p$, and computes $W = g^{cr_1}$, $W_0 = g^{a(r_1+r_2)} g^{bH_2(w)r_1}$ and $W' = g^{br_2}$. It further computes secret shares of r_2 for all leaves of T by running $\{q_v(0)|v \in \text{lvs}(\text{T})\} \leftarrow \text{Share}(\text{T}, r_2)$. Then for each $v \in \text{lvs}(\text{T})$, it computes $W_v = g^{q_v(0)}$ and $D_v = H_1(\text{att}(v))^{q_v(0)}$. Let

$$\text{cph} = (\text{T}, W, W_0, W', \{(W_v, D_v)|v \in \text{lvs}(\text{T})\})$$

TokenGen(sk, w): This algorithm selects $s \leftarrow \mathbb{Z}_p$, and computes $\text{tok}_1 = (g^a g^{bH_2(w)})^s, \text{tok}_2 = g^{cs}$ and $\text{tok}_3 = A^s = g^{(acs-rs)/b}$. Then for each $\text{at}_j \in \text{Atts}$, it computes $A'_j = A_j^s$ and $B'_j = B_j^s$. Let

$$\text{tk} = (\text{Atts}, \text{tok}_1, \text{tok}_2, \text{tok}_3, \{(A'_j, B'_j)|\text{at}_j \in \text{Atts}\}).$$

Search(tk, cph): Given an attribute set Atts included in tk, this algorithm selects a minimum attribute set S that satisfies the access tree T included in cph. If S does not exist, return 0; otherwise, for each $\text{at}_j \in S$, it computes

$$E_v = e(A'_j, W_v)/e(B'_j, D_v) = e(g, g)^{rsq_v(0)},$$

where $\text{att}(v) = \text{at}_j$ for $v \in \text{lvs}(\text{T})$. Then it computes

$$e(g, g)^{rsq_{\text{root}}(0)} \leftarrow \text{Combine}(\text{T}, \{E_v|\text{att}(v) \in S\})$$

so that $E_{\text{root}} = e(g, g)^{rsr_2}$. It returns 1 if $e(W_0, \text{tok}_2) = e(W, \text{tok}_1)E_{\text{root}}e(\text{tok}_3, W')$; otherwise return 0.

This completes the CP-ABKS construction, and its correctness can be verified similar to that of KP-ABKS. To show the security of CP-ABKS, we have Theorem 3 and Theorem 4. The proof of former one is deferred to Appendix C, and the proof of latter one is omitted because it is similar to that of Theorem 2.

Theorem 3: CP-ABKS scheme is selectively secure against chosen-keyword attack in the generic bilinear group model.

Theorem 4: Given the one way hash function H_2 , the CP-ABKS scheme achieves keyword secrecy.

IV. VERIFIABLE ATTRIBUTE-BASED KEYWORD SEARCH

While ABKS enables data owners to grant search capability over outsourced encrypted keywords, it is not adequate for secure data sharing in the cloud computing environment in that the cloud is untrusted, meaning that the cloud might return incomplete/manipulated search result. Therefore, in order to combat this, it requires another layer protection – verifiability, validating whether the cloud has faithfully executed search operations and returned research result honestly.

A. System Model and Threat Model

We consider the system model as illustrated in Figure 1, which is involved four entities: the data owner, who has large amount of data indexed by keywords (i.e. inverted index as in Figure 1) and will store it in the cloud; the cloud, which is to provide data storage and retrieval service and has sufficient storage space and computation resource; the data user, who is to retrieve data owner's encrypted data according to some keyword; and the trusted authority that issues credentials and necessary public/private keys to all




Fig. 1. System model for verifiable attribute-based keyword search over outsourced encrypted data, where keywords X, Y and V, W have different access control policies.

cloud users. We assume that the authentication between any two entities can be appropriately done.

The data owner encrypts its own data (i.e. keywords and their corresponding data files) before outsourcing for the sake of data privacy. The data owner can specify access control policies to determine who can generate appropriate search token to conduct search over encrypted keywords. When the data user wishes to retrieve encrypted data (i.e. encrypted keyword and the associated encrypted data files) from the cloud according to some keyword, he generates a search token and delivers it to the cloud. Upon request, the cloud conducts keyword search over encrypted data and returns the corresponding encrypted data file to the user. To show its faithful execution, the cloud also returns necessary proof so that the user can verify that the returned result is correct.

We assume that the data owner are trusted, and data users (i.e. w/ or w/o proper attributes satisfying the data owner's access control policies) are semi-trusted, meaning they may try to derive other sensitive information beyond their own. We assume that the cloud server is untrusted: he might either conduct search operation with arbitrary input, or return incorrect results which were manipulated.

B. VABKS Definition and Security

Let $D = (KS = \{KS_1, \dots, KS_l\}, MP = \{MP(w) | w \in \cup_{i=1}^l KS_i\}, FS = \{F_1, \dots, F_n\})$ denote the inverted index and the set of data files, where $KS_i, 1 \leq i \leq l$, is the set of keywords that will be encrypted with the same access control policy, $F_j, 1 \leq j \leq n$, is the data file, and $MP(w)$ denote the set of ordered identifiers identifying data files associated with keyword w .

Definition 4: A VABKS scheme consists of the following algorithms:

- $(mk, pm) \leftarrow \text{Init}(1^\ell)$: This algorithm is run by a trusted party (e.g., attribute authority) when initializing the system.
- $sk \leftarrow \text{KeyGen}(mk, I_{\text{KeyGen}})$: This algorithm is run by the trusted party to issue credential sk to a data user.
- $(Au, \text{Index}, D_{\text{cph}}) \leftarrow \text{BuildIndex}(\{I_{\text{Enc}}\}_l, \{I'_{\text{Enc}}\}_n, D)$: This algorithm is run by the data owner to encrypt data collection $D = (KS, MP, FS)$ and obtain encrypted data files D_{cph} , encrypted index Index and auxiliary information Au , where $\{I_{\text{Enc}}\}_l$ is the set of access control policies to encrypt KS and $\{I'_{\text{Enc}}\}_n$ is the set of access control policies to encrypt FS .

- $tk \leftarrow \text{TokenGen}(sk, w)$: This algorithm is run by the data user, who has the credential sk , to issue search token tk on keyword w .
- $(\text{proof}, \text{rslt}) \leftarrow \text{SearchIndex}(Au, \text{Index}, D_{\text{cph}}, tk)$: This algorithm is run by the cloud to perform keyword search over encrypted data on behalf of the data user. It outputs the search result rslt and a proof proof to show that the result is correct.
- $\{0, 1\} \leftarrow \text{Verify}(sk, w, tk, \text{rslt}, \text{proof})$: This algorithm is run by the data user to verify that $(\text{rslt}, \text{proof})$ is valid with respect to search token tk .

The correctness of the VABKS scheme requires that, given $(mk, pm) \leftarrow \text{Init}(1^\ell)$, $sk \leftarrow \text{KeyGen}(mk, I_{\text{KeyGen}})$, for any keyword-based data collection D and keyword w , $(Au, \text{Index}, D_{\text{cph}}) \leftarrow \text{BuildIndex}(\{I_{\text{Enc}}\}_l, \{I'_{\text{Enc}}\}_n, D)$, $tk \leftarrow \text{TokenGen}(sk, w)$ and $(\text{proof}, \text{rslt}) \leftarrow \text{SearchIndex}(Au, \text{Index}, D_{\text{cph}}, tk)$, then $1 \leftarrow \text{Verify}(sk, w, tk, \text{rslt}, \text{proof})$ always holds.

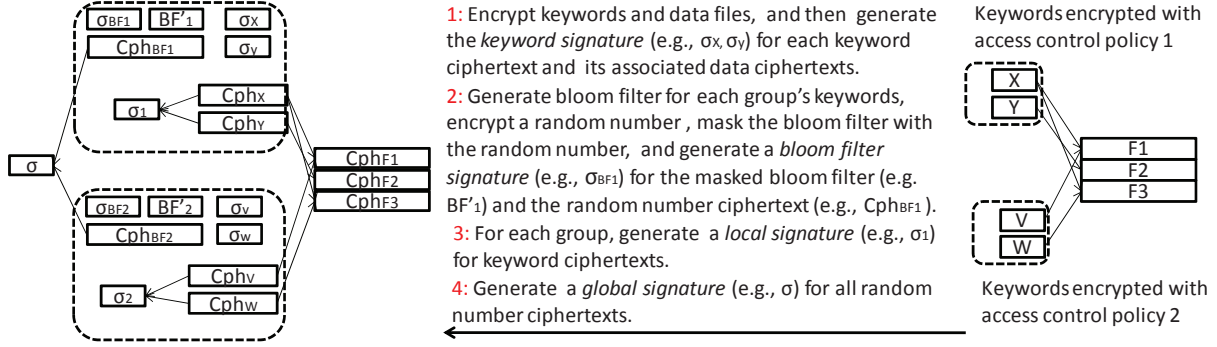
Intuitively, a VABKS scheme should satisfy the following security requirements even in the presence of a probabilistic polynomial-time adversary \mathcal{A} (i.e. the cloud server):

