Extended Functionality in Verifiable Searchable Encryption*

James Alderman**, Christian Janson, Keith M. Martin,
and Sarah Louise Renwick***

Information Security Group, Royal Holloway, University of London
Egham, Surrey, TW20 0EX, United Kingdom
{James.Alderman, Keith.Martin}@rhul.ac.uk
{Christian.Janson.2012, SarahLouise.Renwick.2012}@live.rhul.ac.uk

Abstract. When outsourcing the storage of sensitive data to an (untrusted) remote server, a data owner may choose to encrypt the data beforehand to preserve confidentiality. However, it is then difficult to efficiently retrieve specific portions of the data as the server is unable to identify the relevant information. Searchable encryption has been well studied as a solution to this problem, allowing data owners and other authorised users to generate search queries which the server may execute over the *encrypted* data to identify relevant data portions.
However, many current schemes lack two important properties: verifiability of search results, and expressive queries. We introduce Extended Verifiable Searchable Encryption (eVSE) that permits a user to verify that search results are correct and complete. We also permit verifiable computational queries over keywords and specific data values, that go beyond the standard keyword matching queries to allow functions such as averaging or counting operations. We formally define the notion of eVSE within relevant security models and give a provably secure instantiation.

1 Introduction

It is now common for data owners to outsource their data to public servers providing storage on a pay-as-you-go basis. This can reduce the costs of data storage compared with that of running a private data center (e.g. hardware, construction, air conditioning and security costs), making this a cost effective solution. If the server is not fully trusted and the data is of a sensitive nature, the data owner may wish to encrypt it to ensure confidentiality. This, however, prevents the efficient retrieval of specific portions of the data as the server is unable to identify the relevant information.

* The final publication is available at Springer via http://dx.doi.org/10.1007/978-3-319-29172-7_12.
** Supported by the European Commission under project H2020-644024 "CLARUS" and acknowledges support from BAE Systems Advanced Technology Centre.
*** Supported by Thales UK and EPSRC under a CASE Award.

Searchable Encryption (SE) [11, 15, 18, 20, 21, 23, 25, 32] addresses this issue by indexing the encrypted data in such a way as to allow a server to execute a search query (formed by the data owner or an authorised data user) over the encrypted data and return the identifiers of any file that satisfies the query.

To preserve confidentiality of the data, the server must not learn anything about the underlying data from the encrypted data and the data indexes; namely *ciphertext indistinguishability* and *index indistinguishability*. In the presence of a search query the only information leaked to the server is the search results. *Query indistinguishability* is also a desirable property although, due to the *offline keyword guessing attack* [12], this is not always easy to achieve in the public key setting (where indexes are generated using the data owner's public key).

The majority of existing work on SE focusses on efficiently preserving confidentiality in the presence of an *honest-but-curious* server. This means that the server is trusted to follow the search protocol honestly but may try to infer information about data or search queries that it is unauthorised to know.

Verifiable Searchable Encryption (VSE) [13, 24, 30, 35, 37] assumes a stronger *semi-honest-but-curious* adversarial model in which the server might execute only a fraction of the search, or return a fraction of the search results in order to preserve its resources. To ensure the completeness and correctness of search results in this scenario, it is required that the server is able to prove to the querier that the search was computed honestly.

The current approaches to VSE in the literature do not support a wide range of expressive search queries. We address this issue by extending the types of queries that can be executed and verified by a VSE scheme to include more expressive search queries, as well as some computations. Most VSE schemes in the literature also require that the verification of query results be performed by the entity that issued the query whereas eVSE is publicly and blindly verifiable.

1.1 Our Contributions

We adapt and apply new techniques from the area of Publicly Verifiable Outsourced Computation to VSE in a novel way to enable a wider family of queries, and some types of computations, to be performed over outsourced encrypted data with verifiable query results. In summary, our contributions are:

- More expressive queries: Our scheme supports queries such as boolean formulae involving conjunctions, disjunctions and negations, threshold operations, polynomials, arbitrary CNF and DNF formulae, and fuzzy search[1].
- Evaluation of computations: Our scheme supports the evaluation of some computations over the encrypted data, such as averaging and counting operations. As well as assigning keywords to label data, we propose to also assign keywords representing certain data values that may be computed over (either in the form of single keywords or as a string of keywords encoding binary data, see Section 3.3).

[1] Depending on the choice of underlying ABE scheme; see Section 4.1.

– Blind public verifiability of query results: Any entity is able to verify the correctness and completeness of query results without any knowledge of either the underlying query or the results themselves.

The remainder of this paper is organised as follows. Section 2 gives some background information on SE and verifiable computation. Section 3 formally defines eVSE and its security model, Section 4 gives an instantiation of eVSE and Section 5 concludes the paper, highlighting possible avenues of future research. The Appendix provides more details on the security models and gives a security proof sketch as well as a discussion comparing our scheme with the ones in the literature. Additional details can be found in the full version [3].

2 Background

Searchable encryption (SE) allows data to be outsourced in encrypted form and for keyword search queries to be performed remotely. Methods based on oblivious RAM [19] provide a high level of security (hiding both the access and search patterns) at the expense of slow search times and high communication costs. Song et al. [28] achieve a scheme with fewer rounds of communication, but which leaks the access pattern and requires each word of a document to be encrypted separately, so compression is not possible. Goh [18] introduced meta data (indexes) describing the content of each document, and enabled constant time searches using Bloom filters over the index only. Curtmola et al. [15] extended the system model to allow multiple users to query the data, using broadcast encryption to manage user access privileges. SE schemes that allow many users to upload data can be built using public key encryption, however the data can only be searched by the holder of the corresponding secret key (or a derivative thereof) [11]. Most SE schemes assume an honest-but-curious server model.

Verifiable searchable encryption (VSE) schemes assume a semi-honest-but-curious server model. The first VSE scheme was presented by Chai et al. [13], where they extend the paradigm of searchable symmetric encryption (SSE) [15] to create a verifiable SSE (VSSE) scheme that allows verification of search results from a single keyword equality query. Another approach by [24] extends a public key encryption with keyword search scheme [11] to support verification of search results from a single keyword equality query, where the indexes are created using a public key. Sun et al. [30] and Wang et al. [34] detail VSE schemes with enhanced functionality; verifiable multi-keyword ranked search and verifiable fuzzy keyword search, respectively.

Verifiable Computation (VC) allows a client with limited resources to efficiently outsource a computation to a more powerful server, and to verify the correctness of results. Gennaro et al. [17] considered the use of garbled circuits, whilst Parno et al. [26] introduced *publicly* verifiable computation (PVC) built from key policy attribute based enryption (KP-ABE), where a single client computes an evaluation key for the server and publishes information enabling other clients to outsource computation to the server. *Any* client may verify the correctness of a result. Alderman et al. [2] considered an alternative system model that

used ciphertext policy attribute based encryption (CP-ABE) to allow clients to query computations on data held by the server (or initially outsourced by a client) called *Verifiable Delegable Computation* (VDC). This can naturally be applied to problems like querying on remote data, as well as MapReduce. Data remains statically stored on the server and may be embedded in a server's secret key, whilst the computation of many different functions can be requested by creating ciphertexts using *only* public information. Other notable approaches in the realm of querying remote data can be found in [4–6, 8–10, 14].

