

Conjunctive, Subset, and Range Queries on Encrypted Data

Dan Boneh ^{*}
dabo@cs.stanford.edu

Brent Waters [†]
bwaters@csl.sri.com

Abstract

We construct public-key systems that support comparison queries ($x \geq a$) on encrypted data as well as more general queries such as subset queries ($x \in S$). These systems also support arbitrary conjunctive queries ($P_1 \wedge \dots \wedge P_\ell$) without leaking information on individual conjuncts. In addition, we present a general framework for constructing and analyzing public-key systems supporting queries on encrypted data.

1 Introduction

Queries on encrypted data are easiest to explain with an example. Consider a credit card payment gateway that observes a stream of encrypted transactions, say encrypted under Visa's public key. The gateway needs to flag all transactions satisfying a certain predicate P . Say, all transactions whose value is over \$1000. Storing Visa's secret key on the gateway is a bad idea for both security and privacy concerns. Instead, Visa wishes to give the gateway a token TK_P that enables the gateway to identify transactions satisfying P without learning anything else about these transactions. Of course, generating the token TK_P will require Visa's secret key.

As another example, consider a mail server that receives a stream of email messages encrypted under the recipients public key. If the email message satisfies a certain predicate P the mail server should forward the email to the recipient's pager. If the email satisfies some other predicate P' the server should just discard the email. Otherwise, the server should place the email in the recipient's inbox. The recipient does not want to give the mail server the full private key. Instead, she wants to give the server two tokens TK_P and $TK_{P'}$ enabling the server to test for the predicates P and P' without learning any other information about the email.

Our goal is to build a public-key system that supports a rich set of query predicates. In our payment gateway example one can imagine comparison queries such as ($\text{value} > 1000$) or even conjunctions such as ($\text{value} > 1000$) and ($\text{TransactionTime} > 5\text{pm}$). The gateway should learn no information other than the value of the conjunctive predicate. In case a conjunction $P_1 \wedge P_2$ is false, the gateway should not learn which of the two conjuncts P_1 or P_2 is false. In our second example involving a mail server one can imagine testing for subset queries such as ($\text{sender} \in S$) where S is a set of email addresses. Conjunctive queries such as ($\text{sender} \in S$) and ($\text{subject} = \text{urgent}$) also make sense. Perhaps in the distant future, when highly complex queries on encrypted data are possible, one can imagine running an anti-virus/anti-spam predicate on encrypted emails. The mail server learns nothing about incoming encrypted email other than its spam status.

^{*}Supported by NSF and the Packard Foundation.

[†]Supported by NSF and U.S. Army Research Office under Research Grant No. W911NF-06-1-0316.

Unfortunately, until now, only simple equality queries on encrypted data were possible. Song et al. [19] developed a mechanism for equality tests on data encrypted with a symmetric key system. Boneh et al. [9] constructed equality tests in the public-key settings.

Our results. We present a general framework for analyzing and constructing searchable public-key systems for various families of predicates. We then construct public-key systems that support comparison queries (such as greater-than) and general subset queries. We also support arbitrary conjunctions. We evaluate our results based on ciphertext size and token size. Let $T = \{1, 2, \dots, n\}$ and suppose we encrypt a tuple $x = (x_1, \dots, x_w) \in T^w$. Say x_1 is a transaction value, x_2 is a card expiration date, and so on. The following table summarizes our results at a high level.

Query Type	Source	Ciphertext Size	Token Size
Equality query: $(x_i = a)$ for any $a \in T$	[19, 17, 9, 1]	$O(1)$	$O(1)$
Comparison query: $(x_i \geq a)$ for any $a \in T$	[11, 12] ¹	$O(\sqrt{n})$	$O(\sqrt{n})$
Subset query: $(x_i \in A)$ for any $A \subseteq T$	This paper	$O(n)$	$O(n)$
Equality conjunction: $(x_1 = a_1) \wedge \dots \wedge (x_w = a_w)$	This paper	$O(w)$	$O(w)$
Comparison conjunction: $(x_1 \geq a_1) \wedge \dots \wedge (x_w \geq a_w)$	This paper	$O(nw)$	$O(w)$
Subset conjunction: $(x_1 \in A_1) \wedge \dots \wedge (x_w \in A_w)$	This paper	$O(nw)$	$O(nw)$

Here (a_1, \dots, a_w) is an arbitrary vector that defines a conjunctive equality or a comparison predicate. Similarly, A_1, \dots, A_w are *arbitrary* subsets of $\{1, \dots, n\}$ that define a conjunctive subset query predicate. We emphasize that when a conjunction predicate is false, the system does not leak which of the w conjuncts caused it.

Prior to these results the best systems for comparison and subset queries were the trivial brute-force systems that we discuss in Section 3. For comparison queries these systems generate a ciphertext of size $O(n^w)$ and for subset queries they generate a ciphertext of size $O(2^{nw})$. Note that even without conjunction, namely for $w = 1$, our subset query construction generates ciphertexts that are exponentially shorter than the best known previous solution ($O(n)$ vs. $O(2^n)$).

The main tool used in these constructions is a new primitive we call *Hidden Vector Encryption* or HVE for short. This primitive can be viewed as an extreme generalization of Anonymous Identity Based Encryption (AnonIBE) [9, 1, 13]. We show how HVE implies all the results in the table.

Are there public key systems that support larger classes of predicates? Ultimately, one would like a public-key system that supports searches for any predicate computable by a shallow circuit. Presently, this appears to be a difficult open problem.

Related work. Equality tests on encrypted data were considered in [19, 9]. Equality searches on an encrypted audit log were proposed in [20]. Equality tests in the symmetric key settings are closely related to oblivious RAM techniques [17, 14]. Equality tests in the public key settings are closely related to Anonymous Identity Based Encryption (AnonIBE) [9, 1, 13]. Conjunctive equality queries were first studied in [15]. Equality searches on streaming data that hide the requested predicate were discussed in [18] and [4]. Efficient equality searches in databases were

¹Both papers [11, 12] focus on traitor tracing, but as we observe in Appendix C, their approach directly gives a comparison searching system without conjunctions.

recently presented in [2]. Bethencourt et al. [3] recently gave a construction for efficient range queries in a weaker security model. That is, when the encrypted index falls in the specified range, the search token reveals the index. In addition, their decryption algorithm takes time $O(\lg(n)^w)$ for a conjunctive search over w ranges, each of size n .

2 Definitions

We begin by defining a general framework for queries on encrypted data. Let Σ be a finite set of binary strings. A predicate P over Σ is a function $P : \Sigma \rightarrow \{0, 1\}$. We say that $I \in \Sigma$ satisfies the predicate if $P(I) = 1$.

2.1 Searchable encryption

Let Φ be a set of predicates over Σ . A Φ -**searchable** public key system comprises of the following algorithms:

Setup(λ) A probabilistic algorithm that takes as input a security parameter and outputs a public key PK and secret key SK.

Encrypt(PK, I, M) Encrypts the plaintext pair (I, M) using the public key PK. We view $I \in \Sigma$ as the searchable field, called an **index**, and $M \in \mathcal{M}$ as the data.

GenToken(SK, $\langle P \rangle$) Takes as input a secret key SK and the description of a predicate $P \in \Phi$. It outputs a token TK_P .

Query(TK, C) Takes a token TK for some predicate $P \in \Phi$ as input and a ciphertext C . It outputs a message $M \in \mathcal{M}$ or \perp . Roughly speaking, if C is an encryption of (I, M) then the algorithm outputs M when $P(I) = 1$ and outputs \perp otherwise. The precise requirement is captured in the query correctness property below.

Correctness. The system must satisfy the following **correctness property**:

- **Query correctness:** For all $(I, M) \in \Sigma \times \mathcal{M}$ and all predicates $P \in \Phi$:

Let $(\text{PK}, \text{SK}) \stackrel{R}{\leftarrow} \text{Setup}(\lambda)$, $C \stackrel{R}{\leftarrow} \text{Encrypt}(\text{PK}, I, M)$, and $\text{TK} \stackrel{R}{\leftarrow} \text{GenToken}(\text{SK}, \langle P \rangle)$.

If $P(I) = 1$ then $\text{Query}(\text{TK}, C) = M$.

If $P(I) = 0$ then $\Pr[\text{Query}(\text{TK}, C) = \perp] > 1 - \epsilon(\lambda)$ where $\epsilon(\lambda)$ is a negligible function.

Suppose that given a ciphertext $C \leftarrow \text{Encrypt}(\text{PK}, I, M)$ we are only interested in testing whether a predicate $P(I)$ is satisfied. In this case the message space \mathcal{M} can be set to a singleton, say $\mathcal{M} = \{\text{true}\}$. Algorithm $\text{Query}(\text{TK}, C)$ will return **true** when $P(I) = 1$ and \perp otherwise. A larger message space \mathcal{M} is useful if TK is intended to unlock some $M \in \mathcal{M}$ whenever the predicate $P(I) = 1$. For example, when the transaction value is over \$1000 we may want the payment gateway to obtain more information about the transaction. Otherwise, the gateway should learn nothing.