- **Data secrecy:** Even if \mathcal{A} is given keywords, it still cannot learn any information about encrypted data files. It can be formalized via the conventional chosen-plaintext security game with additional requirement that the two challenge data $D_0 = (KS, MP, FS_0), D_1 = (KS, MP, FS_1)$ have the same KS and MP , and $|FS_0| = |FS_1|$. Otherwise, \mathcal{A} can distinguish FS_0 and FS_1 by observing the search result with the search token for a specific keyword.
- **Selective security against chosen-keyword attack:** Without seeing corresponding search tokens, \mathcal{A} cannot guess which of the given keyword sets was encrypted. It can be extended from the definition of selective security against chosen-keyword attack of ABKS.
- **Keyword secrecy:** Even if \mathcal{A} is given the data, the probability of \mathcal{A} learning the keyword from the keyword ciphertext and the search token is no more than of a random guess within the set of keywords which composes the given data sets. This security definition also can be extended from the keyword secrecy of ABKS.
- **Verifiability:** \mathcal{A} cannot return incorrect search results without being caught. Specifically, after executing algorithm SearchIndex , \mathcal{A} responses with the search result and the proof. If \mathcal{A} returns incorrect and/or incomplete search result, the cheating behavior can be detected by algorithm Verify with overwhelming probability. We formalize this security property via the following verifiability game.

Verifiability Game:

Setup: The challenger runs $(pm, mk) \leftarrow \text{Init}(1^\ell)$. \mathcal{A} selects $D = (KS, MP, FS), \{I_{\text{Enc}}\}_l$ and $\{I'_{\text{Enc}}\}_n$ and sends them to the challenger. The challenger runs $(Au, \text{Index}, D_{\text{cph}}) \leftarrow \text{BuildIndex}(\{I_{\text{Enc}}\}_l, \{I'_{\text{Enc}}\}_n, D)$, and gives $(Au, \text{Index}, D_{\text{cph}})$ to \mathcal{A} .

Phase 1: \mathcal{A} can query the following oracles for polynomially many times.



- 1: Encrypt keywords and data files, and then generate the *keyword signature* (e.g., σ_x, σ_y) for each keyword ciphertext and its associated data ciphertexts.
- 2: Generate bloom filter for each group's keywords, encrypt a random number, mask the bloom filter with the random number, and generate a *bloom filter signature* (e.g., σ_{BF1}) for the masked bloom filter (e.g. $BF'1$) and the random number ciphertext (e.g., Cph_{BF1}).
- 3: For each group, generate a *local signature* (e.g., σ_1) for keyword ciphertexts.
- 4: Generate a *global signature* (e.g., σ) for all random number ciphertexts.




Fig. 2. Basic idea for verifiability, where data files F_1, F_2, F_3 were encrypted to $cph_{F_1}, cph_{F_2}, cph_{F_3}$, keywords X, Y were encrypted to cph_X, cph_Y with access control policy 1, and keywords V, W were encrypted to cph_V, cph_W with access control policy 2 (different from access control policy 1).

- $\mathcal{O}_{KeyGen}(I_{KeyGen})$: The challenger returns to \mathcal{A} credential sk corresponding to I_{KeyGen} .
- $\mathcal{O}_{TokenGen}(I_{KeyGen}, w)$: The challenger generates credential sk with I_{KeyGen} , and returns to \mathcal{A} a search token tk by running algorithm $TokenGen$ with inputs sk and w .
- $\mathcal{O}_{Verify}(I_{KeyGen}, w, tk, rslt, proof)$: The challenger generates credential sk with I_{KeyGen} , returns γ to \mathcal{A} by running $\gamma \leftarrow Verify(sk, w, tk, rslt, proof)$.

Challenge phase: \mathcal{A} selects a non-trivial challenge I_{Enc}^* and a keyword w^* and gives them to the challenger. The challenger selects I_{KeyGen}^* such that $F(I_{KeyGen}^*, I_{Enc}^*) = 1$, generates credential sk^* with I_{KeyGen}^* and returns to \mathcal{A} a search token tk^* by running $tk^* \leftarrow TokenGen(sk, w^*)$.

Guess: \mathcal{A} outputs $(rslt^*, proof^*)$ to the challenger. We say \mathcal{A} wins the game if $1 \leftarrow Verify(sk^*, w^*, tk^*, rslt^*, proof^*)$ and $rslt^* \neq rslt$, where $(rslt, proof)$ is produced by the challenger by running $SearchIndex(Au, Index, tk^*)$.

Definition 5: A VABKS scheme is verifiable if the advantage that any \mathcal{A} wins the verifiability game is negligible in security parameter ℓ .

C. VABKS Construction

It is natural to adopt ABKS as a building block to enable keyword search over encrypted data and grant search capability, and let us focus on how to achieve verifiability for VABKS.

Note that when keywords are to be encrypted, they are separated into multiple groups according to access control policies, so that keywords within each group are unique (meaning non redundant keyword exists in the same group). To achieve verifiability, the data user should verify that (1) the cloud performed search operations on all group's keyword ciphertexts; and (2) the cloud honestly returned the search result for each group, which could be either a null result or one keyword ciphertext and its associated data ciphertexts.

Specifically, we elaborate the basic idea for achieving verifiability in Figure 2, where the data owner generates the signatures and bloom filters as follows:

- The first type of signature, called *keyword signature*, is generated for each keyword ciphertext and its associated data ciphertexts, which is prevent the cloud from returning part of data ciphertexts as the search result.

- For each group, one bloom filter is built from its keywords, which enables data users to *efficiently* check that the queried keyword was indeed not in the group when the cloud returns a null result, without downloading all keyword ciphertexts from the cloud. A random number is selected and encrypted with the same access control policy as keywords, and is used to mask the bloom filter for preserving keyword privacy. The second type of signature, called *bloom filter signature*, is generated for the masked bloom filter and the random number ciphertext for assuring their integrity.
- The third type of signature, called *global signature*, is obtained by signing random number ciphertexts of all groups. It enforces the cloud to perform search operations over all groups' keyword ciphertexts. The reason of selecting random number ciphertext is that it has identical access control policy with that of keyword ciphertext within the same group, and can be used to assure that the cloud only return one result (i.e. null or one keyword ciphertext and its associated ciphertexts) for each group.
- The last type of signature, called *local signature*, is generated from all keyword ciphertexts within the same group. When data users download one group's keyword ciphertexts from the cloud (due to the false positive property of the bloom filter), this signature is to validate the integrity of those keyword ciphertexts.

The VABKS scheme is depicted in Figure 3, which makes use of a symmetric encryption $SE = (KeyGen, Enc, Dec)$, an attribute-based encryption scheme $ABE = (Setup, KeyGen, Enc, Dec)$, both of which are used to encrypt data files, an attribute-based keyword search scheme $ABKS = (Setup, KeyGen, Enc, TokenGen, Search)$, which is to encrypt keywords, and a signature scheme $Sig = (KeyGen, Sign, Verify)$. Here ABE and ABKS are either ciphertext-policy or key-policy, so that we have two variants of VABKS.

In algorithm $Verify$ of Figure 3, when the authorized data user verifies a null result for the group $\{cph_w | w \in KS_i\}$, given a search token for keyword w' , it can happen that $1 \leftarrow BFVerify(\{H'_1, \dots, H'_k\}, BF_i, w')$ but the matched keyword ciphertext was not stored in the cloud. The probability

Init(1^ℓ): Given security parameter ℓ , the attribute authority chooses k hash functions H'_1, \dots, H'_k , which are used to construct m -bit Bloom filter. Let $H : \{0,1\}^\ell \rightarrow \{0,1\}^m$ be a secure pseudorandom generator. It executes $(\text{ABE.pm}, \text{ABE.mk}) \leftarrow \text{ABE.Setup}(1^\ell)$ and $(\text{ABKS.pm}, \text{ABKS.mk}) \leftarrow \text{ABKS.Setup}(1^\ell)$. Let $\text{pm} = (\text{ABE.pm}, \text{ABKS.pm}, H'_1, \dots, H'_k)$ which is public known, and $\text{mk} = (\text{ABE.mk}, \text{ABKS.mk})$.

KeyGen($\text{mk}, I_{\text{KeyGen}}$): The attribute authority runs $\text{ABE.sk} \leftarrow \text{ABE.KeyGen}(\text{ABE.mk}, I_{\text{KeyGen}})$ and $\text{ABKS.sk} \leftarrow \text{ABKS.KeyGen}(\text{ABKS.mk}, I_{\text{KeyGen}})$, and sets $\text{sk} = (\text{ABE.sk}, \text{ABKS.sk})$, which is delivered to the user via authenticated private channel.