3 Extended Verifiable Searchable Encryption

3.1 System Model

We consider a system comprising a *data owner*, a remote storage *server*, and a set of authorised *data users*. The data owner sets up the system to generate a master secret and holds a set of data D (e.g. a database) that they wish to encrypt and outsource to the remote server. The data owner controls which additional users are able to query their encrypted data. Queries may be formulated over these keywords (e.g. to identify records associated with a given set of keywords) as usual in SE, but we also allow *computational queries* of functions in the class NC^1, which consists of Boolean functions computable by circuits of depth $\mathcal{O}(\log n)$ where each gate has a fan-in of two, over encoded data values.

For example consider workgroups within an organisation. The manager or system administrator acts as the data owner for the organisation and outsources a shared database to a remote server. Authorisation is granted by issuing a secret key to each user, which is required when creating a query token QT_Q for a particular query Q. The token is sent to the server who performs the query on the encoded index to generate a result R. We allow *any* entity to verify the correctness and completeness of the result[2], but we restrict the ability to read the value of the result to only authorised data users (holding a retrieval key).

Throughout this work, we assume a strict separation between *queriers* (the data owner and users) and the remote server – the server may not issue queries itself, else it will trivially be able to learn the encoding of the index and queries (legitimate queriers must know this encoding to gain meaningful results).

3.2 Formal Definition

We now formally define a scheme for eVSE. We use the following notation. Data to be outsourced is denoted D and is considered to be a collection of n documents. Prior to outsourcing, the data owner specifies a *pre-index* for D, denoted $\delta(D)$, which assigns a set of descriptive labels to each document e.g. keywords contained in the document or specific data values that may be computed upon. The encoded form of the data, including the descriptive labels,

[2] We also permit the server to verify correctness to avoid the rejection problem, where a server may learn some useful information by observing if results are accepted.

is referred to as the *index* of D, denoted $\mathcal{I}_{\mathcal{D}}$, and is stored by the server. Queries for functions in the class NC^1 are denoted by Q and to make such a query, a data user creates a query token QT_Q for Q, a verification key VK_Q which allows *any* entity to blindly verify the result, R, of the query, and a retrieval key RK_Q which is issued to authorised data users to enable the query result to be learnt.

Definition 1. *An* Extended Verifiable Searchable Encryption (eVSE) scheme *comprises the following algorithms:*

- $(\mathrm{MK}, \mathrm{PP}) \xleftarrow{\$} \mathsf{Setup}(1^\kappa, \mathcal{U})$: Run by the data owner and takes as input the security parameter and a universe of attributes (keywords and data values). It outputs the data owner's master secret key MK that is used for further administrative tasks and public parameters PP, both of which are provided to the remaining algorithms where required.
- $(\mathcal{I}_D, st_s, st_o) \xleftarrow{\$} \mathsf{BuildIndex}(\delta(D), G, \mathrm{MK}, \mathrm{PP})$: Run by the data owner and takes as input the pre-index of the data $\delta(D)$ and the set G of authorised users, and outputs a searchable index $\mathcal{I}_D$ for the data D, as well as a server and data owner state.
- $(SK_{\mathrm{ID}}, st_s) \xleftarrow{\$} \mathsf{AddUser}(\mathrm{ID}, G, \mathrm{MK}, \mathrm{PP})$: Run by the data owner to authorise a user ID to perform queries by issuing them a secret key SK_{ID} and outputs an updated server state.
- $(QT_Q, VK_Q, RK_Q) \xleftarrow{\$} \mathsf{Query}(Q, st_s, st_o, SK_{\mathrm{ID}}, \mathrm{PP})$: Run by a data user using its secret key and both states to generate a query token QT_Q for a query Q, a verification key VK_Q and an output retrieval key RK_Q.
- $R \xleftarrow{\$} \mathsf{Search}(\mathcal{I}_D, QT_Q, st_s, SK_{\mathrm{S}}, \mathrm{PP})$: Run by the server to execute a query given in the query token QT_Q on the index $\mathcal{I}_D$. It generates a result R which can be returned to the querying user or published.
- $r \leftarrow \mathsf{Verify}(R, VK_Q, RT_Q, RK_Q, \mathrm{PP})$: Verification consists of two steps:
  1. $RT_Q \leftarrow \mathsf{BVerif}(R, VK_Q, \mathrm{PP})$: Run by *any* party to verify the correctness and completeness of the result R. It takes the verification key VK_Q and, if the result is accepted, it outputs a retrieval token RT_Q which can be used to learn the result. Otherwise a distinguished failure symbol $RT_Q = \perp$ is returned.
  2. $r \leftarrow \mathsf{Retrieve}(VK_Q, RT_Q, RK_Q, \mathrm{PP})$: Run by a data user to read the value of the result. It takes as input the retrieval token RT_Q, the retrieval key RK_Q and the user's secret key. If the user holds a valid retrieval key for Q and the computation was performed correctly, then it returns the actual result $r = Q(\mathcal{I}_D)$, otherwise it returns $r = \perp$.
- $(st_s, st_o) \xleftarrow{\$} \mathsf{RevokeUser}(\mathrm{ID}, G, \mathrm{MK}, \mathrm{PP})$: Run by the data owner using its master secret key to revoke a user's authorisation to make queries and read results. It does so by updating the server and data owner state.

An eVSE is *correct* if there is a negligible probability that verification does not suceed when all algorithms are run honestly. A formal definition can be found in [3].

3.3 Types of Query

We consider a broader range of verifiable queries than many prior schemes. In particular, we consider two main types:

- **Keyword matching queries**: Queries of this type have formed the basis of most prior work in SE. Suppose there exists a universe (dictionary) of keywords. Each encrypted data item is associated with an *index* of one or more keywords to describe the contents. Queries are formed over the same universe of keywords. In this work, we permit Boolean formulae over sets of keywords (e.g. $((\mathtt{a} \wedge \mathtt{b}) \vee \mathtt{c})$ where $\mathtt{a}, \mathtt{b}, \mathtt{c}$ are keywords). We return an identifier for each file whose associated keywords in the index satisfy this formula. Thus we can perform very expressive search queries over keywords.
- **Computational queries**: Queries of this type are similar to the operations commonly discussed in the context of outsourced computation. We allow statistical queries over keywords (e.g. counting the number of data items that satisfy a keyword matching query), as well as operations over selected data values that have been encoded using additional portions of the keyword universe. It is possible to encode the entire database in such a way as to enable computations over all data fields, but it would usually be more efficient to select a (small) subset of fields that are most useful or most frequently queried. Clearly, keyword matching queries can be seen as a special case of computational queries where the function operator is equality testing.
- **Mixed queries**: Queries of this type combine both the functionalities of the aforementioned query types (e.g. finding the average of data values contained in all documents associated with a particular keyword).

All types of query are performed in a verifiable manner to ensure that results are correct and complete.