Notice that a Φ -searchable system does not provide a *Decrypt* algorithm that uses SK to decrypt a ciphertext C and outputs (I, M) . One can always add this capability by also encrypting (I, M) under a standard public key system. There is no need for the searchable system to explicitly provide this capability.




Figure 1: Tokens for $\sigma_1, \sigma_2, \sigma_3, \sigma_4$ given to the adversary

An example – comparison queries. Before defining security, we first give a motivating example using comparison queries. Let $\Sigma = \{1, \dots, n\}$ for some integer n . For $\sigma \in \{1, \dots, n\}$ let P_σ be the following comparison predicate:

$$P_\sigma(x) = \begin{cases} 1 & \text{if } x \geq \sigma, \\ 0 & \text{otherwise} \end{cases}$$

Let $\Phi_n = \{P_1, \dots, P_n\}$ be the set of all n comparison predicates. Suppose the adversary has the tokens for predicates $P_{\sigma_1}, P_{\sigma_2}, \dots, P_{\sigma_w}$ where $\sigma_1 < \sigma_2 < \dots < \sigma_w$. Lets x, y, z be some integers as in Figure 1. Clearly the adversary can distinguish $Encrypt(\text{PK}, x, m)$ from $Encrypt(\text{PK}, y, m)$ using the token for the predicate P_{σ_2} . However, the adversary should not be able to distinguish $Encrypt(\text{PK}, y, m)$ from $Encrypt(\text{PK}, z, m)$. Indeed, separating an encryption of y from an encryption of z is information that should not be exposed by the tokens at the adversary’s disposal. Our definition of security captures this property using the general framework.

2.2 Security

We define security of a Φ -searchable system \mathcal{E} using a **query security game** that captures the intuition that tokens TK reveal no unintended information about the plaintext. The game gives the adversary a number of tokens and requires that the adversary cannot use these tokens to deduce unintended information. The game proceeds as follows:

- **Setup.** The challenger runs $Setup(\lambda)$ and gives the adversary PK.
- **Query phase 1.** The adversary adaptively outputs descriptions of predicates $P_1, P_2, \dots, P_{q_1} \in \Phi$. The challenger responds with the corresponding tokens $\text{TK}_j \leftarrow GenToken(\text{SK}, \langle P_j \rangle)$. We refer to such queries as **predicate queries**.
- **Challenge.** The adversary outputs two pairs (I_0, M_0) and (I_1, M_1) subject to two restrictions:
 - First, $P_j(I_0) = P_j(I_1)$ for all $j = 1, 2, \dots, q_1$.
 - Second, if $M_0 \neq M_1$ then $P_j(I_0) = P_j(I_1) = 0$ for all $j = 1, 2, \dots, q_1$.

The challenger flips a coin $\beta \in \{0, 1\}$ and gives $C_* \stackrel{R}{\leftarrow} Encrypt(\text{PK}, I_\beta, M_\beta)$ to the adversary.

The two restrictions ensure that the tokens given to the adversary do not trivially break the challenge. The first restriction ensures that tokens given to the adversary do not directly distinguish I_0 from I_1 . The second restriction ensures that the tokens do not directly distinguish M_0 from M_1 .

- **Query phase 2.** The adversary continues to adaptively request tokens for predicates $P_{q_1+1}, \dots, P_q \in \Phi$, subject to the two restrictions above. The challenger responds with the corresponding tokens $\text{TK}_j \leftarrow GenToken(\text{SK}, \langle P_j \rangle)$.

- **Guess** The adversary returns a guess $\beta' \in \{0, 1\}$ of β .

We define the advantage of adversary \mathcal{A} in attacking \mathcal{E} as the quantity $\text{QU Adv}_{\mathcal{A}} = |\Pr[\beta' = \beta] - 1/2|$.

Definition 2.1. We say that a Φ -searchable system \mathcal{E} is **secure** if for all polynomial time adversaries \mathcal{A} attacking \mathcal{E} the function $\text{QU Adv}_{\mathcal{A}}$ is a negligible function of λ .

Another example – equality queries. Let Σ be some finite set. For $\sigma \in \Sigma$ let $P_{\sigma}(x)$ be an equality predicate, namely

$$P_{\sigma}(x) = \begin{cases} 1 & \text{if } x = \sigma, \\ 0 & \text{otherwise} \end{cases}$$

Let $\Phi_{\text{eq}} = \{P_{\sigma} \text{ for all } \sigma \in \Sigma\}$. Then a Φ_{eq} -searchable encryption supports equality queries on ciphertexts. It is easy to see that a secure Φ_{eq} -searchable encryption is also an anonymous IBE system [9, 1, 13] — an Identity Based Encryption system where a ciphertext reveals no useful information about the identity that was used to create it. This should not be too surprising since it was previously shown [9, 1] that anonymous IBE is sufficient for equality searches. A Φ_{eq} -searchable encryption system ($Setup, Encrypt, GenToken, Query$) gives an anonymous IBE as follows:

- $Setup_{\text{IBE}}(\lambda)$ runs $Setup(\lambda)$ and outputs IBE parameters PK and master key SK.
- $Encrypt_{\text{IBE}}(\text{PK}, \mathcal{I}, M)$ where $\mathcal{I} \in \Sigma$ outputs $Encrypt(\text{PK}, \mathcal{I}, M)$.
- $Extract_{\text{IBE}}(\text{SK}, \mathcal{I})$ where $\mathcal{I} \in \Sigma$ outputs $\text{TK}_{\mathcal{I}} \leftarrow GenToken(\text{SK}, \langle P_{\mathcal{I}} \rangle)$.
- $Decrypt_{\text{IBE}}(\text{TK}_{\mathcal{I}}, C)$ outputs $Query(\text{TK}_{\mathcal{I}}, C)$.

The correctness property ensures that if C is the result of $Encrypt(\text{PK}, \mathcal{I}, M)$ then $Query(\text{TK}_{\mathcal{I}}, C)$ will output M since $P_{\mathcal{I}}(\mathcal{I}) = 1$. It is not difficult to see that the Φ_{eq} -security game ensures semantic security for both the message and the identity. Hence, the resulting system is an anonymous IBE.

By considering larger classes of predicates Φ we obtain more general searching capabilities. The challenge is then to build secure encryption schemes that are Φ -searchable for the most general Φ possible.

Chosen ciphertext security. Definition 2.1 easily extends to address chosen ciphertext attacks (CCA), but we do not pursue that here.

2.3 Selective security

We will also need a slightly weaker security definition in which the adversary commits to the search strings I_0, I_1 at the beginning of the game. Everything else remains the same. The game proceeds as follows:

- **Setup.** The adversary outputs two strings $I_0, I_1 \in \Sigma$. The challenger runs $Setup(\lambda)$ and gives the adversary PK.
- **Query phase 1.** The adversary adaptively outputs descriptions of predicates $P_1, P_2, \dots, P_{q_1} \in \Phi$. The only restriction is that

$$P_j(I_0) = P_j(I_1) \text{ for all } j = 1, 2, \dots, q_1 \quad (1)$$

The challenger responds with the corresponding tokens $\text{TK}_j \leftarrow GenToken(\text{SK}, \langle P_j \rangle)$.

- **Challenge.** The adversary outputs two messages $M_0, M_1 \in \mathcal{M}$ subject to the restriction that:

$$\text{if } M_0 \neq M_1 \text{ then } P_j(I_0) = P_j(I_1) = 0 \text{ for all } j = 1, 2, \dots, q_1 \quad (2)$$

The challenger flips a coin $\beta \in \{0, 1\}$ and gives $C_* \stackrel{R}{\leftarrow} \text{Encrypt}(\text{PK}, I_\beta, M_\beta)$ to the adversary.

- **Query phase 2.** The adversary continues to adaptively request query tokens for predicates $P_{q_1+1}, \dots, P_q \in \Phi$, subject to the two restrictions (1) and (2). The challenger responds with the corresponding tokens $\text{TK}_j \leftarrow \text{GenToken}(\text{SK}, \langle P_j \rangle)$.
- **Guess** The adversary returns a guess $\beta' \in \{0, 1\}$ of β .

The advantage of adversary \mathcal{A} in attacking \mathcal{E} is the quantity $\text{sQU Adv}_{\mathcal{A}} = |\Pr[\beta' = \beta] - 1/2|$.

Definition 2.2. We say that a Φ -searchable system \mathcal{E} is **selectively secure** if for all polynomial time adversaries \mathcal{A} attacking \mathcal{E} the function $\text{sQU Adv}_{\mathcal{A}}$ is a negligible functions of λ .