BuildIndex($\{I_{\text{Enc}}\}_l, \{I'_{\text{Enc}}\}_n, D$): If the data owner has no public/private keys for the signature scheme Sig, then it runs $(\text{Sig.sk}, \text{Sig.pk}) \leftarrow \text{Sig.KeyGen}(1^\ell)$, by keeping Sig.sk private and making Sig.pk public. Given $D = (\text{KS} = \{\text{KS}_1, \dots, \text{KS}_l\}, \text{MP} = \{\text{MP}(w) | w \in \cup_{i=1}^l \text{KS}_i\}, \text{FS} = \{F_1, \dots, F_n\})$, the data owner executes as follows:

- 1) Encrypt the data file with hybrid encryption: $\forall F_j \in \text{FS}$, it generates the ciphertext $\text{cph}_{F_j} = (\text{cph}_{\text{sk}_j}, \text{cph}_{\text{SE}_j})$ by running $\text{SE.sk}_j \leftarrow \text{SE.KeyGen}()$, $\text{cph}_{\text{SE}_j} \leftarrow \text{SE.Enc}(\text{SE.sk}_j, F_j)$, and $\text{cph}_{\text{sk}_j} \leftarrow \text{ABE.Enc}(I'_{\text{Enc}_j}, \text{SE.sk}_j)$.
- 2) Encrypt each keyword and generate keyword signature: Given $\text{KS}_i, 1 \leq i \leq l$, for each $w \in \text{KS}_i$, it runs $\text{cph}_w \leftarrow \text{ABKS.Enc}(I_{\text{Enc}_i}, w)$, sets $\text{MP}(\text{cph}_w) = \{\text{ID}_{\text{cph}_{F_j}} | \text{ID}_{F_j} \in \text{MP}(w)\}$, and generates $\sigma_w \leftarrow \text{Sig.Sign}(\text{Sig.sk}, \text{cph}_w || \{\text{cph}_{F_j} | \text{ID}_{\text{cph}_{F_j}} \in \text{MP}(\text{cph}_w)\})$, where identities in $\text{MP}(\text{cph}_w)$ are ordered, and ID_{F_j} and $\text{ID}_{\text{cph}_{F_j}}$ are identities identifying data file F_j and data ciphertext cph_{F_j} , respectively.
- 3) Generate the bloom filter, bloom filter signature and local signature for each group KS_i : Let $\text{BF}_i \leftarrow \text{BFGen}(\{H'_1, \dots, H'_k\}, \text{KS}_i)$, $\text{cph}_{\text{BF}_i} \leftarrow \text{ABE.Enc}(I_{\text{Enc}_i}, M)$ by randomly selecting M from the message space of ABE, compute $\text{BF}'_i = H(M) \otimes \text{BF}_i$ and generate $\sigma_{\text{BF}_i} \leftarrow \text{Sig.Sign}(\text{Sig.sk}, \text{BF}'_i || \text{cph}_{\text{BF}_i})$. Let $\sigma_i \leftarrow \text{Sig.Sign}(\text{Sig.sk}, \text{cph}_{w_1} || \dots || \text{cph}_{|\text{KS}_i|})$.
- 4) Generate the global signature: Let $\sigma = \text{Sig.Sign}(\text{Sig.sk}, \text{cph}_{\text{BF}_1} || \dots || \text{cph}_{\text{BF}_l})$.
- 5) Let $\text{Au} = (\sigma, \sigma_1, \dots, \sigma_l, \sigma_{\text{BF}_1}, \dots, \sigma_{\text{BF}_l}, \{\sigma_w | w \in \cup_{i=1}^l \text{KS}_i\})$, $\text{Index} = (\{\text{cph}_w | w \in \cup_{i=1}^l \text{KS}_i\}, \{\text{MP}(\text{cph}_w) | w \in \cup_{i=1}^l \text{KS}_i\})$ and $\text{D}_{\text{cph}} = (\{\text{cph}_{F_j} | F_j \in \text{FS}\})$.

TokenGen(sk, w): Given sk , the data user generates search token $\text{tk} \leftarrow \text{ABKS.TokenGen}(\text{ABKS.sk}, w)$.

SearchIndex($\text{Au}, \text{Index}, \text{D}_{\text{cph}}, \text{tk}$): Let rslt be an empty set and $\text{proof} = (\sigma)$ initially. Let $\prod = \{\text{cph}_w | w \in \text{KS}_i\}$ be the set of keyword ciphertexts having the identical access control policy. For each $\text{cph}_w \in \prod$, it runs $\gamma \leftarrow \text{ABKS.Search}(\text{cph}_w, \text{tk})$, and if $\gamma = 1$, then adds $(\text{cph}_w, \{\text{cph}_{F_j} | \text{ID}_{\text{cph}_{F_j}} \in \text{MP}(\text{cph}_w)\})$ to rslt and $(\sigma_w, \text{cph}_{\text{BF}_i})$ to proof . If there exist no $\gamma = 1$ after enumerating all cph_w in S , then add $(\text{BF}'_i, \text{cph}_{\text{BF}_i}, \sigma_{\text{BF}_i})$ to the proof .

Verify($\text{sk}, w, \text{tk}, \text{proof}, \text{rslt}$): The data user verifies the search result from the cloud as follows:

- 1) Verify the integrity of random number ciphertexts: $\gamma = \text{Sig.Verify}(\text{Sig.pk}, \sigma, \text{cph}_{\text{BF}_1} || \dots || \text{cph}_{\text{BF}_l})$. If $\gamma = 0$, then returns 0.
- 2) For $i = 1, \dots, l$, it executes as follows to ensure that the cloud returned correct result for each group:
 - Case 1: When the search result with respect to the group is not null, meaning $(\text{cph}_w, \{\text{cph}_{F_j} | \text{ID}_{\text{cph}_{F_j}} \in \text{MP}(\text{cph}_w)\})$ is included in rslt , then it runs $\gamma \leftarrow \text{ABKS.Search}(\text{cph}_w, \text{tk})$ and $\gamma' \leftarrow \text{Sig.Verify}(\text{Sig.pk}, \sigma_w, \text{cph}_w || \{\text{cph}_{F_j} | \text{ID}_{\text{cph}_{F_j}} \in \text{MP}(\text{cph}_w)\})$. If either $\gamma = 0$ or $\gamma' = 0$, then return 0.
 - Case 2: When the search result with respect to the group is null, meaning $(\text{BF}'_i, \text{cph}_{\text{BF}_i}, \sigma_{\text{BF}_i})$ is included in rslt , then it runs $\gamma' \leftarrow \text{Sig.Verify}(\text{Sig.pk}, \sigma_{\text{BF}_i}, \text{BF}'_i || \text{cph}_{\text{BF}_i})$. If $\gamma' = 0$, return 0; otherwise, it executes as follows:
 - If the data user is authorized, it computes $M \leftarrow \text{ABE.Dec}(\text{ABE.sk}, \text{cph}_{\text{BF}_i})$, $\text{BF}_i = H(M) \otimes \text{BF}'_i$, and $\delta \leftarrow \text{BFVerify}(\{H'_1, \dots, H'_k\}, \text{BF}, w)$. If $\delta = 1$, then it downloads $\{\text{cph}_w | w \in \text{KS}_i\}$ and σ_i from the cloud, and runs $\eta \leftarrow \text{Sig.Verify}(\text{Sig.pk}, \sigma_i, \text{cph}_{w_1} || \dots || \text{cph}_{w_{|\text{KS}_i|}})$. If $\eta = 0$, return 0; otherwise it runs $\tau \leftarrow \text{ABKS.Search}(\text{cph}_w, \text{tk})$ for each cph_w . If there exists some $\tau = 1$, return 0.
 - If the data user's attributes do not satisfy the access control policy, then he continues because cph_{BF_i} cannot be decrypted.
- 3) Return 1, meaning that the search result rslt from the cloud is correct.

Fig. 3. VABKS construction

that this event happens equals to the false-positive rate of the Bloom filter. In order to validate the search result in this case, algorithm Verify has to download $\{cph_w | w \in KS_i\}$, and check one by one. Because this incurs high computational and communication cost, we can set the false-positive rate as low as possible by choosing appropriate m and k . For example, in our experiment we set the false-positive rate to be 4.5×10^{-9} .

D. Security Analysis

We show that the proposed VABKS scheme satisfies the security requirements with the following theorems, the formal proofs of which are deferred to Appendix D.

We show that if there exists polynomial time algorithm \mathcal{A} breaking VABKS's data secrecy, then it breaks the assumption that CPA-secure ABE and CPA-secure SE. This is formally achieved via Theorem 5.

Theorem 5: If the ABE is CPA-secure and SE is CPA-secure, the VABKS scheme achieves data secrecy.

We show that if there exists polynomial time algorithm \mathcal{A} breaking VABKS's selective security against chosen keyword attack, then it breaks the assumption that ABKS achieves the selective security against chosen keyword attack, given that ABE is CPA-secure and H is a secure pseudorandom generator. This is formally achieved via Theorem 6.