3.4 Security Model

We now formalise several notions of security as a series of cryptographic games. The adversary against each notion is modelled as a probabilistic polynomial time (PPT) algorithm $\mathcal{A}$ run by a challenger, with input parameters chosen to represent the knowledge of a real attacker as well as the security parameter κ. The adversary algorithm may maintain state and be multi-stage; we refer to each stage as $\mathcal{A}$ for ease of notation. The notation $\mathcal{A}^{\mathcal{O}}$ denotes the adversary being provided with oracle access to the following algorithms: $\mathsf{BuildIndex}(\cdot, \cdot, \mathrm{MK}, \mathrm{PP})$, $\mathsf{AddUser}(\cdot, \cdot, \mathrm{MK}, \mathrm{PP})$, $\mathsf{Query}(\cdot, \cdot, \cdot, \cdot, \mathrm{PP})$ and $\mathsf{Search}(\cdot, \cdot, \cdot, \cdot, \mathrm{PP})$. We assume that oracle queries are performed in a logical order such that all required information is generated from previous queries. For each game, we define the *advantage* and *security* of $\mathcal{A}$ as:

Definition 2. *The* advantage *of a PPT adversary $\mathcal{A}$ is defined as follows, where $\mathbf{X} \in \{PubVerif, IndPriv, QueryPriv\}$:*

$$Adv^{\mathbf{X}}_{\mathcal{A}}(e\mathcal{VSE}, 1^{\kappa}) = \Pr[\mathbf{Exp}^{\mathbf{X}}_{\mathcal{A}}[e\mathcal{VSE}, 1^{\kappa}] = 1].$$

Game 1 $\mathbf{Exp}_{\mathcal{A}}^{PubVerif}\,[e\mathcal{VSE}, 1^{\kappa}]$:

$(Q, \delta(D^{\star})) \leftarrow \mathcal{A}(1^{\kappa})$
$(\mathrm{PP}, \mathrm{MK}) \leftarrow \mathsf{Setup}(1^{\kappa}, \mathcal{U})$
$G \leftarrow \emptyset$
$\mathrm{ID} \xleftarrow{\$} \mathsf{Users}$
$(SK_{\mathrm{ID}}, st_s) \leftarrow \mathsf{AddUser}(\mathrm{ID}, G, \mathrm{MK}, \mathrm{PP})$
$(\mathcal{I}_{D^{\star}}, st_s, st_o) \leftarrow \mathsf{BuildIndex}(\delta(D^{\star}), G, \mathrm{MK}, \mathrm{PP})$
$(QT_Q, VK_Q, RK_Q) \leftarrow \mathsf{Query}(Q, st_s, st_o, SK_{\mathrm{ID}}, \mathrm{PP})$
$R^{\star} \leftarrow \mathcal{A}^{\mathcal{O}}(QT_Q, VK_Q, RK_Q, \mathcal{I}_{D^{\star}}, \mathrm{PP})$
$RT_Q \leftarrow \mathsf{BVerif}(R^{\star}, VK_Q, \mathrm{PP})$
$r \leftarrow \mathsf{Retrieve}(VK_Q, RT_Q, RK_Q, \mathrm{PP})$
if $(r \neq \perp)$ **and** $(r \neq Q(\mathcal{I}_{D^{\star}}))$ **then return** 1
else return 0

An eVSE scheme is secure against Game $\mathbf{X}$ *if for all PPT adversaries* $\mathcal{A}$, $Adv_{\mathcal{A}}^{\mathbf{X}}(e\mathcal{VSE}, 1^{\kappa}) \leq \mathrm{negl}(\kappa)$ *where* negl *is a negligible function.*

Public Verifiability. In Game 1, we capture the notion of *public verifiability* such that a server may not cheat by returning an incorrect result without being detected. This is a selective notion of security where, at the beginning of the game, the adversary chooses the challenge query and pre-index. The challenger then initialises the system, runs AddUser for a randomly chosen ID from the userspace, runs BuildIndex for the challenge pre-index to create the index, and finally runs Query. The adversary is given the resulting parameters, as well as access to the above specified oracle queries, and outputs $R^{\star}$, which it believes to be an incorrect result that will, nevertheless, be accepted by the verifier. The challenger runs the verification steps on this output. The adversary wins if verification succeeds, yet the result is not $Q(\mathcal{I}_{D^{\star}})$.

Index Privacy and Query Privacy. In Appendix A, we provide notions of *index indistinguishability* against a selective chosen keyword attack and *query privacy*, which ensure that no information regarding the keywords is leaked from the index or query tokens respectively.

4 Construction

4.1 Overview

We base our instantiation on a CP-ABE scheme. As shown by Alderman et al. [2], CP-ABE can be used to verifiably request computations to be performed on data held by a server, referred to as VDC. In VDC, a trusted Key Distribution Center (KDC) initialises the system and issues a CP-ABE decryption key to the server pertaining to the data it holds. We use a similar technique, but have the data owner act as the KDC (so the data need not be revealed to an external KDC, as in VDC). The index for a set of data is a CP-ABE decryption key for a set of attributes encoding the pre-index, and is sent to the server. The method of encoding is described in Section 4.2.

We consider the family $\mathcal{B}$ of Boolean functions closed under complement – that is, if $F \in \mathcal{B}$ then $\overline{F}$, where $\overline{F}(x) = F(x) \oplus 1$, is also in $\mathcal{B}$. A function $F : \{0,1\}^n \to \{0,1\}$ is *monotonic* if $x \leqslant y$ implies $F(x) \leqslant F(y)$, where $x = (x_1, \ldots, x_n) \leq y = (y_1, \ldots, y_n)$ if and only if $x_i \leqslant y_i$ for all i. For a monotonic function F, the set $\mathbb{A}_F = \{x : F(x) = 1\}$ defines a monotonic access structure.

A query Q is represented as a Boolean function of keywords and computational data points. If a monotonic CP-ABE scheme is used then queries can be comprised of `AND` and `OR` gates (and negation can inefficiently be handled by including both a positively and negatively labelled attribute in the universe and requiring the presence of exactly one of them). A non-monotonic CP-ABE scheme enables queries formed from `AND`, `OR` and `NOT` gates, which is a universal set of gates, and fuzzy CP-ABE enables fuzzy keyword search. We can achieve all functions in the class NC^1, which includes common arithmetic and comparison operators useful in queries. An n-bit result can be formed by performing n Boolean queries, each of which returns the i^{th} bit of the output.

The query token for a Boolean function $Q \in \mathcal{B}$ comprises two CP-ABE ciphertexts for access structures representing Q and $\overline{Q} \in \mathcal{B}$ respectively. To perform the search, the server attempts to decrypt each ciphertext under the secret key (associated with the pre-index) and outputs the result. Each decryption succeeds if and only if the query evaluates to True on the index. Any entity may perform the blind verification operation using the verification key to learn only whether the operation was performed correctly or not. Only entities holding the retrieval token can read the value of the result.

4.2 Data Encoding

Defining the Index. Suppose the data D to be outsourced comprises n documents. We now discuss how to form a *pre-index* $\delta(D)$, which represents the keywords and data fields that may be queried over.

Let $\mathcal{D}$ be a dictionary of keywords that describe the documents. $\mathcal{D}$ alone suffices for keyword matching queries but for computational queries, we also need to be able to encode data values such that they can be input to queries represented as access structures encoding Boolean functions.

For each data field x that may be input to a computational query, let the maximum size of the data value be m_x bits. We define m_x additional attributes $A_{x,1}, A_{x,2}, \ldots, A_{x,m_x}$, and define the universe $\mathcal{C} = \bigcup_{x \in D} \cup_{i=1}^{m_x} A_{x,i}$ to be the union of these attributes over all data fields. Let y be a value stored in the data field x and let the binary representation of y be $y_1, \ldots, y_{m_x}$. We view y as a *characteristic tuple* of an attribute set $A_y \subseteq \mathcal{C}$, where $A_y = \{A_{x,i} : y_i = 1\}$ – we include an attribute for position i in the set if and only if the i^{th} bit of y is 1.