3 The Trivial Construction

Let Σ be a finite set of binary strings. We build a Φ -searchable public key system \mathcal{E}_{TR} , for *any* set of (polynomial time computable) predicates Φ . We refer to this system as the brute force Φ -searchable system.

The brute force system. Let $\mathcal{E} = (\text{Setup}', \text{Encrypt}', \text{Decrypt}')$ be a public-key system. Let $\Phi = \{P_1, P_2, \dots, P_t\}$ The Φ -searchable system \mathcal{E}_{TR} is defined as follows:

Setup(λ) Run $\text{Setup}'(\lambda)$ t times to obtain

$$\text{PK} \leftarrow (\text{PK}_1, \dots, \text{PK}_t) \quad \text{and} \quad \text{SK} \leftarrow (\text{SK}_1, \dots, \text{SK}_t)$$

Output PK and SK.

Encrypt(PK, I, M) For $j = 1, \dots, t$ define:

$$C_j \stackrel{R}{\leftarrow} \begin{cases} \text{Encrypt}'(\text{PK}_j, M) & \text{if } P_j(I) = 1, \\ \text{Encrypt}'(\text{PK}_j, \perp) & \text{otherwise.} \end{cases}$$

Output $C \leftarrow (C_1, \dots, C_t)$. Note that the length of C is linear in n .

GenToken(SK, $\langle P \rangle$) Here $\langle P \rangle$ (the description of a predicate P) is the index j of P in Φ .
Output $\text{TK} \leftarrow (j, \text{SK}_j)$.

Query(TK, C) Let $C = (C_1, \dots, C_t)$ and $\text{TK} = (j, \text{SK}_j)$. Output $\text{Decrypt}'(\text{SK}_j, C_j)$.

The following lemma proves security of this construction. The proof is a straightforward hybrid argument and is given in Appendix A.

Lemma 3.1. *The system \mathcal{E}_{TR} above is a secure Φ -searchable encryption system assuming \mathcal{E} is a semantically secure public key system against chosen plaintext attacks.*

3.1 A third example — conjunctive comparison predicates

Suppose $\Sigma = \{1, \dots, n\}^w$ for some n, w . Let $\Phi_{n,w}$ be the set of n^w predicates

$$P_{a_1 \dots a_w}(x_1, \dots, x_w) = \begin{cases} 1 & \text{if } x_j \geq a_j \text{ for all } j = 1, \dots, w, \\ 0 & \text{otherwise} \end{cases}$$

for all $\bar{a} = (a_1 \dots a_w) \in \{1, \dots, n\}^w$. Then $|\Phi_{n,w}| = n^w$.

The trivial system in this case produces ciphertexts of length $O(n^w)$. Essentially, the system uses a unary encoding of the w columns and assigns a private key to each cell in this n by w matrix. We will construct a much better system in Section 6.

4 Background on pairings and complexity assumptions

Our goal is to construct Φ -searchable systems for a large class of predicates Φ that is much better than the trivial construction. To do so we will make use of bilinear maps.

4.1 Bilinear groups of composite order

We review some general notions about bilinear maps and groups, with an emphasis on groups of *composite order*. We follow [10] in which composite order bilinear groups were first introduced.

Let \mathcal{G} be an algorithm called a *group generator* that takes as input a security parameter $\lambda \in \mathbb{Z}^{>0}$ and outputs a tuple $(p, q, \mathbb{G}, \mathbb{G}_T, e)$ where p, q are two distinct primes, \mathbb{G} and \mathbb{G}_T are two cyclic groups of order $n = pq$, and e is a function $e : \mathbb{G}^2 \rightarrow \mathbb{G}_T$ satisfying the following properties:

- (Bilinear) $\forall u, v \in \mathbb{G}, \forall a, b \in \mathbb{Z}, e(u^a, v^b) = e(u, v)^{ab}$.
- (Non-degenerate) $\exists g \in \mathbb{G}$ such that $e(g, g)$ has order n in \mathbb{G}_T .

We assume that the group action in \mathbb{G} and \mathbb{G}_T as well as the bilinear map e are all computable in polynomial time in λ . Furthermore, we assume that the description of \mathbb{G} and \mathbb{G}_T includes generators of \mathbb{G} and \mathbb{G}_T respectively.

To summarize, \mathcal{G} outputs the description of a group \mathbb{G} of order $n = pq$ with an efficiently computable bilinear map. We will use the notation $\mathbb{G}_p, \mathbb{G}_q$ to denote the respective subgroups of order p and order q of \mathbb{G} and we will use the notation $\mathbb{G}_{T,p}, \mathbb{G}_{T,q}$ to denote the respective subgroups of order p and order q of \mathbb{G}_T .

4.2 The bilinear Diffie-Hellman assumption

First we review the standard Bilinear Diffie-Hellman assumption, but in groups of composite order. For a given group generator \mathcal{G} define the following distribution $P(\lambda)$:

$$\begin{aligned} & (p, q, \mathbb{G}, \mathbb{G}_T, e) \xleftarrow{R} \mathcal{G}(\lambda), \quad n \leftarrow pq, \quad g_p \xleftarrow{R} \mathbb{G}_p, \quad g_q \xleftarrow{R} \mathbb{G}_q \\ & a, b, c \xleftarrow{R} \mathbb{Z}_n \\ & \bar{Z} \leftarrow ((n, \mathbb{G}, \mathbb{G}_T, e), g_q, g_p, g_p^a, g_p^b, g_p^c) \\ & T \leftarrow e(g_p, g_p)^{abc} \\ & \text{Output } (\bar{Z}, T) \end{aligned}$$

For an algorithm \mathcal{A} , define \mathcal{A} 's advantage in solving the composite bilinear Diffie-Hellman problem for \mathcal{G} as:

$$\text{cBDH Adv}_{\mathcal{G},\mathcal{A}}(\lambda) := \left| \Pr[\mathcal{A}(\bar{Z}, T) = 1] - \Pr[\mathcal{A}(\bar{Z}, R) = 1] \right|$$

where $(\bar{Z}, T) \xleftarrow{R} P(\lambda)$ and $R \xleftarrow{R} \mathbb{G}_{T,p}$.

Definition 4.1. We say that \mathcal{G} satisfies the composite bilinear Diffie-Hellman assumption (cBDH) if for any polynomial time algorithm \mathcal{A} we have that $\text{cBDH Adv}_{\mathcal{G},\mathcal{A}}(\lambda)$ is a negligible function of λ .

4.3 The composite 3-party Diffie-Hellman assumption

Our construction makes use of an additional assumption in composite bilinear groups. For a given group generator \mathcal{G} define the following distribution $P(\lambda)$:

$$\begin{aligned} & (p, q, \mathbb{G}, \mathbb{G}_T, e) \xleftarrow{R} \mathcal{G}(\lambda), \quad n \leftarrow pq, \quad g_p \xleftarrow{R} \mathbb{G}_p, \quad g_q \xleftarrow{R} \mathbb{G}_q \\ & R_1, R_2, R_3 \xleftarrow{R} \mathbb{G}_q \\ & a, b, c \xleftarrow{R} \mathbb{Z}_n \\ & \bar{Z} \leftarrow ((n, \mathbb{G}, \mathbb{G}_T, e), g_q, g_p, g_p^a, g_p^b, g_p^{ab} \cdot R_1, g_p^{abc} \cdot R_2) \\ & T \leftarrow g_p^c \cdot R_3 \\ & \text{Output } (\bar{Z}, T) \end{aligned}$$

For an algorithm \mathcal{A} , define \mathcal{A} 's advantage in solving the composite 3-party Diffie-Hellman problem for \mathcal{G} as:

$$\text{C3DH Adv}_{\mathcal{G},\mathcal{A}}(\lambda) := \left| \Pr[\mathcal{A}(\bar{Z}, T) = 1] - \Pr[\mathcal{A}(\bar{Z}, R) = 1] \right|$$

where $(\bar{Z}, T) \xleftarrow{R} P(\lambda)$ and $R \xleftarrow{R} \mathbb{G}$.

Definition 4.2. We say that \mathcal{G} satisfies the composite 3-party Diffie-Hellman assumption (C3DH) if for any polynomial time algorithm \mathcal{A} we have that $\text{C3DH Adv}_{\mathcal{G},\mathcal{A}}(\lambda)$ is a negligible function of λ .

The assumption is formed around the intuition that it is hard to test for Diffie-Hellman tuples in the order p subgroup if the elements to be tested have a random order q subgroup component.

5 Hidden Vector Encryption

We construct a Φ -searchable encryption system for a general class of equality predicates. We call such systems Hidden Vector Systems or HVEs for short. We then show in Section 6 that our HVE system leads to comparison and subset queries far more efficient than the trivial system.