Theorem 6: If ABE is CPA-secure, H is a secure pseudorandom generator and the ABKS achieves selective security against chosen keyword attack, the VABKS scheme also achieves security against chosen keyword attack.

We show that if there exists polynomial time algorithm \mathcal{A} breaking keyword secrecy of the VABKS, then it breaks the assumption that ABKS achieves keyword secrecy, given that ABE is CPA-secure and H is a secure pseudorandom generator. This is formally achieved via Theorem 7.

Theorem 7: If ABE is CPA-secure, H is a secure pseudorandom generator and the ABKS achieves keyword secrecy, the VABKS scheme also achieves keyword secrecy.

We show that if there exists polynomial time algorithm \mathcal{A} breaking verifiability of the VABKS then it breaks the Sig's unforgeability. This is formally achieved via Theorem 8.

Theorem 8: If Sig is a secure signature, the VABKS construction achieves the verifiability.

V. PERFORMANCE EVALUATION

We evaluate efficiency of the ABKS schemes in terms of both asymptotic complexity and actual execution time, and efficiency of the VABKS scheme in terms of actual execution time. Asymptotic complexity is measured in terms of four kinds of operations: H_1 denotes the operation of mapping a bit-string to an element of G , Pair denotes the pairing operation, E denotes the exponentiation operation in G , and E_T denotes the exponentiation operation in G_T . We ignore multiplication and hash operations (other than H_1) because they are much more efficient than the above operations (see benchmark of [37]).

We implemented ABKS and VABKS in the JAVA language based on the Java Pairing Based Cryptography library (jPBC)

[37]. In our implementations, we instantiated the bilinear map with Type A pairing, which offers a level of security that is equivalent to 1024-bit DLOG [37]. For both CP-VABKS and KP-VABKS, we instantiated the symmetric encryption scheme with AES, and the signature scheme with DSA signature scheme provided by JDK1.6. We instantiated ABKS, ABE with CP-ABKS, CP-ABE [24] for CP-ABKS, and KP-ABKS, KP-ABE [23] for KP-VABKS, respectively. Finally, we set the example access control policy as "at₁ AND ... AND at_N."

A. Efficiency of ABKS

Asymptotic Complexity of the ABKS Schemes. Table I describes the asymptotic complexity of the ABKS schemes. We observe that in the CP-ABKS scheme, the complexity of KeyGen is almost the same as that of Enc. However, in the KP-ABKS scheme, KeyGen is more expensive than Enc. This discrepancy should serve as a factor when deciding whether to use CP-ABKS or KP-ABKS. In both schemes, the two Search algorithms incur almost the same cost.

ABKS Performance. To evaluate the performance of the ABKS schemes, we ran the experiments on a client machine with Linux OS, 2.93GHz Intel Core Duo CPU (E7500), and 2GB RAM. We varied N , the number of attributes that are involved in the example access control policy, from 1 to 100. We ran each experiment for 10 times so as to average the execution time. Table II shows the execution time of the two ABKS schemes.

From Table II we observe that for both schemes, the keyword encryption algorithm Enc (run by the data owner) is more expensive than that of the keyword search algorithm Search (run by the cloud) with the same N . However, we note that the former algorithm is executed only once for each keyword, whereas the latter algorithm will be performed as many times as needed. This means that even though the keyword search algorithm is relatively light-weight, it can be a performance bottleneck when executed many times. This confirms that the data users should outsource the keyword search operations to the cloud (i.e., taking advantage of the cloud's computational resources).

B. Efficiency of VABKS with Real Data

To demonstrate the feasibility of VABKS in practice, we evaluated it using some real data. The algorithms run by the data owner and the data users (i.e. BuildIndex, TokenGen and Verify) were executed on a client machine with Linux OS, 2.93GHz Intel Core Duo CPU (E7500), and 2GB RAM. The algorithm run by the cloud (i.e., SearchIndex) was executed on a server machine (a laptop) with Windows 7, Intel i5 2.60GHz CPU, and 8GB RAM.

We vary the access control policy ranging from 1 to 50 attributes with step-length 10. In each round, we encrypted all keywords with the identical access control policy. The data consists of 2,019 distinct keywords extracted from 670 PDF documents (papers) from the ACM Digital Library with a total size of 778.1MB. We choose $k = 28$ and $m = 10KB$ so that

KP- ABKS	KeyGen	Enc	TokenGen	Search
	Operations	$3NE + NH_1$	$(S + 4)E + SH_1$	$(2N + 2)E$
Output size	$2N G $	$(S + 3) G $	$(2N + 2) G $	
CP- ABKS	KeyGen	Enc	TokenGen	Search
	Operations	$(2S + 2)E + SH_1$	$(2N + 4)E + NH_1$	$(2S + 4)E$
Output size	$(4S + 1) G $	$(2N + 2) G $	$(2S + 2) G $	

TABLE I

ASYMPTOTIC COMPLEXITIES OF CP-ABKS AND KP-ABKS, WHERE S IS THE NUMBER OF A DATA USER'S ATTRIBUTES AND N IS THE NUMBER OF ATTRIBUTES THAT ARE INVOLVED IN A DATA OWNER'S ACCESS CONTROL POLICY (I.E., THE NUMBER OF LEAVES IN THE ACCESS TREE).

KP- ABKS	KeyGen	0.088	0.786	1.539	2.316	3.081	3.863	4.646	5.396	6.156	6.928	7.705
	Enc	0.108	0.539	1.016	1.492	1.983	2.434	2.939	3.414	3.873	4.358	4.820
TokenGen	0.073	0.331	0.627	0.917	1.211	1.504	1.778	2.081	2.364	2.673	2.965	
Search	0.049	0.275	0.480	0.711	0.947	1.182	1.425	1.665	1.911	2.153	2.405	
S(N)	1	10	20	30	40	50	60	70	80	90	100	
CP- ABKS	KeyGen	0.107	0.686	1.275	1.901	2.525	3.151	3.779	4.405	5.014	5.640	6.262
	Enc	0.121	0.681	1.304	1.923	2.546	3.169	3.787	4.415	5.060	5.673	6.28
TokenGen	0.088	0.349	0.673	0.932	1.228	1.513	1.823	2.082	2.399	2.665	2.97	
Search	0.061	0.329	0.493	0.728	0.97	1.202	1.441	1.688	1.929	2.182	2.424	

TABLE II

EXECUTION TIME (SECOND) OF THE ALGORITHMS IN THE KP-ABKS AND CP-ABKS SCHEMES, WHERE N IS THE NUMBER OF ATTRIBUTES INVOLVED IN THE EXAMPLE ACCESS CONTROL POLICY. THE NUMBER OF DATA USER'S ATTRIBUTES IS ALSO SET TO N , NAMELY $S = N$ IN THE EXPERIMENTS.




Fig. 4. Performance of the CP-VABKS and KP-VABKS schemes, where N is the number of attributes involved in the example access control policy. The number of data user's attributes is also set to N , namely $S = N$ in the experiments.

$\frac{m}{n} = \frac{10 \cdot 8 \cdot 1024}{2019} \approx 40$ and the false-positive rate is around 4.5×10^{-9} .

Figure 4(a) shows the execution time of BuildIndex that was run by the data owner. We can observe that with the same attribute/policy complexity, CP-VABKS cost more time than that of KP-VABKS when running algorithm BuildIndex. Figure 4(b) plots the execution time of the algorithms run by the data user and the cloud. As algorithm SearchIndex conducts search over keyword ciphertexts one by one sequentially, in the experiments, we generated the search token which matched the 1010th keyword ciphertext. That is, when testing the algorithm SearchIndex, it needs to conduct search operations over 1010 keyword ciphertexts. We can observe that the execution time of TokenGen and Verify is really small, compared with keyword search algorithm SearchIndex. This again confirms that the data user should outsource keyword search operations to the cloud. Figure 4(c) plots the size of index, including 2,019 keyword ciphertexts, bloom filter and signatures. We also see that CP-VABKS consumes more storage space than KP-VABKS while with the same attribute/policy complexity. The reason is that in CP-VABKS scheme, it needs to store extra access trees when serializing CP-ABKS ciphertexts to the storage space.

VI. CONCLUSION

We have introduced a novel cryptographic scheme called verifiable attribute-based keyword search, which is a new tool for secure cloud computing over outsourced encrypted data. This scheme can be used for secure cloud computing: Given the keyword-based data collection, data owners is allowed to control the search capability and use of its outsourced encrypted data according to their access control policy, while authorized data users can outsource the often costly search operations to the cloud and forces the cloud to faithfully execute search operations. Performance evaluation shows that the new tool is practical.