Finally, to enable the index for all n documents to be encoded within a single CP-ABE key (and hence for computations to be performed simultaneously on all documents), and to ensure that the correct index data is used for each query, we must encode a labelling of the document that each attribute pertains to. We define our attribute universe $\mathcal{U}$ for the CP-ABE scheme to be $\mathcal{U} = \{\mathcal{D} \cup \mathcal{C}\} \times [n]$. Thats is, we take n copies of $\mathcal{D}$ and $\mathcal{C}$. Each element of $\{\mathcal{D} \cup \mathcal{C}\}$ describes a

particular keyword or data value, and each copy relates to a different document in D - if we index each copy of an attribute $w \in \{\mathcal{D} \cup \mathcal{C}\}$ as $\{w_i\}_{i=1}^n$, then w_i denotes the presence of w in document i. In practice, it may be desirable to use a 'large universe' CP-ABE scheme, wherein arbitrary textual strings are mapped to attributes (group elements), e.g. using a hash function H. Thus, for a keyword or data value w in document i, the attribute could be defined as $\mathsf{H}(w||i)$.[3]

The pre-index of the data D is a set of attributes $\delta(D) \subseteq \mathcal{U}$. The index that is outsourced will be a CP-ABE key generated over this attribute set.

Hiding the Index. In general, CP-ABE schemes do not hide the attributes within the decryption key. This is usually expected behaviour since CP-ABE is often used to cryptographically enforce access control policies and it is natural to assume that an entity is aware of their access rights.

However, in this setting we are using CP-ABE not to protect objects from unauthorised access, but instead to prove the outcome of a function evaluation. The keys in our setting are formed over attributes encoding the index of outsourced data, as opposed to encoding access rights. Since the server should not learn any information about the data, *including* the index, we must implement a mechanism by which the decryption key hides the associated attributes.

In many CP-ABE schemes, the public parameters comprise an ordered set of group elements [36], each associated with an attribute from the universe; that is, $\forall i \in \mathcal{U}$, choose $t_i \xleftarrow{\$} \mathbb{Z}_p$, then form the encoded attribute set $\{g^{t_i}\}_{i \in \mathcal{U}}$. Thus, given a key (or ciphertext) that comprises g^{t_i}, it is possible, based on the ordering of this set, to determine the attribute $i \in \mathcal{U}$ it relates to. In addition, the attributes may be listed in the clear, and attached to keys and ciphertexts to indicate which group elements should be applied at each point. Clearly, this is unsuitable for our requirement for a hidden index.

To this end, we first apply a random permutation to $\mathcal{U}$ such that the position of the group elements within the ordered set does not reveal the attribute string (unless the permutation is known). We then use a symmetric encryption scheme to encrypt each attribute $x \in \mathcal{U}$ under a key k, and then instantiate the CP-ABE scheme on this universe of *encrypted* attributes. Thus, without knowledge of the key k, the server should be unable to determine the attribute string x. We assume that only the keywords or data items being computed over are considered sensitive, and not the logical makeup of the Boolean function (in terms of gates).

4.3 Formal Details

The data owner initialises the system and encodes the data as an index which is pushed to the server. Each (authorised) user will be issued with a personalised secret key enabling them to form queries. To make a query Q, a user chooses a random message from the message space $\mathcal{M}$ to act as a verification token, and

[3] In this case, it may be possible to avoid the use of symmetric encryption in our construction by letting the secret k be the key for this cryptographic hash function.

encrypt this using the CP-ABE scheme under the access structure encoding Q. The server attempts to decrypt the ciphertext and recovers the chosen message if and only if $Q(\mathcal{I}_D) = 1$. By the indistinguishability security of the CP-ABE scheme, the server learns nothing about the message if $Q(\mathcal{I}_D) = 0$ since this corresponds to an access structure not being satisfied. Thus, if a server returns the correct message, the user is assured that the query evaluated to 1 on the data. If, however, $Q(\mathcal{I}_D) = 0$, then decryption will return $\perp$. This is insufficient for verification purposes since the server can return $\perp$ to convince a user of a false negative search result. Thus, the user must, in fact, produce two CP-ABE ciphertexts. As above, one corresponds to the function Q, whilst the other corresponds to $\overline{Q}$, the complement query of Q. Hence, the server's key will decrypt *exactly one* ciphertext and the returned message will distinguish whether Q or $\overline{Q}$ was satisfied, and therefore the value of $Q(\mathcal{I}_D)$. A well-formed response (d_0, d_1) from a server, therefore, satisfies the following:

$$(d_0, d_1) = \begin{cases} (m_0, \perp), & \text{if } Q(\mathcal{I}_D) = 1 \\ (\perp, m_1), & \text{if } Q(\mathcal{I}_D) = 0. \end{cases} \tag{1}$$

Public Verifiability is achieved by publishing a token comprising a one-way function g applied to both plaintexts. Any entity can apply g to the server's response and compare with this token to check correctness. To achieve blind verification, a random bit b permutes the order of the ciphertexts. Thus, verifiers that do not know b cannot determine whether a plaintext is associated with Q or $\overline{Q}$.

Our adversarial model allows the adversary (and hence servers in our system) to hold more than one key (for multiple datasets); we must ensure that a key cannot produce a valid looking response to a query on a different index. We achieve this by labelling each pre-index with a label $l(\delta(D))$ and define an attribute for each label. Then, for a pre-index $\delta(D)$, the decryption key is formed over the attribute set $(\delta(D) \cup l(\delta(D)))$. Recall that encoded data stored on the server's side is a collection of n documents, which we label $D_1, \ldots, D_n$. When making a query $Q(\mathcal{I}_D)$, a sub-query Q_i may be formed for each document (e.g. to check if a given keyword is contained in each document). In this case, the encryption algorithm takes the access structure encoding of the conjunction $(D_i \wedge l(\delta(D)))$ for $i \in [n]$. A valid result can only be formed by applying the sub-query to the specified document, which is also labelled by $D_i \in \mathcal{D}$ – decryption succeeds if and only if the function is satisfied *and* the label $l(\delta(D))$ is matched in the key and ciphertext. Note that a key for a different pre-index will not include the correct label. Inputs to the Query algorithm are assumed to be in this form.

Let $\mathcal{CPABE}$ = (ABE.Setup, ABE.KeyGen, ABE.Encrypt, ABE.Decrypt) define a CP-ABE encryption scheme over the universe $\mathcal{U}$. Let $\mathcal{SE}$= (SE.KeyGen, SE.Encrypt, SE.Decrypt) be an authenticated symmetric encryption scheme secure [7] in the sense of IND-CPA. Let $\mathcal{BE}$ = (BE.KeyGen, BE.Encrypt, BE.Add, BE.Decrypt) be a broadcast encryption scheme that retains IND-CPA security against a coalition of revoked users. Finally, let g be a one-way function and let Π and ϕ be pseudo-random permutations (PRPs) (which pad their inputs if required). Then Algorithms 1–8 define an eVSE scheme for a class of queries $\mathcal{Q}$.