5.1 HVE Definition

Let Σ be a finite set and let $*$ be a special symbol not in Σ . Define $\Sigma_* = \Sigma \cup \{*\}$. The star $*$ plays the role of a wildcard or “don’t care” value.² In our subset and range query applications we typically set $\Sigma = \{0, 1\}$.

For $\sigma = (\sigma_1, \dots, \sigma_\ell) \in \Sigma_*^\ell$ define a predicate P_σ^{HVE} over Σ^ℓ as follows. For $x = (x_1, \dots, x_\ell) \in \Sigma^\ell$ set:

$$P_\sigma^{\text{HVE}}(x) = \begin{cases} 1 & \text{if for all } i = 1, \dots, \ell : (\sigma_i = x_i \text{ or } \sigma_i = *), \\ 0 & \text{otherwise} \end{cases}$$

In other words, the vector x matches σ in all the coordinates where σ is not $*$.

Let $\Phi_{\text{HVE}} = \{P_\sigma^{\text{HVE}} \text{ for all } \sigma \in \Sigma_*^\ell\}$. We refer to ℓ as the **width** of the HVE.

Definition 5.1. A *Hidden Vector System (HVE)* over Σ^ℓ is a selectively secure Φ_{HVE} -searchable encryption system.

The case $\ell = 1$ degenerates to the example discussed in Section 2.2 where we showed equivalence to anonymous IBE [9, 1, 13]. For larger ℓ we obtain a more general concept that is much harder to build. In particular, the wildcard character $*$ — which is essential for the applications we have in mind — makes it challenging to construct a Φ_{HVE} -searchable system. We construct an HVE with the following parameters:

$$\text{CT-size} = O(\ell) \quad \text{and} \quad \text{TK-size} = O(\text{weight}(\sigma))$$

where $\text{weight}(\sigma = (\sigma_1, \dots, \sigma_\ell))$ is the number of coordinates where $\sigma_i \neq *$.

5.2 Construction

For our particular HVE construction we will let $\Sigma = \mathbb{Z}_m$ for some integer m . We set $\Sigma_* = \mathbb{Z}_m \cup \{*\}$. We describe an HVE where the payload M is in a subset \mathcal{M} of \mathbb{G}_T where $|\mathcal{M}|/n^{1/2}$ is negligible. This is not a serious restriction since the payload M is typically a short symmetric message key. Our HVE system works as follows:

Setup(λ) The setup algorithm first chooses random primes $p, q > m$ and creates a bilinear group \mathbb{G} of composite order $n = pq$, as specified in Section 4.1. Next, it picks random elements

$$(u_1, h_1, w_1), \dots, (u_\ell, h_\ell, w_\ell) \in \mathbb{G}_p^3, \quad g, v \in \mathbb{G}_p, \quad g_q \in \mathbb{G}_q.$$

and an exponent $\alpha \in \mathbb{Z}_p$. It keeps all these as the secret key SK.

It then chooses $3\ell + 1$ random blinding factors in \mathbb{G}_q :

$$(R_{u,1}, R_{h,1}, R_{w,1}), \dots, (R_{u,\ell}, R_{h,\ell}, R_{w,\ell}) \in \mathbb{G}_q \text{ and } R_v \in \mathbb{G}_q.$$

For the public key, PK, it publishes the description of the group \mathbb{G} and the values

$$g_q, \quad V = vR_v, \quad A = e(g, v)^\alpha, \quad \begin{pmatrix} U_1 = u_1R_{u,1}, & H_1 = h_1R_{h,1}, & W_1 = w_1R_{w,1} \\ & \vdots & \\ U_\ell = u_\ell R_{u,\ell}, & H_\ell = h_\ell R_{h,\ell}, & W_\ell = w_\ell R_{w,\ell} \end{pmatrix}$$

The message space \mathcal{M} is set to be a subset of \mathbb{G}_T of size less than $n^{1/4}$.

²We note that we use the symbol Σ differently than we did in the context of Section 2.1.

Encrypt(PK, $\mathcal{I} \in \mathbb{Z}_m^\ell$, $M \in \mathcal{M} \subseteq \mathbb{G}_T$) Let $\mathcal{I} = (\mathcal{I}_1, \dots, \mathcal{I}_\ell) \in \mathbb{Z}_m^\ell$. The encryption algorithm works as follows:

- choose a random $s \in \mathbb{Z}_n$ and random Z , $(Z_{1,1}, Z_{1,2}), \dots, (Z_{\ell,1}, Z_{\ell,2}) \in \mathbb{G}_q$. (The algorithm picks random elements in \mathbb{G}_q by raising g_q to random exponents from \mathbb{Z}_n .)
- Output the ciphertext:

$$C = \left(C' = MA^s, C_0 = V^s Z, \begin{pmatrix} C_{1,1} = (U_1^{\mathcal{I}_1} H_1)^s Z_{1,1}, & C_{1,2} = W_1^s Z_{1,2} \\ \vdots \\ C_{\ell,1} = (U_\ell^{\mathcal{I}_\ell} H_\ell)^s Z_{\ell,1}, & C_{\ell,2} = W_\ell^s Z_{\ell,2} \end{pmatrix} \right)$$

GenToken(SK, $\mathcal{I}_* \in \Sigma_*^\ell$) The key generation algorithm will take as input the secret key and an ℓ -tuple $\mathcal{I}_* = (\mathcal{I}_1, \dots, \mathcal{I}_\ell) \in \{\mathbb{Z}_m \cup \{*\}\}^\ell$. Let S be the set of all indexes i such that $\mathcal{I}_i \neq *$. To generate a token for the predicate $P_{\mathcal{I}_*}^{\text{HVE}}$ choose random $(r_{i,1}, r_{i,2}) \in \mathbb{Z}_p^2$ for all $i \in S$ and output:

$$\text{TK} = \left(\mathcal{I}_*, K_0 = g^\alpha \prod_{i \in S} (u_i^{\mathcal{I}_i} h_i)^{r_{i,1}} w_i^{r_{i,2}}, \forall i \in S : K_{i,1} = v^{r_{i,1}}, K_{i,2} = v^{r_{i,2}} \right)$$

Query(TK, C) Using the notation in the description of *Encrypt* and *GenToken* do:

- First, compute

$$M \leftarrow C' / \left(e(C_0, K_0) / \prod_{i \in S} e(C_{i,1}, K_{i,1}) e(C_{i,2}, K_{i,2}) \right) \quad (3)$$

- If $M \notin \mathcal{M}$ output \perp . Otherwise, output M .

Correctness Before proving security we first show that the system satisfies the correctness property defined in Section 2.1. Let (\mathcal{I}, M) be a pair in $\Sigma^\ell \times \mathcal{M}$ and let $B_* \in \Sigma_*^\ell$. This B_* defines a predicate P_{B_*} in Φ_{HVE} .

Let $(\text{PK}, \text{SK}) \stackrel{R}{\leftarrow} \text{Setup}(\lambda)$, $C \stackrel{R}{\leftarrow} \text{Encrypt}(\text{PK}, \mathcal{I}, M)$, and $\text{TK} \stackrel{R}{\leftarrow} \text{GenToken}(\text{SK}, B_*)$.

- If $P_{B_*}(\mathcal{I}) = 1$ then a simple calculation shows that $\text{Query}(\text{TK}, C) = M$. This uses in a crucial way the fact that $e(h_p, h_q) = 1$ for all $h_p \in \mathbb{G}_p$ and $h_q \in \mathbb{G}_q$.
- If $P_{B_*}(\mathcal{I}) = 0$ the following lemma shows that $\Pr[\text{Query}(\text{TK}, C) = \perp] \geq 1 - \frac{|\mathcal{M}|}{n}$ where the probability is over the random bits used to create the ciphertext. Hence, if $|\mathcal{M}|$ is sufficiently small compared to n then the probability that $\Pr[\text{Query}(\text{TK}, C) \neq \perp]$ is negligible.