REFERENCES

- [1] D. X. Song, D. Wagner, and A. Perrig, "Practical techniques for searches on encrypted data," in *Proc. of IEEE S&P*, pp. 44–, IEEE Computer Society, 2000.
- [2] E.-J. Goh, "Secure indexes." Cryptology ePrint Archive, Report 2003/216, 2003. <http://eprint.iacr.org/2003/216/>.
- [3] Y.-C. Chang and M. Mitzenmacher, "Privacy preserving keyword searches on remote encrypted data," in *Proc. of ACNS*, pp. 442–455, Springer-Verlag, 2005.
- [4] R. Curtmola, J. Garay, S. Kamara, and R. Ostrovsky, "Searchable symmetric encryption: improved definitions and efficient constructions," in *Proc. of*, pp. 79–88, ACM, 2006.
- [5] M. Chase and S. Kamara, "Structured encryption and controlled disclosure," in *Proc. of ASIACRYPT*, pp. 577–594, 2010.

- [6] K. Kurosawa and Y. Ohtaki, “Uc-secure searchable symmetric encryption,” in *Proc. of FC*, pp. 285–298, Springer Berlin / Heidelberg.
- [7] S. Kamara and K. Lauter, “Cryptographic cloud storage,” in *Proc. of FC*, pp. 136–149, 2010.
- [8] S. Kamara, C. Papamanthou, and T. Roeder, “Cs2: A searchable cryptographic cloud storage system.” Microsoft Technical Report, 2011. <http://research.microsoft.com/apps/pubs/?id=148632>.
- [9] S. Kamara, C. Papamanthou, and T. Roeder, “Dynamic searchable symmetric encryption,” in *Proc. of ACM CCS*, pp. 965–976, 2012.
- [10] D. Boneh, G. D. Crescenzo, R. Ostrovsky, and G. Persiano, “Public key encryption with keyword search,” in *Proc. of EUROCRYPT*, pp. 506–522, 2004.
- [11] B. R. Waters, D. Balfanz, G. Durfee, and D. K. Smetters, “Building an encrypted and searchable audit log,” in *Proc. of NDSS*, 2004.
- [12] M. Bellare, A. Boldyreva, and A. O’Neill, “Deterministic and efficiently searchable encryption,” in *Proc. of CRYPTO*, pp. 535–552, Springer-Verlag, 2007.
- [13] J. Baek, R. Safavi-Naini, and W. Susilo, “Public key encryption with keyword search revisited,” in *Proc. of IACM CCSA*, pp. 1249–1259, 2008.
- [14] J. Camenisch, M. Kohlweiss, A. Rial, and C. Sheedy, “Blind and anonymous identity-based encryption and authorised private searches on public key encrypted data,” in *Proc. of PKC*, pp. 196–214, Springer-Verlag, 2009.
- [15] P. Golle, J. Staddon, and B. Waters, “Secure conjunctive keyword search over encrypted data,” in *Proc. of ACNS*, pp. 31–45, Springer-Verlag, 2004.
- [16] E. Shi, J. Bethencourt, H. T.-H. Chan, D. X. Song, and A. Perrig, “Multi-dimensional range query over encrypted data,” in *Proc. of IEEE S&P*, pp. 350–364, 2007.
- [17] D. Boneh and B. Waters, “Conjunctive, subset, and range queries on encrypted data,” in *Proc. of TCC*, pp. 535–554, Springer-Verlag, 2007.
- [18] M. Li, S. Yu, N. Cao, and W. Lou, “Authorized private keyword search over encrypted data in cloud computing,” in *Proc. of ICDCS*, pp. 383–392, IEEE Computer Society, 2011.
- [19] F. Bao, R. H. Deng, X. Ding, and Y. Yang, “Private query on encrypted data in multi-user settings,” in *Proc. of ISPEC*, pp. 71–85, Springer-Verlag, 2008.
- [20] T. Okamoto and K. Takashima, “Hierarchical predicate encryption for inner-products,” in *Proc. of ASIACRYPT*, pp. 214–231, Springer-Verlag, 2009.
- [21] J. Katz, A. Sahai, and B. Waters, “Predicate encryption supporting disjunctions, polynomial equations, and inner products,” in *Proc. of EUROCRYPT*, pp. 146–162, Springer-Verlag, 2008.
- [22] A. Sahai and B. Waters, “Fuzzy identity-based encryption,” in *Proc. of EUROCRYPT*, pp. 457–473, 2005.
- [23] V. Goyal, O. Pandey, A. Sahai, and B. Waters, “Attribute-based encryption for fine-grained access control of encrypted data,” in *Proc. of ACM CCS*, pp. 89–98, ACM, 2006.
- [24] J. Bethencourt, A. Sahai, and B. Waters, “Ciphertext-policy attribute-based encryption,” in *Proc. of IEEE S&P*, pp. 321–334, IEEE Computer Society, 2007.
- [25] T. Okamoto and K. Takashima, “Fully secure functional encryption with general relations from the decisional linear assumption,” in *Proc. of CRYPTO*, pp. 191–208, 2010.
- [26] A. B. Lewko and B. Waters, “New proof methods for attribute-based encryption: Achieving full security through selective techniques,” in *Proc. of CRYPTO*, pp. 180–198, 2012.
- [27] M. Chase, “Multi-authority attribute based encryption,” in *Proc. of TCC*, pp. 515–534, Springer-Verlag, 2007.
- [28] M. Chase and S. S. Chow, “Improving privacy and security in multi-authority attribute-based encryption,” in *Proc. of ACM CCS*, pp. 121–130, ACM, 2009.
- [29] S. Benabbas, R. Gennaro, and Y. Vahlis, “Verifiable delegation of computation over large datasets,” in *Proc. of CRYPTO*, pp. 111–131, Springer-Verlag, 2011.
- [30] C. Papamanthou, E. Shi, and R. Tamassia, “Signatures of correct computation.” Cryptology ePrint Archive, Report 2011/587, 2011. <http://eprint.iacr.org/>.
- [31] D. Fiore and R. Gennaro, “Publicly verifiable delegation of large polynomials and matrix computations, with applications,” in *Proc. of ACM CCS*, pp. 501–512, ACM, 2012.
- [32] Q. Chai and G. Gong, “Verifiable symmetric searchable encryption for semi-honest-but-curious cloud servers,” in *Proc. of ICC*, pp. 917–922, 2012.
- [33] D. Boneh, X. Boyen, and E.-J. Goh, “Hierarchical identity based encryption with constant size ciphertext,” in *Proc. of EUROCRYPT*, pp. 440–456, 2005.
- [34] J. Katz and Y. Lindell, *Introduction to Modern Cryptography*. Chapman and Hall/CRC Press, 2007.
- [35] B. H. Bloom, “Space/time trade-offs in hash coding with allowable errors,” *Commun. ACM*, vol. 13, pp. 422–426, July 1970.
- [36] E. Shen, E. Shi, and B. Waters, “Predicate privacy in encryption systems,” in *Proc. of TCC*, pp. 457–473, Springer-Verlag, 2009.
- [37] “The java pairing based cryptography library. <http://gas.dia.unisa.it/projects/jpbc/>.”

APPENDIX A PROOF OF THEOREM 1

Proof: We show that if there is a polynomial-time adversary \mathcal{A} that wins the SCKA game with advantage μ , then there is a challenger algorithm that solves the DL problem with advantage $\mu/2$. Given a DL instance $(g, h, f, f^{r_1}, g^{r_2}, Q)$, where $g, f, h, Q \leftarrow G$ and $r_1, r_2 \leftarrow \mathbb{Z}_p$, the challenger simulates the SCKA game as follows.

Setup: The challenger sets $g^a = h$ and $g^c = f$ where a and c are unknown, selects $d \leftarrow \mathbb{Z}_p$ and computes $g^b = f^d = g^{cd}$ by implicitly defining $b = cd$. Let H_2 be an one-way hash function and $\text{pm} = (e, g, p, h, f^d, f)$ and $\text{mk} = (d)$.

\mathcal{A} selects an attribute set Atts^* and gives it to the challenger. The random oracle $\mathcal{O}_{H_1}(\text{at}_j)$ is defined as follows:

- If at_j has not been queried before,
 - if $\text{at}_j \in \text{Atts}^*$, select $\beta_j \leftarrow \mathbb{Z}_p$, add $(\text{at}_j, \alpha_j = 0, \beta_j)$ to \mathcal{O}_{H_1} , and return g^{β_j} ;
 - otherwise, select $\alpha_j, \beta_j \leftarrow \mathbb{Z}_p$, add $(\text{at}_j, \alpha_j, \beta_j)$ to \mathcal{O}_{H_1} , and return $f^{\alpha_j} g^{\beta_j}$.
- If at_j has been queried before, retrieve (α_j, β_j) from \mathcal{O}_{H_1} and return $f^{\alpha_j} g^{\beta_j}$.

Phase 1: \mathcal{A} can adaptively query the following oracles for polynomially-many times and the challenger keeps a keyword list L_{kw} , which is empty initially.

$\mathcal{O}_{\text{KeyGen}}(\mathbb{T})$: \mathcal{A} gives an access tree \mathbb{T} to the challenger. If $F(\text{Atts}^*, \mathbb{T}) = 1$, then the challenger aborts; otherwise, the challenger generates attributes as follows.