Alg. 1 $(\mathrm{MK}, \mathrm{PP}) \leftarrow \mathsf{Setup}(1^\kappa, \mathcal{U})$

$mk \leftarrow \mathsf{BE.KeyGen}(1^\kappa)$
$k \leftarrow \mathsf{SE.KeyGen}(1^\kappa)$
for $i \in \mathcal{U}$ **do**
    $u_i \leftarrow \mathsf{SE.Encrypt}(i, k)$
$\mathcal{U}' \leftarrow \{u_i\}_{i \in \mathcal{U}}$
$\tilde{\mathcal{U}} \leftarrow \Pi(\mathcal{U}')$
$(MSK_{\mathrm{ABE}}, MPK_{\mathrm{ABE}}) \leftarrow \mathsf{ABE.Setup}(1^\kappa, \tilde{\mathcal{U}})$
$\mathrm{PP} \leftarrow (MPK_{\mathrm{ABE}}, \tilde{\mathcal{U}})$
$\mathrm{MK} \leftarrow (MSK_{\mathrm{ABE}}, mk, k, \Pi)$

Alg. 2 $(\mathcal{I}_D, st_s, st_o) \leftarrow \mathsf{BuildIndex}(\delta(D), G, \mathrm{MK}, \mathrm{PP})$

$\mathcal{I}_D \leftarrow \mathsf{ABE.KeyGen}((\delta(D) \cup l(\delta(D))), MSK_{\mathrm{ABE}}, MPK_{\mathrm{ABE}})$
$j \xleftarrow{\$} \{0,1\}^\kappa$
$st_s \leftarrow \mathsf{BE.Encrypt}(G, j, mk)$
$st_o \leftarrow j$

Alg. 3 $(SK_{\mathrm{ID}}, st_s) \leftarrow \mathsf{AddUser}(\mathrm{ID}, G, \mathrm{MK}, \mathrm{PP})$

$uk_{\mathrm{ID}} \leftarrow \mathsf{BE.Add}(\mathrm{ID}, \mathrm{mk})$
if ID is a user **then** $SK_{\mathrm{ID}} \leftarrow (uk_{\mathrm{ID}}, k, \Pi)$
else $SK_{\mathrm{ID}} \leftarrow uk_{\mathrm{ID}}$
$st_s \leftarrow \mathsf{BE.Encrypt}(G \cup \mathrm{ID}, \mathrm{j}, \mathrm{mk})$

Alg. 4 $(QT_Q, VK_Q, RK_Q) \leftarrow \mathsf{Query}(Q = \{Q_i\}, st_s, st_o, SK_{\mathrm{u}}, \mathrm{PP})$

$\widetilde{j} \leftarrow \mathsf{BE.Decrypt}(st_s, uk_{\mathrm{ID}})$
if $(\widetilde{j} \neq st_o)$ **then return** $\perp$
for $i = 1$ **to** $|Q|$ **do**
    $(m_{0_i}, m_{1_i}) \xleftarrow{\$} \mathcal{M} \times \mathcal{M}$
    $b_i \xleftarrow{\$} \{0,1\}$
    $c_{b_i} \leftarrow \mathsf{ABE.Encrypt}(m_{b_i}, Q_i, MPK_{\mathrm{ABE}})$
    $c_{1-b_i} \leftarrow \mathsf{ABE.Encrypt}(m_{1-b_i}, \overline{Q_i}, MPK_{\mathrm{ABE}})$
    $QT_{Q_i} \leftarrow (c_{b_i}, c_{1-b_i})$
    $\gamma_i \leftarrow \phi_j(c_{b_i} \| c_{1-b_i})$
    $VK_{Q_i} \leftarrow (g(m_{0_i}), g(m_{1_i}))$
    $RK_{Q_i} \leftarrow b_i$
$QT_Q \leftarrow \{\gamma_i\}$, $VK_Q \leftarrow \{VK_{Q_i}\}$, $RK_Q \leftarrow \{RK_{Q_i}\}$

Alg. 5 $R \leftarrow \mathsf{Search}(\mathcal{I}_D, QT_Q = \{\gamma_i\}, st_s, SK_{\mathrm{S}}, \mathrm{PP})$

$\widetilde{j} \leftarrow \mathsf{BE.Decrypt}(st_s, uk_S)$
if $(\widetilde{j} \neq st_s)$ **then return** $\perp$
for $i = 1$ **to** $|Q|$ **do**
    $(c_{b_i} \| c_{1-b_i}) \leftarrow \phi_{\widetilde{j}}^{-1}(\gamma_i)$
    $d_{b_i} \leftarrow \mathsf{ABE.Decrypt}(c_{b_i}, \mathcal{I}_D, MPK_{\mathrm{ABE}})$
    $d_{1-b_i} \leftarrow \mathsf{ABE.Decrypt}(c_{1-b_i}, \mathcal{I}_D, MPK_{\mathrm{ABE}})$
    $R_i = (d_{b_i}, d_{1-b_i})$
$R = \{R_i\}$

Alg. 6 $RT_Q \leftarrow \mathsf{BVerif}(R = \{(d_i, d'_i)\}, VK_Q = \{(VK_i, VK'_i)\}, \mathrm{PP})$

1: **for** $i = 1$ **to** $|Q|$ **do**
2: **if** $VK_i = g(d_i)$ **then** $RT_{Q_i} = d_i$
3: **else if** $VK'_i = g(d'_i)$ **then** $RT_{Q_i} = d'_i$
4: **else** $RT_{Q_i} = \perp$
5: $RT_Q = \{RT_{Q_i}\}$

Alg. 7 $r \leftarrow \mathsf{Retrieve}(VK_Q = \{(g(m_{b_i}), g(m_{1-b_i}))\}, RT_Q = \{RT_{Q_i}\}, RK_Q = \{b_i\}, \mathrm{PP})$

1: **for** $i = 1$ **to** $|Q|$ **do**
2: **if** $g(RT_{Q_i}) = g(m_0)$ **then** $r_i = 1$
3: **else if** $g(RT_{Q_i}) = g(m_1)$ **then** $r_i = 0$
4: **else** $r_i = \perp$
5: $r = \{r_i\}$

Alg. 8 $(st_s, st_o) \leftarrow \mathsf{RevokeUser}(\mathrm{ID}, G, \mathrm{MK}, \mathrm{PP})$

1: $j' \xleftarrow{\$} \{0,1\}^\kappa$
2: $st_s \leftarrow \mathsf{BE.Encrypt}(G \setminus \mathrm{ID}, \mathrm{j}', \mathrm{mk})$
3: $st_o \leftarrow j'$

Theorem 1. *Given a selective* IND-CPA *secure CP-ABE scheme, an authenticated symmetric encryption scheme and a broadcast encryption scheme, both secure in the sense of* IND-CPA*, pseudo-random permutations Π and ϕ, and a one-way function g. Let $e\mathcal{VSE}$ be the extended verifiable searchable encryption scheme defined in algorithms 1–8. Then $e\mathcal{VSE}$ is secure in the sense of Public Verifiability, Index Privacy and Query Privacy.*

In Appendix A.1 we provide a proof sketch. Full proofs can be found in [3]. In Appendix B we discuss the trade-off between efficiency and functionality of our scheme. Note that we can add additional contextual access control following Alderman et al. [1] by replacing ϕ with a *key assignment scheme*.