Lemma 5.2. *With the notation as above, If $P_{B_*}(\mathcal{I}) = 0$ then $\Pr[\text{Query}(\text{TK}, C) = \perp] \geq 1 - \frac{|\mathcal{M}|}{n}$ where the probability is over the random bits used to create the ciphertext.*

Proof. Let $\mathcal{I} = (\mathcal{I}_1, \dots, \mathcal{I}_\ell) \in \Sigma$ and let $B_* = (B_1, \dots, B_\ell) \in \Sigma_*^\ell$. Let S be the set of all indexes i such that B_i is not a wildcard $*$ at index i . Since $P_{B_*}(\mathcal{I}) = 0$ we know that there is some $i \in S$ such that $B_i \neq \mathcal{I}_i$. Then the decryption equation (3) contains a factor

$$e(C_0, K_0) / e(C_{i,1}, K_{i,1}) e(C_{i,2}, K_{i,2}) = e(v, u_i)^{(B_i - \mathcal{I}_i) \cdot sr_{i,1}}$$

which is a uniformly distributed value in $\mathbb{G}_{T,p}$ and is independent of the rest of the equation. Since the message space is of size $n^{1/4}$ and the size of $\mathbb{G}_{T,p}$ is approximately $n^{1/2}$, the probability that we get a false positive is negligible in n as required. \square

We note that in practice there is no need to use a small message space $\mathcal{M} \subseteq \mathbb{G}_T$ to determine if decryption succeeded. We only use \mathcal{M} to simplify the description of the system. In practice, one could do the following. The encryptor first picks a random $k \in \mathbb{G}_T$ and derives two independent b -bit symmetric keys (k_0, k_1) from k . It encrypts the payload M using a symmetric encryption system under key k_0 to obtain C_1 . Next, it runs our $Encrypt(\text{PK}, \mathcal{I}, k)$ to obtain C . The final ciphertext is the tuple (C, C_1, k_1) . Now, our algorithm $Query$ works as follows. It first recovers a k' from C using the given token TK. Next, it derives (k'_0, k'_1) from k' and outputs \perp if $k'_1 \neq k_1$. Otherwise, it outputs the decryption of C_1 under k'_0 using a symmetric system. Lemma 5.2 shows that the false error probability is now $1/2^b$. Alternatively, if the symmetric encryption system provides authenticated encryption, then one could decide if $Query$ produced the right value based on whether symmetric decryption succeeded.

Extensions In our description above we limited the index space Σ to be \mathbb{Z}_m . We can expand this space to all of $\{0, 1\}^*$ by taking a large enough m to contain the range of a collision-resistant hash function. Then $Encrypt(\text{PK}, \mathcal{I} \in (\{0, 1\}^*)^\ell, M \in \mathbb{G}_T)$ first hashes all the coordinates of \mathcal{I} into \mathbb{Z}_m using the collision resistant hash and then applies the $Encrypt$ algorithm described above.

5.3 Proof of Security

We prove our scheme selectively secure (as defined in Section 2.3) under the composite 3-party Diffie-Hellman assumption and the bilinear Diffie-Hellman assumption. We give the high-level arguments of the proof in this section and defer the proofs of some lemmas to the appendix.

Suppose the adversary commits to vectors $L_0, L_1 \in \Sigma^\ell$ at the beginning of the game. Let X be the set of indexes i such that $L_{0,i} = L_{1,i}$ and \bar{X} be the set of indexes i such that $L_{0,i} \neq L_{1,i}$.

The proof uses a sequence of $2\ell + 2$ games to argue that the adversary cannot win the original security game of Section 2.3 which we denote by G . We begin by slightly modifying the game G into a game G' . Games G and G' are identical except for how the challenge ciphertext is generated. In G' if $M_0 \neq M_1$ then the adversary multiplies the challenge ciphertext component C' by a random element of $\mathbb{G}_{T,p}$. The rest of the ciphertext is generated as usual. Additionally, if $M_0 = M_1$ then the challenge ciphertext is generated correctly.

Lemma 5.3. *Assume that the Bilinear Diffie-Hellman assumption holds. Then for any polynomial time adversary \mathcal{A} the difference of advantage of \mathcal{A} in game G and game G' is negligible.*

The proof is in Appendix B.1.

Next, we define a game \tilde{G} . In this game the adversary will give two challenge messages, M_0, M_1 . If $M_0 \neq M_1$ then the challenger outputs a random element of \mathbb{G}_T as the C' component of the challenge ciphertext. The rest of ciphertext is constructed as normal. If $M_0 = M_1$ the challenger outputs the challenge ciphertext as normal.

Lemma 5.4. *Assume that the Composite 3-party Diffie-Hellman assumption holds. Then for any polynomial time adversary \mathcal{A} the difference of advantage of \mathcal{A} in game G' and game \tilde{G} is negligible.*

The proof is in Appendix B.2.

Finally, we define two sequences of hybrid games G_j and G'_j for $j = 1, \dots, |\overline{X}|$. We define the game G_j as follows. Let \tilde{X} be a set containing the first j indexes in \overline{X} . The challenger creates the challenge ciphertext components C_0 and $C_{i,1}, C_{i,2}$ as normal for all $i \notin \tilde{X}$. However, for all $i \in \tilde{X}$ the challenger creates $C_{i,1}, C_{i,2}$ as completely random group elements in \mathbb{G} . Additionally, if $M_0 \neq M_1$ then C' is replaced by a completely random element from \mathbb{G}_T (otherwise it is created as normal).

We define a game G'_j as follows. Let \tilde{X} be a set containing the first j indexes in \overline{X} and let δ be the $(j+1)$ -th index in \overline{X} . In the challenge ciphertext the challenger creates C_0 and $C_{i,1}, C_{i,2}$ as normal for all $i \notin \tilde{X}$ and $i \neq \delta$. For all $i \in \tilde{X}$ the challenger creates $C_{i,1}, C_{i,2}$ as completely random group elements in \mathbb{G} . Finally, the challenger chooses a random s' and creates

$$C_{\delta,1} = (u_p^{T_\delta} h_p)^{s'} g_q^{z_{\delta,1}}, \quad C_{\delta,2} = g_p^{s'} g_q^{z_{\delta,2}}.$$

Additionally, if $M_0 \neq M_1$ then C' is replaced by a completely random element from \mathbb{G}_T (otherwise it is created as normal).

Observe that for all i in \tilde{X} the challenge ciphertext contains no information about $L_{\beta,i}$. Therefore the adversary's advantage in game $G_{|\overline{X}|}$ is 0. Additionally, game G_0 is equivalent to \tilde{G} . We state the following two lemmas whose proofs are given in Appendix B.3 and B.4.

Lemma 5.5. *Assume the Composite 3-party Diffie-Hellman assumption holds. Then for all j and any polynomial time adversary \mathcal{A} the difference of advantage of \mathcal{A} in game G_j and game G'_j is negligible.*

Lemma 5.6. *Assume the Composite 3-party Diffie-Hellman assumption holds. Then for all j and any polynomial time adversary \mathcal{A} the difference of advantage of \mathcal{A} in game G'_j and game G_{j+1} is negligible.*

It now follows that if the Composite 3-party Diffie-Hellman and Bilinear Diffie-Hellman assumptions hold then no polynomial-time adversary can break our scheme with non-negligible advantage. This follows from the sequence of hybrid games starting with the original game G :

$$G, \tilde{G}, G'_0, G_1, G_{1'}, G_2, G_{2'}, \dots, G_{|\overline{X}|}.$$

The adversary's advantage in the game $G_{|\overline{X}|}$ is 0 and the difference in adversary's advantage between any two consecutive hybrid games is negligible by the lemmas above. Hence, no polynomial adversary can win game G with non-negligible advantage.

6 Applications of HVE

We show how HVE leads to efficient systems for subset queries and conjunctive comparison queries. Throughout the section we let $\Sigma_{01} = \{0, 1\}$ and $\Sigma_{01*} = \{0, 1, *\}$.

Conjunctive comparison queries. In Section 3.1 we defined conjunctive comparison queries and the predicate family $\Phi_{n,w}$. We use HVE to build a $\Phi_{n,w}$ -searchable encryption system with ciphertext size $O(nw)$ and token size $O(w)$.

Let $(Setup_{HVE}, Encrypt_{HVE}, GenToken_{HVE}, Query_{HVE})$ be a secure HVE over Σ_{01}^{nw} . Thus, the width of this HVE is $\ell = nw$. We construct a $\Phi_{n,w}$ -searchable system as follows:

for all $i = 1, \dots, w$. Then σ is an elements of $(2^T)^w$. Define the predicate $P_\sigma : T^w \rightarrow \{0, 1\}$ as follows:

$$P_\sigma((x_1, \dots, x_w)) = \begin{cases} 1 & \text{if } x_i \in A_i \text{ for all } i = 1, \dots, w, \\ 0 & \text{otherwise} \end{cases}$$

Let $\Phi = \{ P_\sigma \text{ for all } \sigma \in (2^T)^w \}$. Note that Φ is huge — its size is 2^{nw} .

The Φ -searchable system is as follows:

- $Encrypt(\text{PK}, I, M)$ where $I = (x_1, \dots, x_w) \in T^w$. Build a vector $\sigma(I) = (\sigma_{i,j}) \in \Sigma_{01}^{nw}$ as:

$$\sigma_{i,j} = \begin{cases} 1 & \text{if } x_i = j, \\ 0 & \text{otherwise} \end{cases} \quad (6)$$

Then output $Encrypt_{HVE}(\text{PK}, \sigma(I), M)$. The ciphertext size is $O(nw)$ as was the case for comparison queries.