Define the following two procedures to determine the polynomial for each node of \mathbb{T} :

- $\text{PolySat}(\mathbb{T}_v, \text{Atts}^*, \lambda_v)$: Given secret λ_v , this procedure determines the polynomial for each node of \mathbb{T}_v rooted at v when $F(\text{Atts}^*, \mathbb{T}_v) = 1$. It works as follows: Suppose the threshold value of node v is k_v , it sets $q_v(0) = \lambda_v$ and picks $k_v - 1$ coefficients randomly to fix the polynomial q_v . For each child node v' of v , recursively call $\text{PolySat}(\mathbb{T}_{v'}, \text{Atts}^*, \lambda_{v'})$ where $\lambda_{v'} = q_v(\text{Index}(v'))$.
- $\text{PolyUnsat}(\mathbb{T}_v, \text{Atts}^*, g^{\lambda_v})$: Given element $g^{\lambda_v} \in G$ where the secret λ_v is unknown, this procedure determines the polynomial for each node of \mathbb{T}_v rooted at v when $F(\text{Atts}^*, \mathbb{T}_v) = 0$ as follows. Suppose the threshold value of the node v is k_v . Let V be the empty set. For each child node v' of v , if $F(\text{Atts}^*, \mathbb{T}_{v'}) = 1$, then set $V = V \cup \{v'\}$. Because $F(\text{Atts}^*, \mathbb{T}_v) = 0$, then

$|V| < k_v$. For each node $v' \in V$, it selects $\lambda_{v'} \leftarrow \mathbb{Z}_p$, and sets $q_v(\text{Index}(v')) = \lambda_{v'}$. Finally it fixes the remaining $k_v - |V|$ points of q_v randomly to define q_v and makes $g^{q_v(0)} = g^{\lambda_v}$. For each child node v' of v ,

- if $F(\text{Atts}^*, T_{v'}) = 1$, then run $\text{PolySat}(T_{v'}, \text{Atts}^*, q_v(\text{Index}(v')))$, where $q_v(\text{Index}(v'))$ is known to the challenger;
- otherwise, call $\text{PolyUnsat}(T_{v'}, \text{Atts}^*, g^{\lambda_{v'}})$, where $g^{\lambda_{v'}} = g^{q_v(\text{Index}(v'))}$ is known to the challenger.

With the above two procedures, the challenger runs $\text{PolyUnsat}(T, \text{Atts}^*, g^a)$, by implicitly defining $q_{\text{root}}(0) = a$. Then for each $v \in \text{lvs}(T)$, the challenger gets $q_v(0)$ if $\text{att}(v) \in \text{Atts}^*$, and gets $g^{q_v(0)}$ otherwise. Because $cq_v(0)$ is the secret share of ac , due to the linear property, the challenger generates credentials for each $v \in \text{lvs}(T)$ as follows:

- If $\text{att}(v) = \text{at}_j$ for some $\text{at}_j \in \text{Atts}^*$: Select $t \leftarrow \mathbb{Z}_p$, set $A_v = f^{q_v(0)} g^{\beta_j t} = g^{cq_v(0)} H_1(\text{att}(v))^t$ and $B_v = g^t$;
- If $\text{att}(v) \notin \text{Atts}^*$ (assuming $\text{att}(v) = \text{at}_j$): Select $t' \leftarrow \mathbb{Z}_p$, set $A_v = (g^{q_v(0)})^{\frac{-\beta_j}{\alpha_j}} (f^{\alpha_j} g^{\beta_j})^{t'}$ and $B_v = g^{q_v(0) \frac{-1}{\alpha_j}} g^{t'}$. Note that (A_v, B_v) is a valid credential because

$$\begin{aligned} B_v &= g^{q_v(0) \frac{-1}{\alpha_j}} g^{t'} = g^{t' - \frac{q_v(0)}{\alpha_j}} \\ A_v &= g^{q_v(0) \frac{-\beta_j}{\alpha_j}} (f^{\alpha_j} g^{\beta_j})^{t'} \\ &= f^{q_v(0) \frac{-\beta_j}{\alpha_j}} (f^{\alpha_j} g^{\beta_j})^{\frac{-q_v(0)}{\alpha_j}} (f^{\alpha_j} g^{\beta_j})^{t'} \\ &= f^{q_v(0)} (f^{\alpha_j} g^{\beta_j})^{t' - \frac{q_v(0)}{\alpha_j}} \\ &= g^{cq_v(0)} H_1(\text{att}(v))^{t' - \frac{q_v(0)}{\alpha_j}} \end{aligned}$$

by implicitly letting $t = t' - \frac{q_v(0)}{\alpha_j}$. Note also that \mathcal{A} cannot construct A_v and B_v without knowing α_j, β_j .

Eventually, the challenger returns $\text{sk} = \{(A_v, B_v) | v \in \text{lvs}(T)\}$ to \mathcal{A} .

$\mathcal{O}_{\text{TokenGen}}(T, w)$: The challenger runs $\mathcal{O}_{\text{KeyGen}}(T)$ to get $\text{sk} = (T, \{A_v, B_v | v \in \text{lvs}(T)\})$, computes $\text{tk} \leftarrow \text{TokenGen}(\text{sk}, w)$, and returns tk to \mathcal{A} . If $F(\text{Atts}, T) = 1$, the challenger adds w to the keyword List L_{kw} .

Challenge phase: \mathcal{A} chooses two keywords w_0 and w_1 of equal length, such that $w_0, w_1 \notin L_{kw}$. The challenger outputs cph^* as:

- Select $\lambda \leftarrow \{0, 1\}$.
- For each $\text{at}_j \in \text{Atts}^*$, set $W_j = (g^{r_2})^{\beta_j}$.
- Set $W' = f^{r_1}$, $W = Q(f^{r_1})^{dH_2(w_\lambda)}$, and $W_0 = g^{r_2}$.
- Set $\text{cph}^* = (\text{Atts}^*, W', W, W_0, \{W_j | \text{at}_j \in \text{Atts}^*\})$ and return cph^* to \mathcal{A} .

We note that if $Q = h^{r_1+r_2}$, then cph^* is indeed a legitimate ciphertext for keyword w_λ . The reason is that $W' = f^{r_1} = g^{cr_1}$, $W = Q(f^{r_1})^{dH_2(w_\lambda)} = Qg^{r_1cdH_2(w_\lambda)} = g^{a(r_1+r_2)} g^{br_1H_2(w_\lambda)}$, $W_0 = g^{r_2}$, and for $\text{at}_j \in \text{Atts}^*$, $W_j = (g^{r_2})^{\beta_j} = H_1(\text{at}_j)^{r_2}$.

Phase 2: \mathcal{A} continues to query the oracles as in Phase 1. The only restriction is that (T, w_0) and (T, w_1) cannot be the input to $\mathcal{O}_{\text{TokenGen}}$ if $F(\text{Atts}^*, T) = 1$.

Guess: Finally, \mathcal{A} outputs a bit λ' and gives it to the challenger. If $\lambda' = \lambda$, then the challenger outputs $Q = h^{r_1+r_2}$; otherwise, it outputs $Q \neq h^{r_1+r_2}$.

This completes the simulation. In the challenge phase, if $Q = h^{r_1+r_2}$, then cph^* is a valid ciphertext of w_λ , so the probability of \mathcal{A} outputting $\lambda = \lambda'$ is $\frac{1}{2} + \mu$. If Q is an element randomly selected from G , then cph^* is not a valid ciphertext of w_λ . The probability of \mathcal{A} outputting $\lambda = \lambda'$ is $\frac{1}{2}$ since W is a random element in G . Therefore, the probability of the challenger correctly guessing $Q \stackrel{?}{=} h^{r_1+r_2}$ with the DL instance $(g, h, f, f^{r_1}, g^{r_2}, Q)$ is $\frac{1}{2}(\frac{1}{2} + \mu + \frac{1}{2}) = \frac{1}{2} + \frac{\mu}{2}$. That is, the challenger solves the DL problem with advantage $\mu/2$ if \mathcal{A} wins the SCKA game with an advantage μ . ■

APPENDIX B PROOF OF THEOREM 2

Proof: We construct a challenger that exploits the keyword secrecy game as follows:

Setup: The challenger selects $a, b, c \leftarrow \mathbb{Z}_p$, $f \leftarrow G$. Let H_2 be an one-way hash function and $\text{pm} = (e, g, g^a, g^b, g^c, f)$ and $\text{mk} = (a, b, c)$.

The random oracle $\mathcal{O}_{H_1}(\text{at}_j)$ is simulated as follows: If at_j has not been queried before, the challenger selects $\alpha_j \leftarrow \mathbb{Z}_p$, adds (at_j, α_j) to \mathcal{O}_{H_1} , and returns g^{α_j} ; otherwise, the challenger retrieves α_j from \mathcal{O}_{H_1} and returns g^{α_j} .

Phase 1: \mathcal{A} can adaptively query the following oracles for polynomially-many times.