5 Conclusion

With this work we have begun to consider the application of VC techniques in the setting of searchable encryption. On the searchable encryption side, this enables additional functionality in the form of computational queries (e.g. computing the average of outsourced data fields that are linked to a specific set of keywords), whilst on the VC side, this introduces additional privacy concerns regarding the outsourced data and computations. The choice of using VC techniques based on ABE stems from the natural correspondence between attributes and keywords in an index. However, future work should investigate other forms of VC to achieve different classes of functionality and (especially) improve efficiency.

In future work, we would like to consider a model whereby multiple data owners can store data on a server without each having to initialise their own scheme. In practice, this could result in the Key Distribution Center from VDC [2] setting up the system and publishing public parameters that any data owner can use, but enabling each data owner to generate their own CP-ABE decryption keys for the data they hold.

References

1. J. Alderman, C. Janson, C. Cid, and J. Crampton. Access control in publicly verifiable outsourced computation. In F. Bao, S. Miller, J. Zhou, and G. Ahn, editors, *Proceedings of the 10th ACM Symposium on Information, Computer and Communications Security, ASIA CCS '15, Singapore, April 14-17, 2015*, pages 657–662. ACM, 2015.
2. J. Alderman, C. Janson, C. Cid, and J. Crampton. Hybrid publicly verifiable computation. *IACR Cryptology ePrint Archive*, 2015:320, 2015.
3. J. Alderman, C. Janson, K. M. Martin, and S. L. Renwick. Extended functionality in verifiable searchable encryption. *IACR Cryptology ePrint Archive*, 2015.
4. D. Apon, J. Katz, E. Shi, and A. Thiruvengadam. Verifiable oblivious storage. In H. Krawczyk, editor, *Public-Key Cryptography - PKC 2014 - 17th International Conference on Practice and Theory in Public-Key Cryptography, Buenos Aires, Argentina, March 26-28, 2014. Proceedings*, volume 8383 of *Lecture Notes in Computer Science*, pages 131–148. Springer, 2014.
5. M. Backes, M. Barbosa, D. Fiore, and R. M. Reischuk. ADSNARK: nearly practical and privacy-preserving proofs on authenticated data. In *2015 IEEE Symposium on Security and Privacy, SP 2015, San Jose, CA, USA, May 17-21, 2015*, pages 271–286. IEEE Computer Society, 2015.
6. M. Backes, D. Fiore, and R. M. Reischuk. Verifiable delegation of computation on outsourced data. In A. Sadeghi, V. D. Gligor, and M. Yung, editors, *2013 ACM SIGSAC Conference on Computer and Communications Security, CCS'13, Berlin, Germany, November 4-8, 2013*, pages 863–874. ACM, 2013.
7. M. Bellare and C. Namprempre. Authenticated encryption: Relations among notions and analysis of the generic composition paradigm. *J. Cryptology*, 21(4):469–491, 2008.
8. E. Ben-Sasson, A. Chiesa, D. Genkin, and E. Tromer. Fast reductions from rams to delegatable succinct constraint satisfaction problems: extended abstract. In R. D. Kleinberg, editor, *Innovations in Theoretical Computer Science, ITCS '13, Berkeley, CA, USA, January 9-12, 2013*, pages 401–414. ACM, 2013.
9. S. Benabbas, R. Gennaro, and Y. Vahlis. Verifiable delegation of computation over large datasets. In Rogaway [27], pages 111–131.
10. N. Bitansky, R. Canetti, A. Chiesa, and E. Tromer. From extractable collision resistance to succinct non-interactive arguments of knowledge, and back again. In S. Goldwasser, editor, *Innovations in Theoretical Computer Science 2012, Cambridge, MA, USA, January 8-10, 2012*, pages 326–349. ACM, 2012.
11. D. Boneh, G. Di Crescenzo, R. Ostrovsky, and G. Persiano. Public key encryption with keyword search. In *Advances in Cryptology - EUROCRYPT 2004*, volume 3027 of *Lecture Notes in Computer Science*, pages 506–522. Springer, 2004.
12. J. W. Byun, H. S. Rhee, H. Park, and D. H. Lee. Off-line keyword guessing attacks on recent keyword search schemes over encrypted data. In *Secure Data Management, Third VLDB Workshop, SDM 2006*, volume 4165 of *Lecture Notes in Computer Science*, pages 75–83. Springer, 2006.
13. Q. Chai and G. Gong. Verifiable symmetric searchable encryption for semi-honest-but-curious cloud servers. In *Proceedings of IEEE International Conference on Communications, ICC 2012*, pages 917–922. IEEE, 2012.

14. K. Chung, Y. T. Kalai, F. Liu, and R. Raz. Memory delegation. In Rogaway [27], pages 151–168.
15. R. Curtmola, J. A. Garay, S. Kamara, and R. Ostrovsky. Searchable symmetric encryption: improved definitions and efficient constructions. In *3th ACM Conference on Computer and Communications Security,*, pages 79–88. ACM, 2006.
16. Z. Fu, J. Shu, X. Sun, and N. Linge. Smart cloud search services: verifiable keyword-based semantic search over encrypted cloud data. *Consumer Electronics, IEEE Transactions on*, 60(4):762–770, 2014.
17. R. Gennaro, C. Gentry, and B. Parno. Non-interactive verifiable computing: Outsourcing computation to untrusted workers. In T. Rabin, editor, *CRYPTO*, volume 6223 of *Lecture Notes in Computer Science*, pages 465–482. Springer, 2010.
18. E. Goh. Secure indexes. *IACR Cryptology ePrint Archive*, 2003:216, 2003.
19. O. Goldreich and R. Ostrovsky. Software protection and simulation on oblivious rams. *Journal of the Association for Computing Machinery*, 43:431–473, 1996.
20. S. Kamara, C. Papamonthou, and T. Roeder. Dynamic searchable symmetric encryption. In *Conference on Computer and Communications Security*, pages 965–976. ACM, 2012.
21. J. Katz, A. Sahai, and B. Waters. Predicate encryption supporting disjunctions, polynomial equations, and inner products. In *Advances in Cryptology - EUROCRYPT 2008*, volume 4965 of *Lecture Notes in Computer Science*, pages 146–162. Springer, 2008.
22. K. Kurosawa and Y. Ohtaki. How to update documents verifiably in searchable symmetric encryption. In *Cryptology and Network Security - 12th International Conference, CANS 2013*, volume 8257 of *Lecture Notes in Computer Science*, pages 309–328. Springer, 2013.
23. J. Li, Q. Wang, C. Wang, N. Cao, K. Ren, and W. Lou. Fuzzy keyword search over encrypted data in cloud computing. In *INFOCOM 2010. 29th IEEE International Conference on Computer Commu- nications, Joint Conference of the IEEE Computer and Communications Societies*, pages 441–445. IEEE, 2010.
24. P. Liu, J. Wang, H. Ma, and H. Nie. Efficient verifiable public key encryption with keyword search based on KP-ABE. In *Ninth International Conference on Broadband and Wireless Computing, Communication and Applications, BWCCA 2014*, pages 584–589. IEEE, 2014.
25. D. J. Park, K. Kim, and P. J. Lee. Public key encryption with conjunctive field keyword search. In *Information Security Applications, 5th International Workshop*, volume 3325 of *Lecture Notes in Computer Science*, pages 73–86. Springer, 2004.
26. B. Parno, M. Raykova, and V. Vaikuntanathan. How to delegate and verify in public: Verifiable computation from attribute-based encryption. In R. Cramer, editor, *TCC*, volume 7194 of *Lecture Notes in Computer Science*, pages 422–439. Springer, 2012.
27. P. Rogaway, editor. *Advances in Cryptology - CRYPTO 2011 - 31st Annual Cryptology Conference, Santa Barbara, CA, USA, August 14-18, 2011. Proceedings*, volume 6841 of *Lecture Notes in Computer Science*. Springer, 2011.
28. D. X. Song, D. Wagner, and A. Perrig. Practical techniques for searches on encrypted data. In *IEEE Symposium on Security and Privacy, Berkeley, California, USA*, pages 44–55. IEEE, 2000.
29. E. Stefanov, C. Papamonthou, and E. Shi. Practical dynamic searchable encryption with small leakage. In *21st Annual Network and Distributed System Security Symposium, NDSS 2014*. The Internet Society, 2014.