- $GenToken(\text{SK}, \langle P_\alpha \rangle)$ where $\alpha = (A_1, \dots, A_w)$. Define $\sigma_*(\alpha) = (\sigma_{i,j}) \in \Sigma_{01*}^{nw}$ as follows:

$$\sigma_{i,j} = \begin{cases} 0 & \text{if } j \notin A_i, \\ * & \text{otherwise} \end{cases} \quad (7)$$

Output $\text{TK}_\alpha \stackrel{R}{\leftarrow} GenToken_{HVE}(\text{SK}, \sigma_*(\alpha))$. The token size is $O(nw)$, which is bigger than tokens for comparison queries.

- $Setup$ and $Query$ are the same algorithms from the HVE system, as for comparison queries.

It is easiest to see how this works in the one dimensional setting, namely $w = 1$. We encrypt a value $x \in T$ using an HVE vector

$$\sigma(x) = \begin{array}{|c|c|c|c|c|c|c|} \hline & 1 & & x & & & n \\ \hline 0 & \cdots & 0 & 1 & 0 & \cdots & 0 \\ \hline \end{array} \in \{0, 1\}^n$$

Consider a predicate P_A where, for example, $A = \{2, 3, n\} \subseteq T$. We generate a token for P_A by calling $GenToken_{HVE}(\text{SK}, \sigma_*(A))$ using the HVE vector

$$\sigma_*(A) = \begin{array}{|c|c|c|c|c|c|c|} \hline & 1 & 2 & 3 & 4 & 5 & n \\ \hline 0 & * & * & 0 & 0 & \cdots & 0 & * \\ \hline \end{array} \in \{*, 1\}^n$$

The main point is that $x \in A$ if and only if $P_{\sigma_*(A)}^{\text{HVE}}(\sigma(x)) = 1$. Therefore, correctness and security follow from the properties of the HVE. We obtain a secure system for subset queries for *arbitrary subsets*.

Theorem 6.2. (*Setup, Encrypt, GenToken, Query*) is a selectively secure Φ -searchable system assuming (*Setup_{HVE}, Encrypt_{HVE}, GenToken_{HVE}, Query_{HVE}*) is an HVE over Σ_{01}^{nw} .

Note that the trivial system of Section 3 for subset queries produces ciphertexts of size $O(2^n)$. The construction above generates ciphertexts of size $O(n)$.

Subset queries on large domains using Bloom filters. So far we considered subset queries over a domain of size n . In Section 1 we presented examples where one wishes to test a subset relation over a large domain. For example, we discussed email filtering queries of type $(\text{sender} \in S)$ where S is a set of email addresses. To use our construction one would first hash email addressees to a set $\{1, \dots, n\}$ for some n , using a publicly known hash function, and then use the HVE for small domain.

Unfortunately, by hashing into a small domain there is some chance for false positives, namely *Query* may output M even though $(\text{sender} \notin S)$. False positives result from hash collisions. The false positive probability can be reduced by a standard application of Bloom filters [5]. Instead of using one hash function, we use multiple functions $H_1, \dots, H_d : \{0, 1\}^* \rightarrow T$. Again, consider the one-dimensional case, namely $w = 1$. To encrypt a word $W \in \{0, 1\}^*$ the encryptor creates a vector $\sigma(W) \in \{0, 1\}^n$ that contains a ‘1’ at positions $H_1(W), \dots, H_d(W)$ and ‘0’ everywhere else. The encryptor then runs $\text{Encrypt}(\text{PK}, \sigma(W), M)$.

To generate a token for a set $A = \{W_1, \dots, W_s\}$ the *GenToken* algorithm builds a vector $\sigma_*(A) \in \{0, *\}^n$ that contains $*$ at positions $H_i(W_j)$, for all $i = 1, \dots, d$ and $j = 1, \dots, s$, and contains ‘0’ everywhere else. By choosing n and d appropriately, the false positive probability can be made arbitrarily small.

Another subset query application. In our subset query application we identified a ciphertext with an element x and a user’s token with a set A . This allowed us to test whether $x \in A$. We observe that we can easily apply HVE to achieve the opposite semantics where a user’s key is associated with an element x and the ciphertext with a set A . This could be used by a gateway to test if a particular user was one of the (possibly) many receivers of an email. We expect there to be several other applications that one can build with HVE.

7 Extensions

Privacy for search queries. In some cases one may want the token TK_P not to identify which predicate P is being queried. For example, in the anti-spam example from the introduction, the user may not want to reveal his anti-spam predicate to the server. A similar problem was studied by Ostrovsky and Skeith [18] and is related to Private Information Retrieval [16]. For public-key systems supporting comparison queries this is clearly not possible since, given TK_P the server can identify the threshold in P with a simple binary search. It is an open problem to convert our system to a symmetric-key system where TK_P does not expose P . One approach is to simply keep the public key secret from the server; however, is not sufficient in our system.

Validating ciphertexts. Throughout the paper we assumed that the encryptor is honestly creating ciphertexts as specified by the encryption system. For some applications discussed in the introduction (e.g. spam filtering) this may not be the case. By creating malformed ciphertexts an attacker may generate false-positive or false-negatives for the server using the tokens.

Fortunately, in many settings, such as a payment gateway or spam filter, this is easily avoidable. One technique is to do the following. The authority who has SK will also publish a regular public-key PK_1 and ask the encryptor to (i) encrypt the message S with both the searchable system and with PK_1 . The resulting ciphertext is the pair $C = (\text{Encrypt}(\text{PK}, S, 0), \text{Encrypt}_{\text{PKE}}(\text{PK}_1, S))$. When the authority (e.g. visa) receives a ciphertext $C = (C_0, C_1)$ it recovers S from C_1 . It then

uses SK to test that C_0 was in fact created from the message S . If not then the transaction is rejected immediately. Similarly, for anti-spam, if this test fails the email would be immediately rejected as spam. In doing so, the authority ensures that a malformed ciphertext did not fool the server.

Hence, Φ -searchable systems might also provide an algorithm $Test(C, S, M, SK)$ that outputs true when C was generated from the message (S, M) and false otherwise. Our HVE system supports this type of test. Alternatively, one could require the encryptor to prove that his ciphertext is well formed, for example to prove that C_0 is consistent with C_1 . This can be done using non-interactive proof techniques [6, 7], but as mentioned above, often there is no need for this.

8 Conclusion

In public key systems supporting queries on encrypted data a secret key can produce tokens for testing any supported query predicate. The token lets anyone test the predicate on a given ciphertext without learning any other information about the plaintext. We presented a general framework for analyzing security of searching on encrypted data systems. We then constructed systems for comparisons and subset queries as well as conjunctive versions of these predicates.

The underlying tool behind these new constructions is a primitive we call HVE. The one-dimensional version of HVE (namely $\ell = 1$) is essentially an Anonymous IBE system. For large ℓ we obtain a new concept that is extremely useful for a large variety of searching predicates. We note that by setting $\ell = 1$ in our HVE construction we obtain a new simple anonymous IBE system secure without random oracles.

This work poses many challenging open problems. For example, the best non-conjunctive (i.e. $w = 1$) comparison system we currently have requires ciphertexts of size $O(\sqrt{n})$ where n is the domain size. In principal it should be possible to improve this to $O(\log n)$, but this is currently a wide open problem that will require new ideas. Similarly, for non-conjunctive subset queries the best we have requires ciphertexts of size $O(n)$. Again, can this be improved to $O(\log n)$? Our results mostly focus on conjunction. Are there similar results for disjunctive queries? More generally, what other classes of predicates can we search on?

Acknowledgments

We thank Amit Sahai and Alice Silverberg for helpful comments about this work.

References

- [1] Michel Abdalla, Mihir Bellare, Dario Catalano, Eike Kiltz, Tadayoshi Kohno, Tanja Lange, John Malone-Lee, Gregory Neven, Pascal Paillier, and Haixia Shi. Searchable encryption revisited: Consistency properties, relation to anonymous ibe, and extensions. In *CRYPTO*, pages 205–222, 2005.
- [2] Mihir Bellare, Alexandra Boldyreva, and Adam O’Neill. Efficiently-searchable and deterministic asymmetric encryption. <http://eprint.iacr.org/2006/186>, 2006.