$\mathcal{O}_{\text{KeyGen}}(T)$: The challenger generates $\text{sk} \leftarrow \text{KeyGen}(T, \text{mk})$ and returns sk to \mathcal{A} . It adds T to the list L_{KeyGen} , which is initially empty.

$\mathcal{O}_{\text{TokenGen}}(T, w)$: The challenger runs $\mathcal{O}_{\text{KeyGen}}(T)$ to obtain $\text{sk} = (T, \{A_v, B_v | v \in \text{lvs}(T)\})$, computes $\text{tk} \leftarrow \text{TokenGen}(\text{sk}, w)$, and returns tk to \mathcal{A} .

Challenge Phase: \mathcal{A} selects an attribute set Atts^* . The challenger chooses an access control policy that is represented as T^* such that $F(\text{Atts}^*, T^*) = 1$, computes $\text{sk}^* \leftarrow \text{KeyGen}(\text{mk}, T^*)$. By taking as input Atts^* and sk^* , it selects w^* from keyword space uniformly at random, and computes cph^* and tk^* with Enc and TokenGen . Atts^* should satisfy the requirement defined in the keyword secrecy game.

Guess: Finally, \mathcal{A} outputs a keyword w' and gives it to the challenger. The challenger computes $\text{cph}' \leftarrow \text{Enc}(\text{Atts}, w')$ and if $\text{Search}(\text{tk}^*, \text{cph}') = 1$, then \mathcal{A} wins the game.

This finishes the simulation. Suppose \mathcal{A} has already attempted q distinct keywords before outputting w' , we can see that the probability of \mathcal{A} winning the keyword secrecy game is at most $\frac{1}{|\mathcal{M}|-q} + \epsilon$. This is because the size of the remaining keyword space is $|\mathcal{M}|-q$, and as the H_2 is an one way secure hash function, meaning deriving w^* from $H_2(w^*)$ is at most a negligible probability ϵ . Therefore, given q distinct keywords \mathcal{A} has attempted, the probability of \mathcal{A} winning the keyword secrecy game is at most $\frac{1}{|\mathcal{M}|-q} + \epsilon$. Thus, our scheme achieves keyword secrecy as in Definition ??.

APPENDIX C
PROOF OF THEOREM 3 ON CP-ABKS

Proof: We show that the CP-ABE scheme is selectively secure against chosen-keyword attack in the generic bilinear group model, where H_1 is modeled as a random oracle and H_2 is a one-way hash function.

In the SCKA game, \mathcal{A} attempts to distinguish $g^{a(r_1+r_2)}g^{br_1H_2(w_0)}$ from $g^{a(r_1+r_2)}g^{br_1H_2(w_1)}$. Given $\theta \leftarrow \mathbb{Z}_p$, the probability of distinguishing $g^{a(r_1+r_2)}g^{br_1H_2(w_0)}$ from g^θ is equal to that of distinguishing g^θ from $g^{a(r_1+r_2)}g^{br_1H_2(w_1)}$. Therefore, if \mathcal{A} has advantage ϵ in breaking the SCKA game, then it has advantage $\epsilon/2$ in distinguishing $g^{a(r_1+r_2)}g^{br_1H_2(w_0)}$ from g^θ . Thus, let us consider a modified game where \mathcal{A} can distinguish $g^{a(r_1+r_2)}$ from g^θ . The modified SCKA game is described as follows:

Setup: The challenger chooses $a, b, c \leftarrow \mathbb{Z}_p$ and sends public parameters (e, g, p, g^a, g^b, g^c) to \mathcal{A} . \mathcal{A} chooses an access tree T^* , which is sent to the challenger.

$H_1(\text{at}_j)$ is simulated as follows: If at_j has not been queried before, the challenger chooses $\alpha_j \leftarrow \mathbb{Z}_p$, adds (at_j, α_j) to O_{H_1} and returns g^{α_j} ; otherwise the challenger returns g^{α_j} by retrieving α_j from O_{H_1} .

Phase 1: \mathcal{A} can query $\mathcal{O}_{\text{KeyGen}}$ and $\mathcal{O}_{\text{TokenGen}}$ as follows:

$\mathcal{O}_{\text{KeyGen}}(\text{Atts})$: The challenger selects $r^{(t)} \leftarrow \mathbb{Z}_p$ and computes $A = g^{(ac+r^{(t)})/b}$. For each attribute $\text{at}_j \in \text{Atts}$, the challenger chooses $r_j^{(t)} \leftarrow \mathbb{Z}_p$, computes $A_j = g^{r^{(t)}}g^{\alpha_j r_j^{(t)}}$ and $B_j = g^{r_j^{(t)}}$, and returns $(\text{Atts}, A, \{(A_j, B_j)|\text{at}_j \in \text{Atts}\})$.
 $\mathcal{O}_{\text{TokenGen}}(\text{Atts}, w)$: The challenger queries $\mathcal{O}_{\text{KeyGen}}(\text{Atts})$ to get $\text{sk} = (\text{Atts}, A, \{(A_j, B_j)|\text{at}_j \in \text{Atts}\})$ and returns $\text{tk} = (\text{Atts}, \text{tok}_1, \text{tok}_2, \text{tok}_3, \{(A'_j, B'_j)|\text{at}_j \in \text{Atts}\})$ where $\text{tok}_1 = (g^a g^{bH_2(w)})^s$, $\text{tok}_2 = g^{cs}$, $\text{tok}_3 = A^s$, $A'_j = A_j^s$ and $B'_j = B_j^s$ by selecting $s \leftarrow \mathbb{Z}_p$. If $F(\text{Atts}, T^*) = 1$, the challenger adds w to the keyword List L_{kw} .

Challenge phase: Given two keywords w_0, w_1 of equal length where $w_0, w_1 \notin L_{kw}$, the challenger chooses $r_1, r_2 \leftarrow \mathbb{Z}_p$, and computes secret shares of r_2 for each leaves in T^* . The challenger selects $\lambda \leftarrow \{0, 1\}$. If $\lambda = 0$, it outputs

$$W = g^{cr_1}, W_0 = g^\theta, W' = g^{br_2}, \\ \{(W_v = g^{q_v(0)}, D_v = g^{\alpha_j q_v(0)})|v \in \text{lvs}(T^*), \text{att}(v) = \text{at}_j\}$$

by selecting $\theta \in \mathbb{Z}_p$; otherwise it outputs

$$W = g^{cr_1}, W_0 = g^{a(r_1+r_2)}, W' = g^{br_2}, \\ \{(W_v = g^{q_v(0)}, D_v = g^{\alpha_j q_v(0)})|v \in \text{lvs}(T^*), \text{att}(v) = \text{at}_j\}.$$

Phase 2: This is the same as in the SCKA game.

We can see that if \mathcal{A} can construct $e(g, g)^{\delta a(r_1+r_2)}$ for some g^δ that can be composed from the oracle outputs he has already queried, then \mathcal{A} can use it to distinguish g^θ from $g^{a(r_1+r_2)}$. Therefore, we need to show that \mathcal{A} can construct $e(g, g)^{\delta a(r_1+r_2)}$ for some g^δ with a negligible probability. That is, \mathcal{A} cannot gain non-negligible advantage in the SCKA game.

In the generic group model, ψ_0 and ψ_1 are random injective maps from \mathbb{Z}_p into a set of p^3 elements. Then the probability of \mathcal{A} guessing an element in the image of ψ_0 and ψ_1 is negligible.

a	$r_j^{(t)}$	$s(ac + r^{(t)})/b$	cr_1
b	$r^{(t)} + \alpha_j r_j^{(t)}$	$s(r_j^{(t)})$	$q_v(0)$
c	$(ac + r^{(t)})/b$	$s(r^{(t)} + \alpha_j r_j^{(t)})$	$\alpha_j q_v(0)$
α_j	cs	$s(a + bH_2(w))$	br_2

TABLE III
POSSIBLE TERMS FOR QUERYING GROUP ORACLE G_T

Recall that $G = \{\psi_0(x)|x \in \mathbb{Z}_p\}$ and $G_T = \{\psi_1(x)|x \in \mathbb{Z}_p\}$. Hence, let us consider the probability of \mathcal{A} constructing $e(g, g)^{\delta a(r_1+r_2)}$ for some $\delta \in \mathbb{Z}_p$ from the oracle outputs he has queried.

We list all terms that can be queried to the group oracle G_T in Table III. Let us consider how to construct $e(g, g)^{\delta a(r_1+r_2)}$ for some δ . Because r_1 only appears in the term cr_1 , δ should contain c in order to construct $e(g, g)^{\delta a(r_1+r_2)}$. That is, let $\delta = \delta'c$ for some δ' and \mathcal{A} wishes to construct $e(g, g)^{\delta'ac(r_1+r_2)}$. Therefore, \mathcal{A} needs to construct $\delta'acr_2$, which will use terms br_2 and $(ac+r^{(t)})/b$. Because $(br_2)(ac+r^{(t)})/b = acr_2 + r^{(t)}r_2$, \mathcal{A} needs to cancel $r^{(t)}r_2$, which needs to use the terms $\alpha_j, r^{(t)} + \alpha_j r_j^{(t)}, q_v(0)$ and $\alpha_j q_v(0)$ because $q_v(0)$ is the secret share of r_2 according to T^* . However, it is impossible to construct $r^{(t)}r_2$ with these terms because $r^{(t)}r_2$ only can be reconstructed if the attributes corresponding to $r_j^{(t)}$ of $r^{(t)} + \alpha_j r_j^{(t)}$ satisfies the access tree T^* .