30. W. Sun, B. Wang, N. Cao, M. Li, W. Lou, Y. T. Hou, and H. Li. Verifiable privacy-preserving multi-keyword text search in the cloud supporting similarity-based ranking. *IEEE Transactions on Parallel Distributed Systems*, 25(11):3025–3035, 2014.
31. W. Sun, S. Yu, W. Lou, T. Hou, and H. Li. Protecting your right: Verifiable attribute-based keyword search with fine-grainedowner-enforced search authorization in the cloud. *Parallel and Distributed Systems, IEEE Transactions on*, (99), 2013.
32. C. Wang, N. Cao, J. Li, and W. Lou. Secure ranked keyword search over encrypted cloud data. In *International Conference on Distributed Computing Systems, ICDCS 2010*, pages 253–262. IEEE Computer Society, 2010.
33. C. Wang, N. Cao, K. Ren, and W. Lou. Enabling secure and efficient ranked keyword search over outsourced cloud data. *IEEE Transactions Parallel Distributed Systems*, 23(8):1467–1479, 2012.
34. J. Wang, H. Ma, J. Li, H. Zhu, S. Ma, and X. Chen. Efficient verifiable fuzzy keyword search over encrypted data in cloud computing. *Computer Science Information Systems*, 10(2):667–684, 2013.
35. J. Wang, H. Ma, Q. Tang, J. Li, H. Zhu, S. Ma, and X. Chen. Efficient verifiable fuzzy keyword search over encrypted data in cloud computing. *Computer Science Information Systems*, 10(2):667–684, 2013.
36. B. Waters. Ciphertext-policy attribute-based encryption: An expressive, efficient, and provably secure realization. In D. Catalano, N. Fazio, R. Gennaro, and A. Nicolosi, editors, *Public Key Cryptography*, volume 6571 of *Lecture Notes in Computer Science*, pages 53–70. Springer, 2011.
37. Q. Zheng, S. Xu, and G. Ateniese. VABKS: verifiable attribute-based keyword search over outsourced encrypted data. In *2014 IEEE Conference on Computer Communications, INFOCOM 2014*, pages 522–530. IEEE, 2014.

A Security Models

Index Privacy. In Game 2, we formalise the notion of index indistinguishability against a selective chosen keyword attack, which ensures no information regarding the keywords is leaked from the index. Firstly the adversary outputs two sets of attributes ($D_0, D_1 \subseteq \mathcal{U}$) that they wish to be challenged on, with the restriction that $|D_0| = |D_1|$ (this is required as the CP-ABE used to produce the index does not conceal the index length). The challenger runs Setup to produce the public and secret parameters. The challenger selects a bit $b \in \{0, 1\}$ uniformly at random to select which set of attributes to encode into the index. Before the index is created, the challenger needs to create the pre-index from the set of attributes D_b (line 4 of Game 2). This is done using an Encode mechanism that takes the elements of D_b as input and outputs the pre-index $\delta(D_b)$. Encode is not required in our instantiation as the pre-indexes can be chosen directly from $\tilde{U}$ as the user knows the mapping from $\mathcal{U}$ to $\mathcal{U}'$ and the permutation Π; the adversary however does not. The challenger then runs $\mathsf{BuildIndex}$ using $\delta(D_b)$ to produce the index $\mathcal{I}_{D_b}$, which is given to $\mathcal{A}$. The adversary is then given PP and oracle access, with the restriction that the query results are identical for each index $\mathcal{I}_{D_0}, \mathcal{I}_{D_1}$, i.e. if $R_0 \leftarrow \mathsf{Search}(\mathcal{I}_{D_0}, QT_Q, st_s, SK_{\mathrm{S}}, \mathrm{PP})$ and $R_1 \leftarrow \mathsf{Search}(\mathcal{I}_{D_1}, QT_Q, st_s, SK_{\mathrm{S}}, \mathrm{PP})$ then we need $R_0 = R_1$. After this query

Game 2 $\mathbf{Exp}_{\mathcal{A}}^{IndPriv}\,[e\mathcal{VSE}, 1^\kappa]$:

$(D_0, D_1, Q) \leftarrow \mathcal{A}(1^\kappa, \mathcal{U})$
if $(|D_0| \neq |D_1|)$ **then return** $\perp$
$b \xleftarrow{\$} \{0,1\}$
$(\mathrm{MK}, \mathrm{PP}) \leftarrow \mathsf{Setup}(1^\kappa, \mathcal{U})$
$G \leftarrow \emptyset$
$\mathrm{ID} \xleftarrow{\$} \mathsf{Users}$
$(SK_{\mathrm{ID}}, st_s) \leftarrow \mathsf{AddUser}(\mathrm{ID}, G, \mathrm{MK}, \mathrm{PP})$
$\delta(D_b) \leftarrow \mathsf{Encode}(D_b)$
$(\mathcal{I}_{D_b}, st_s, st_o) \leftarrow \mathsf{BuildIndex}(\delta(D_b), G, \mathrm{MK}, \mathrm{PP})$
$b' \leftarrow \mathcal{A}^{\mathcal{O}}(\mathcal{I}_{D_b}, st_s, \mathrm{PP})$
return $(b' == b)$

Game 3 $\mathbf{Exp}_{\mathcal{A}}^{QueryPriv}\,[e\mathcal{VSE}, 1^\kappa]$:

$(Q_0, Q_1) \leftarrow \mathcal{A}(1^\kappa, \mathcal{U})$
if $(\mathcal{G}_{Q_0} \neq \mathcal{G}_{Q_1})$ **then return** $\perp$
$b \xleftarrow{\$} \{0,1\}$
$(\mathrm{MK}, \mathrm{PP}) \leftarrow \mathsf{Setup}(1^\kappa, \mathcal{U})$
$G \leftarrow \emptyset$
$\mathrm{ID} \xleftarrow{\$} \mathsf{Users}$
$(SK_{\mathrm{ID}}, st_s) \leftarrow \mathsf{AddUser}(\mathrm{ID}, G, \mathrm{MK}, \mathrm{PP})$
$\delta(D_b) \xleftarrow{\$} \tilde{\mathcal{U}}$
$(\mathcal{I}_D, st_s, st_o) \leftarrow \mathsf{BuildIndex}(\delta(D), G, \mathrm{MK}, \mathrm{PP})$
$\tilde{Q}_b \leftarrow \mathsf{Encode}(Q_b)$
$(QT_{Q_b}, VK_{Q_b}, RK_{Q_b}) \leftarrow \mathsf{Query}(\tilde{Q}_b, st_s, st_o, SK_{\mathrm{ID}}, \mathrm{PP})$
$b' \leftarrow \mathcal{A}^{\mathcal{O}}(QT_{Q_b}, VK_{Q_b}, RK_{Q_b}, \mathcal{I}_D, st_s, \mathrm{PP})$
return $(b' == b)$

phase, $\mathcal{A}$ outputs a guess b' and wins the game if the comparison operator $==$ returns 1 which indicates that $b' = b$. Hence $\mathcal{A}$ wins the game if they can identify which attribute set (D_0 or D_1) was encoded into the index $\mathcal{I}_{D_b}$.