- [3] J. Bethencourt, H. Chan, A. Perrig, E. Shi, and D. Song. Anonymous multi-attribute encryption with range query and conditional decryption. Technical report, C.M.U, 2006. CMU-CS-06-135.
- [4] John Bethencourt, Dawn Song, and Brent Waters. New constructions and practical applications for private stream searching. In *Proceeding of 2006 IEEE Symposium on Security and Privacy*, 2006.
- [5] Burton H. Bloom. Space/time trade-offs in hash coding with allowable errors. *Communications of the ACM*, 13:422–426, 1970.
- [6] Manuel Blum, Paul Feldman, and Silvio Micali. Non-interactive zero-knowledge and its applications (extended abstract). In *STOC*, pages 103–112, 1988.
- [7] Manuel Blum, Alfredo De Santis, Silvio Micali, and Giuseppe Persiano. Noninteractive zero-knowledge. *SIAM J. Comput.*, 20(6):1084–1118, 1991.
- [8] Dan Boneh and Xavier Boyen. Efficient selective-ID identity based encryption without random oracles. In *Proceedings of Eurocrypt 2004*, LNCS, pages 223–238. Springer-Verlag, 2004.
- [9] Dan Boneh, Giovanni Di Crescenzo, Rafail Ostrovsky, and Giuseppe Persiano. Public key encryption with keyword search. In *Proceedings of Eurocrypt '04*, 2004.
- [10] Dan Boneh, Eu-Jin Goh, and Kobbi Nissim. Evaluating 2-dnf formulas on ciphertexts. In Joe Kilian, editor, *Proceedings of Theory of Cryptography Conference 2005*, volume 3378 of LNCS, pages 325–342. Springer, 2005.
- [11] Dan Boneh, Amit Sahai, and Brent Waters. Fully collusion resistant traitor tracing with short ciphertexts and private keys. In *Eurocrypt '06*, 2006.
- [12] Dan Boneh and Brent Waters. A fully collusion resistant broadcast trace and revoke system with public traceability. In *ACM Conference on Computer and Communication Security (CCS)*, 2006.
- [13] Xavier Boyen and Brent Waters. Anonymous hierarchical identity-based encryption (without random oracles). In *Crypto '06*, 2006.
- [14] O. Goldreich and R. Ostrovsky. Software protection and simulation by oblivious rams. *JACM*, 1996.
- [15] Philippe Golle, Jessica Staddon, and Brent R. Waters. Secure conjunctive keyword search over encrypted data. In *ACNS*, pages 31–45, 2004.
- [16] Eyal Kushilevitz and Rafail Ostrovsky. Replication is not needed: Single database, computationally-private information retrieval. In *FOCS*, pages 364–373, 1997.
- [17] Rafail Ostrovsky. *Software protection and simulation on oblivious RAMs*. PhD thesis, M.I.T, 1992. Preliminary version in STOC 1990.
- [18] Rafail Ostrovsky and William Skeith. Private searching on streaming data. In *Proceedings of Crypto 2005*, LNCS. Springer, 2005.

- [19] Dawn Song, David Wagner, and Adrian Perrig. Practical techniques for searches on encrypted data. In *Proceedings of the 2000 IEEE symposium on Security and Privacy (S&P 2000)*, 2000.
- [20] Brent Waters, Dirk Balfanz, Glenn Durfee, and Dianna Smetters. Building an encrypted and searchable audit log. In *Proceedings of NDSS '04*, 2004.

A Proof of Lemma 3.1

We prove that the trivial system presented in Section 3 is secure.

Proof. Showing that $\text{QU Adv}_{\mathcal{A}}$ is negligible is a straight forward hybrid argument. Let \mathcal{A} be an adversary playing the query security game. For $i = 1, \dots, n + 1$ we define experiment number i as follows:

- The challenger runs $\text{Setup}(\lambda)$ to obtain

$$\text{PK} \leftarrow (\text{PK}_1, \dots, \text{PK}_n) \quad \text{and} \quad \text{SK} \leftarrow (\text{SK}_1, \dots, \text{SK}_n)$$

It gives PK to \mathcal{A} . Next, \mathcal{A} is given the tokens for any predicates of its choice.

- Then \mathcal{A} outputs two pairs (I_0, M_0) and (I_1, M_1) subject to the restrictions of the query security game challenge phase. For $j = 1, \dots, n$ the challenger constructs the following ciphertexts:

$$C_j \stackrel{R}{\leftarrow} \begin{cases} \text{Encrypt}'(\text{PK}_j, M_0) & \text{if } P_j(I_0) = 1 \text{ and } j \geq i, \\ \text{Encrypt}'(\text{PK}_j, M_1) & \text{if } P_j(I_1) = 1 \text{ and } j < i, \\ \text{Encrypt}'(\text{PK}_j, \perp) & \text{otherwise} \end{cases}$$

The challenger gives $C \leftarrow (C_1, \dots, C_n)$ to \mathcal{A} .

- The adversary continues to adaptively request query tokens subject to the restrictions of the query security game. Finally, \mathcal{A} outputs a bit $\beta' \in \{0, 1\}$ which we denote by $\text{EXP}_{\text{QU}}^{(i)}[\mathcal{A}]$.

A standard argument shows that

$$2 \cdot \text{QU Adv}_{\mathcal{A}} = \left| \text{EXP}_{\text{QU}}^{(1)}[\mathcal{A}] - \text{EXP}_{\text{QU}}^{(n+1)}[\mathcal{A}] \right| \leq \sum_{i=1}^n \left| \text{EXP}_{\text{QU}}^{(i)}[\mathcal{A}] - \text{EXP}_{\text{QU}}^{(i+1)}[\mathcal{A}] \right|$$

But $\left| \text{EXP}_{\text{QU}}^{(i)}[\mathcal{A}] - \text{EXP}_{\text{QU}}^{(i+1)}[\mathcal{A}] \right|$ is clearly negligible assuming \mathcal{E} is semantically secure against chosen plaintext attacks. \square

B Proofs for HVE Construction

The adversary commits to vectors $L_0, L_1 \in \Sigma^\ell$ at the beginning of the game. Let X be the set of indexes i such that $L_{0,i} = L_{1,i}$ and \bar{X} be the set of indexes i such that $L_{0,i} \neq L_{1,i}$. The adversary can issue predicate queries to request a token for any predicate P_L^{HVE} where $L \in \Sigma_*^\ell$. Let S be all the indexes for which L is not a wildcard. We distinguish between three types of queries.

Type 1 For all $S \cap \bar{X} = \emptyset$. That is the predicate does not check any of the indexes in which the challenge tuples differ. These queries can only be made if in the eventual challenge stage $M_0 = M_1$.

Type 2 Case 1 is not met and there exists an $i \in S \cap \bar{X}$ such that $L_i \neq L_{0,i}$ and $L_i \neq L_{1,i}$

Type 3 Case 1 and Case 2 are both not met and there exists an $i \in S \cap X$ such that and $L_i \neq L_{0,i} = L_{1,i}$.

These cases are mutually exclusive (by definition) and complete.

B.1 Proof of Lemma 5.3

We prove our lemma by supposing that a poly-time adversary \mathcal{A} has non-negligible difference ϵ between its advantage in game G and its advantage in game G' . We build a simulator that plays the Bilinear Diffie-Hellman game with advantage ϵ .

The challenger first creates Bilinear Diffie-Hellman challenge as:

$$\begin{aligned} (p, q, \mathbb{G}, \mathbb{G}_T, e) &\stackrel{R}{\leftarrow} \mathcal{G}(\lambda), \quad n \leftarrow pq, \quad g_p \stackrel{R}{\leftarrow} \mathbb{G}_p, \quad g_q \stackrel{R}{\leftarrow} \mathbb{G}_q \\ a, b, c &\stackrel{R}{\leftarrow} \mathbb{Z}_n \\ \bar{Z} &\leftarrow ((n, \mathbb{G}, \mathbb{G}_T, e), g_q, g_p, g_p^a, g_p^b, g_p^c) \\ T &\leftarrow e(g_p, g_p)^{abc} \end{aligned}$$

It then randomly decides whether to give $(\bar{Z}, T' = T)$ or $(\bar{Z}, T' = R)$ where R is a random element in \mathbb{G}_T .

We create the following simulation:

Init The attacker gives the simulator two identities L_0, L_1 . The challenger then flips the coin β internally.

Setup The simulator first chooses random $(R_{u,1}, R_{h,1}, R_{w,1}), \dots, (R_{u,\ell}, R_{h,\ell}, R_{w,\ell}) \in \mathbb{G}_q^3, R_v \in \mathbb{G}_q$ and random $t_1, x_1, y_1, \dots, t_\ell, x_\ell, y_\ell \in \mathbb{Z}_n$.

The simulator first publishes the group description and $g_q, V = g_p R_v$. It lets $A = e(g_p^a, g_p^b)$.

Finally, for all i it creates:

$$U_i = (g_p^b)^{t_i} R_{u,i}, H_i = (g_p^b)^{-t_i L_{\beta,i}} g_p^{y_i} R_{h,i}, W_i = g_p^{x_i} R_{w,i}$$

We observe that the parameters are distributed identically to the real scheme.

Query 1 The adversary will make private key queries to the simulator. The way they are handled depends upon the type of query. Suppose the adversary queries for a key \mathcal{I} , where the set of indexes that are non-wild cards is denoted as S .

Type 1 If the adversary issue a Type 1 query, the simulator simply aborts and takes a random guess. The reason for this is by our definition if a type one query is made then the challenge messages M_0, M_1 will be equal. However, in this case the games G and G' are identical, so there can be no difference in the adversary's advantage when he makes this type of a challenge. Therefore, we can just take a random guess.