Therefore, we can conclude that \mathcal{A} gains a negligible advantage in the modified game, which means that \mathcal{A} gains a negligible advantage in the SCKA game. This completes the proof. \blacksquare

APPENDIX D
PROOFS OF THEOREM 5, THEOREM 6, THEOREM 7 AND
THEOREM 8 ON VABKS

A. Proof of Theorem 5 on VABKS

Proof: We show that if there exists a polynomial-time algorithm \mathcal{A} breaks VABKS's data secrecy with the advantage ρ , then we can break either CPA security for ABE or CPA security for SE with the advantage $\frac{\rho}{n^2}$ where n is the number of data files to be encrypted.

The challenger proceeds the conventional CPA security game with \mathcal{A} . In the challenge phase, suppose \mathcal{A} presents two data collections $D_0 = (\text{KS}, \text{MP}, \text{FS}_0 = \{F_{01}, \dots, F_{0n}\})$, $D_1 = (\text{KS}, \text{MP}, \text{FS}_1 = \{F_{11}, \dots, F_{1n}\})$, $\{I_{\text{Enc}}\}_l$ and $\{I'_{\text{Enc}}\}_n$. The challenge selects $\lambda \leftarrow \{0, 1\}$ and encrypts FS_λ with the ABE and $\{I'_{\text{Enc}}\}_n$.

Now let us consider the advantage of \mathcal{A} correctly guessing λ . As we know, given two messages, the advantage of distinguishing which message was encrypted by the hybrid encryption of ABE and SE is equal. Therefore, given two sets of data files FS_0 and FS_1 , if the advantage of distinguishing which data set was encrypted is ρ , then the advantage of distinguishing which data file was encrypted is $\frac{\rho}{n^2}$ by selecting one data file from FS_0 and one from FS_1 .

Therefore, we can see that if \mathcal{A} breaks VABKS's data secrecy of with a non-negligible advantage ρ , then the advantage

of breaking CPA security for ABE or CPA security for SE is $\frac{\rho}{n^2}$ – a non-negligible probability, which contracts the assumption that ABE is CPA-secure and SE is CPA-secure. ■

B. Proof of Theorem 6 on VABKS

Proof: We show that if there exists a polynomial-time algorithm \mathcal{A} breaks the selective security against chosen-keyword attack of ABKS with the advantage ρ , then we can break the selective security against chosen-keyword attack game of ABKS with the advantage of $\frac{\rho}{l^2}$, given that ABE is CPA-secure and H is a secure pseudorandom generator.

The challenger proceeds selective security against chosen-keyword attack game with \mathcal{A} . In the challenge phase, suppose \mathcal{A} presents two data collections $D_0 = (KS_0 = \{KS_{01}, \dots, KS_{0l}\}, MP, FS)$, and $\{I'_{Enc}\}_n$. The challenge selects $\lambda \leftarrow \{0, 1\}$ and encrypts KS with ABKS, and generates BF'_i, cph_{BF_i} and σ_i for each keyword group.

Since ABE is CPA-secure and H is a secure pseudorandom generator, the probability of \mathcal{A} inferring λ via BF'_i, cph_{BF_i} is negligible. Then let us consider the advantage of \mathcal{A} correctly guessing λ from keyword ciphertexts. As we know, given two keywords, the advantage of distinguishing which keyword was encrypted by ABKS is equal. Therefore, given two keyword sets KS_0 and KS_1 , if the advantage of distinguishing which keyword set was encrypted is ρ , then the advantage of distinguishing which keyword was encrypted is $\frac{\rho}{l^2}$ by selecting one keyword from KS_0 and one from KS_1 .

Therefore, we can see that if \mathcal{A} breaks VABKS's selective security against chosen-keyword attack with a non-negligible advantage ρ , then the advantage of breaking ABKS's selective security against chosen-keyword attack is $\frac{\rho}{l^2}$ – a non-negligible probability, which contracts the assumption that ABKS achieve selective security against chosen-keyword attack, given that ABE is CPA-secure and H is a secure pseudorandom generator. ■

C. Proof of Theorem 7 on VABKS

Proof: We show that if there exists polynomial time algorithm \mathcal{A} breaking VABKS's keyword secrecy, then it breaks the assumption that ABKS achieves keyword secrecy.

Suppose \mathcal{A} presents a data collection $D = (KS = \{KS_1, \dots, KS_l\}, MP, FS)$, $\{I_{Enc}\}_l$ and $\{I'_{Enc}\}_n$. The challenger simulates the keyword secrecy game as in B, where the keyword space consists of keywords specified by FS . We can see that the probability of \mathcal{A} inferring the keyword from a search token and corresponding keyword ciphertext is equal to that of ABKS. Therefore, if in VABKS \mathcal{A} guesses the keyword from the search token and corresponding keyword ciphertext with the probability more than $\frac{1}{|\mathcal{M}|-q} + \epsilon$ after guessing q distinct keywords, then the probability of guessing the keyword from the search token and keyword ciphertext in ABKS is more than $\frac{1}{|\mathcal{M}|-q} + \epsilon$ after guessing q distinct keywords, which contracts the assumption that ABKS achieves keyword secrecy. ■

D. Proof of Theorem 8 on VABKS

Proof: We show that under the assumptions that Sig is unforgeable, any polynomial-time adversary \mathcal{A} presents an incorrect search result and succeeds in the verification with negligible probability.

The challenger proceeds the verifiability game, where \mathcal{A} provides the keyword-based data $D = (KS = \{KS_1, \dots, KS_l\}, MP = \{MP(w) | w \in \cup_{i=1}^l KS_i\}, FS = \{F_1, \dots, F_n\})$, $\{I_{Enc}\}_l$ and $\{I'_{Enc}\}_n$. The challenger runs $(Au, Index, D_{cph}) \leftarrow \text{BuildIndex}(\{I_{Enc}\}_l, \{I'_{Enc}\}_n, D)$, and gives $(Au, Index, D_{cph})$ to \mathcal{A} .

In the challenge phase, with w^* and I_{Enc}^* from \mathcal{A} , the challenger selects I_{KeyGen}^* such that $F(I_{KeyGen}^*, I_{Enc}^*) = 1$ where I_{Enc}^* is selected by \mathcal{A} , generates credential sk^* with I_{KeyGen}^* and returns to \mathcal{A} a search token tk^* by running $tk^* \leftarrow \text{TokenGen}(sk, w^*)$. \mathcal{A} returns $(rslt^*, proof^*)$ to the challenger.

Suppose that $(rslt^*, proof^*)$ succeeds in the verification. That is, $1 \leftarrow \text{Verify}(sk^*, w^*, tk^*, rslt^*, proof^*)$. Let us consider the probability of \mathcal{A} cheating with incorrect search result.

First, we claim that the global signature σ and random keyword ciphertexts $cph_{BF_1}, \dots, cph_{BF_l}$ are included in $proof^*$ without being manipulated; otherwise we can break the unforgeability of Sig.

Second, let us consider the search result within each group with respect to access control policies, i.e. $i = 1, \dots, l$:

- If there exists no keyword ciphertext matched the search token tk^* , then we claim that \mathcal{A} cannot cheat the challenger with some keyword ciphertext and data ciphertexts in order to make VABKS.Verify output 1. The reason is that \mathcal{A} cannot forge a keyword signature σ_w for the keyword ciphertext and data ciphertexts; otherwise, we can break the unforgeability of Sig.
- If there exists a keyword ciphertext matched the search token tk^* , then we claim that \mathcal{A} cannot cheat the challenger with a null search result in order to make VABKS.Verify output 1. Suppose \mathcal{A} returns a null result and the proof $(BF'_i, cph_{BF_i}, \sigma_{BF_i})$. Since BF'_i cannot be manipulated due to σ_{BF_i} , the unmasked bloom filter indicates that w^* is a member within the group. The challenger downloads $cph_{w_1}, \dots, cph_{w_{|KS_i|}}$ and σ_i without being manipulated; otherwise we break the Sig's unforgeability. Then the challenger can conduct the search operation with each keyword ciphertext, and VABKS.Verify will output 0. That is, if there exists keyword ciphertext matched the search token, \mathcal{A} returns a null result, then it cannot make VABKS.Verify output 1.

To sum up, in order to make VABKS.Verify output 1, \mathcal{A} has to faithfully execute search operations and return the search result honestly; otherwise, we will break Sig's unforgeability. ■