Query Privacy. The queries themselves should not leak any information about the corresponding keywords that make up the query. Our construction of the queries leaks the gates, but not the keywords themselves. This notion of query indistinguishability against a selective chosen query attack is formalised in Game 3. The game runs similarly to that of Game 2, subject to the following restrictions: the challenge queries (Q_0, Q_1) must use the same gates. We denote the gate structure of a query Q by $\mathcal{G}_Q$, and hence require that $\mathcal{G}_{Q_0} = \mathcal{G}_{Q_1}$.

A.1 Security Proofs

Proof (Public Verifiability). Here we provide a proof sketch; full details can be found in [3]. We start by assuming that $\mathcal{A}_{eVSE}$ is an adversary with non-negligible advantage δ. We begin by defining the following three games:

- **Game A.** This is the selective Public Verifiability game as defined in Game 1.
- **Game B.** This is the same as **Game A** with the modification that in Query, we no longer return an encryption of m_0 and m_1.

Instead, we choose another random message $m' \neq m_0, m_1$ and, if $Q(\mathcal{I}_D) = 1$, we replace c_1 by $\mathsf{ABE.Encrypt}(\overline{Q}, m', MPK_{\text{ABE}})$. Otherwise, we replace c_0 by $\mathsf{ABE.Encrypt}(Q, m', MPK_{\text{ABE}})$.

– **Game C.** This is the same as **Game B** with the exception that instead of choosing a random message m', we implicitly set m' to be the challenge input w in the one-way function game.

We show that an adversary with non-negligible advantage against the selective Public Verifiability game can be used to construct an adversary that may invert the one-way function g.

We begin by showing that there is a negligible distinguishing advantage between **Game A** and **Game B**. We construct an adversary $\mathcal{A}_{ABE}$ that creates an eVSE instance by executing algorithms 1–8 and uses $\mathcal{A}_{eVSE}$ as a sub-routine to break the selective IND-CPA security of the CP-ABE scheme. The advantage of our constructed adversary is $Adv_{\mathcal{A}_{ABE}} \geqslant \frac{\delta}{2}$. Hence, if $\mathcal{A}_{eVSE}$ has advantage δ at distinguishing these games then $\mathcal{A}_{ABE}$ can win the sIND-CPA game for CP-ABE with non-negligible probability. Thus since we assumed the CP-ABE scheme to be secure, we conclude that $\mathcal{A}_{eVSE}$ cannot distinguish the games with non-negligible probability. The transition from **Game B** to **Game C** is simply to set the value of m'_i to no longer be random but instead to correspond to the challenge w in the one-way function inversion game. We argue that the adversary has no distinguishing advantage between these games since the new value is independent of anything else in the system bar the verification key $g(w)$ and hence looks random to an adversary with no additional information. Finally we show that using $\mathcal{A}_{eVSE}$ in **Game C**, $\mathcal{A}_{ABE}$ can invert the one-way function g – that is, given a challenge $z = g(w)$ we can recover w. Now, if $\mathcal{A}_{eVSE}$ is successful, it will output a forgery comprising the plaintext encrypted under the unsatisfied query (Q or $\overline{Q}$). By construction, this will be w and $\mathcal{A}_{ABE}$ can therefore forward this result to $\mathcal{C}$ in order to invert the one-way function with the same non-negligible probability that $\mathcal{A}_{eVSE}$ has against the public verifiability game.

We conclude that if the ABE scheme is sIND-CPA secure and the one-way function is hard-to-invert, then $e\mathcal{VSE}$ as defined by Algorithms 1–8 is secure in the sense of selective Public Verifiability. □

The remaining proofs can be found in the full version [3].

B Discussion

Our scheme extends the expressiveness of queries that can be achieved in VSE. No other VSE schemes to our knowledge are able to perform the range of search queries or include negation of keywords in their search queries. Additionally our scheme leaks neither the access or the search pattern to the server whilst executing a search. Our combination of search queries with computational queries is also a novel functionality in the field of VSE.

The search time and size of the queries are both linear in n (the amount of data items stored on the remote server). Due to this eVSE may be more suited to smaller databases to prevent these features from being prohibitively expensive. The VSE scheme of [13] has a search time that is linear in the number

Table 1: Comparison of Schemes

Scheme	Data type	Query type	Publicly Verifiable	Leakage	Computations
[33]	Static	Ranked equality	No	AP,SP	No
[22]	Dynamic	Equality	Yes	AP	No
[30]	Static	Conjunctive, Disjunctive	No	AP	No
[31]	Dynamic	Conjunctive	No	AP	No
[29]	Dynamic	Equality	No	AP, SP	No
[37]	Static	Equality	No	AP	No
[34]	Static	Fuzzy	No	AP, SP	No
[16]	Static	Semantic	No	AP, SP	No
[13]	Static	Equality	No	AP, SP	No
Our scheme	Static	Conjunctive, Disjunctive, Arbitrary CNF/DNF formulae, NC^1	Yes	None	Yes

of letters in the queried keyword (which is usually much smaller than n). This faster search is achieved using a tree-based index, however only a single keyword equality search can be performed. Another scheme built using ABE [37] is able to achieve multi-level access, where users can be restricted to searching only certain parts of the database. Keywords are grouped with respect to their access control policies, and the search time is linear in the number of groups. This scheme also only achieves a single keyword equality search. The scheme of [35] achieves verifiable fuzzy keyword search with a search time that is linear in the size of the fuzzy keyword set (which varies depending on the level of fuzziness required i.e. searching for data items that contain keywords of edit distance two will require a larger fuzzy keyword set than searching for keywords with an edit distance of one from the queried keyword [23]). Again, this is likely to be less than n. In terms of the number of rounds of communication required per search, our scheme is optimal requiring only one round of communication. The size of the search results in our scheme is also linear in n. Most VSE schemes in the literature return results of a size that is linear in the number of data items that match the query, however this method leaks the access pattern which in turn may leak information about the query. Our scheme hides the access pattern as all search results are of the same form, regardless of what query was submitted.

In terms of security, as illustrated in our security games, our scheme achieves public verifiability, index privacy and query privacy (in terms of the keywords searched for), which is comparable to other VSE schemes that have been discussed. Overall, our scheme sacrifices efficiency when compared to existing VSE schemes, but gains much increased functionality and query expressiveness.

Table 1 gives a comparison between our scheme and those in the literature.