Type 2 and Type 3 We handle Type 2 and Type 3 queries in the same manner. The primary intuition is that neither a Type 2 or Type 3 query can be used to distinguish the challenge ciphertext. Suppose the adversary queries for a vector \mathcal{I} and let γ be an (arbitrary) index where $\mathcal{I}_\gamma \neq L_{\beta,\gamma}$.

The simulator first chooses random $r_{i,1}, r_{i,2} \in \mathbb{Z}_n \forall i \in S$. Next it creates:

$$K_0 = \left(\prod_{i \in S} (g_p^b)^{r_{i,1}(\mathcal{I}_i - L_{\beta,i})t_i} g_p^{r_{i,1}y_i} g_p^{r_{i,2}x_i} \right)$$

Additionally, it creates:

$$\forall i \in S / \{\gamma\} : K_{i,1} = (g_p^a)^{r_{i,1}}, K_{i,2} = (g_p^a)^{r_{i,2}}$$

Finally, it creates:

$$K_{\gamma,1} = g_p^{r_{\gamma,1}} (g_p^a)^{-1/(\mathcal{I}_\gamma - L_{\beta,\gamma})} K_{\gamma,2} = g_p^{r_{\gamma,2}}$$

The argument for the well-formness of the keys is similar to that of the Boneh-Boyen [8] Identity-Based Encryption system.

Challenge The adversary first gives the simulator messages M_0, M_1 . If $M_0 = M_1$ we can abort the simulation and take a random guess for the reason given above.

The simulator chooses random $Z \in \mathbb{G}_q (Z_{1,1}, Z_{1,2}), \dots, (Z_{\ell,1}, Z_{\ell,2}) \in \mathbb{G}_q^2$ (this can be done since the simulator has g_q). and outputs the challenge as follows:

$$C' = M_\beta T' C_0 = g^c Z, \quad \forall : C_{i,1} = (g^c)^{y_i} Z_{i,1}, C_{i,2} = (g^c)^{x_i} Z_{i,2}.$$

If T' is forms a tuple, then the simulator is playing game G , otherwise it is playing game G' .

Query Phase 2 Same as Query Phase 1.

Guess The adversary outputs a guess β' . If $\beta = \beta'$ output 0 otherwise output 1. By our assumption the probability that the adversary guesses β correctly in game G_j has a non-negligible ϵ difference from that of it guessing it correctly in game G'_j . However, it is in game G' if and only if the challenger gave the simulator R instead of T . Therefore the simulator has advantage ϵ in the Bilinear Diffie-Hellman game. □

B.2 Proof of Lemma 5.4

We begin by reviewing an assumption called the Bilinear Subgroup Decision problem that was introduced by Boneh, Sahai, and Waters [11] and is implied by the Composite 3-Party Diffie-Hellman assumption.

For a given group generator \mathcal{G} define the following distribution $P(\lambda)$:

$$\begin{aligned} (p, q, \mathbb{G}, \mathbb{G}_T, e) &\stackrel{R}{\leftarrow} \mathcal{G}(\lambda), \quad n \leftarrow pq, \quad g_p \stackrel{R}{\leftarrow} \mathbb{G}_p, \quad g_q \stackrel{R}{\leftarrow} \mathbb{G}_q \\ \bar{Z} &\leftarrow ((n, \mathbb{G}, \mathbb{G}_T, e), \quad g_q, g_p,) \\ T &\leftarrow \mathbb{G}_{T,p} \end{aligned}$$

For an algorithm \mathcal{A} , define \mathcal{A} 's advantage in solving the Bilinear subgroup assumption for \mathcal{G} as:

$$\text{BSD Adv}_{\mathcal{G},\mathcal{A}}(\lambda) := \left| \Pr[\mathcal{A}(\bar{Z}, T) = 1] - \Pr[\mathcal{A}(\bar{Z}, R) = 1] \right|$$

where $(\bar{Z}, T) \xleftarrow{R} P(\lambda)$ and $R \xleftarrow{R} \mathbb{G}$.

Definition B.1. We say that \mathcal{G} satisfies the Bilinear Subgroup Decision assumption if for any polynomial time algorithm \mathcal{A} we have that $\text{BSD Adv}_{\mathcal{G},\mathcal{A}}(\lambda)$ is a negligible function of λ .

It is easy to see that the Composite 3-Party Diffie-Hellman assumption implies the Bilinear Subgroup Decision assumption.³ For simplicity we will use the Decision Subgroup assumption directly in our proof. We suppose that there exist an adversary with non-negligible difference in advantage ϵ between winning the game G' and the game \tilde{G} .

We build a simulator that takes in a Bilinear Subgroup challenge (\bar{Z}, T') . The simulation proceeds as follows.

Init The attacker gives the simulator two identities L_0, L_1 . The challenger then flips the coin β internally.

Setup The simulator setups up the parameters as would the real setup algorithm. All the simulator needs to do this is g_p and g_q from the assumption.

Query Phase 1 The simulator answers queries as the real authority would. One small difference is that the simulator chooses exponents from \mathbb{Z}_n instead of \mathbb{Z}_p . However, this doesn't change anything since the both the simulator and a real authority will raise element from \mathbb{G}_p to the exponents.

Challenge The adversary first gives the simulator messages M_0, M_1 . If $M_0 = M_1$ then the adversary simply encrypts the message to the identity L_β . Otherwise, the simulator creates the challenge ciphertext of message M_β to L_β exactly as normal with the exception that C' is multiplied by T' .

If T' is forms a tuple, then the simulator is playing game G' , otherwise it is playing game \tilde{G} .

Query Phase 2 Same as Query Phase 1.

Guess The adversary outputs a guess β' . If $\beta = \beta'$ output 0 otherwise output 1. By our assumption the probability that the adversary guesses β correctly in game G' has a non-negligible ϵ difference from that of it guessing it correctly in game \tilde{G} . However, it is in game \tilde{G} if and only if the challenger gave the simulator R instead of T . Therefore the simulator has advantage ϵ in the Bilinear Subgroup Decision game which implies an advantage of ϵ in the Composite 3-Party Diffie-Hellman game. □

³One first reverses the labellings of p, q in the Composite 3-Party Diffie-Hellman assumption. Next, we can use the pairing to create an element that will be a random in $\mathbb{G}_{T,p}$ if and only if we were give a well formed tuple. Otherwise the element is random one in \mathbb{G}_T .

B.3 Proof of Lemma 5.5

We prove our lemma by supposing that a poly-time adversary \mathcal{A} has non-negligible difference ϵ between its advantage in game G_j and its advantage in game G'_j for some index j . We build a simulator that plays the Composite 3-Party Diffie-Hellman game with advantage ϵ .

The challenger first creates a 3-Party challenge as:

$$\begin{aligned} (p, q, \mathbb{G}, \mathbb{G}_T, e) &\stackrel{R}{\leftarrow} \mathcal{G}(\lambda), \quad n \leftarrow pq, \quad g_p \stackrel{R}{\leftarrow} \mathbb{G}_p, \quad g_q \stackrel{R}{\leftarrow} \mathbb{G}_q \\ R_1, R_2, R_3 &\stackrel{R}{\leftarrow} \mathbb{G}_q \\ a, b, c &\stackrel{R}{\leftarrow} \mathbb{Z}_n \\ \bar{Z} &\leftarrow ((n, \mathbb{G}, \mathbb{G}_T, e), g_q, g_p, g_p^a, g_p^b, \Gamma = g_p^{ab} \cdot R_1, Y = g_p^{abc} \cdot R_2) \\ T &\leftarrow g_p^c \cdot R_3 \end{aligned}$$

It then randomly decides whether to give $(\bar{Z}, T' = T)$ or $(\bar{Z}, T' = R)$ where R is a random element in \mathbb{G} .

We create the following simulation:

Init The attacker gives the simulator two identities L_0, L_1 . The challenger then flips the coin β internally.

Setup Let δ be the $j + 1$ -th index in \bar{X} .

The simulator first chooses random $(R_{u,1}, R_{h,1}, R_{w,1}), \dots, (R_{u,\ell}, R_{h,\ell}, R_{w,\ell}) \in \mathbb{G}_q^3$ and random $t_1, x_1, y_1, \dots, t_\ell, x_\ell, y_\ell \in \mathbb{Z}_n$.

The simulator first publishes the group description and $g_q, V = \Gamma$. It picks a random $\alpha' \in \mathbb{Z}_n$ and lets $A = e(\Gamma, g_p)^{\alpha'}$. It next creates

$$U_\delta = (g_p^b)^{t_\delta} R_{u,\delta}, H_\delta = (g_p^b)^{-t_\delta L_{\beta,\delta}} g_p^{y_\delta} R_{h,\delta}, W_\delta = g_p^{x_\delta} R_{w,\delta}$$

Finally, for all $i \neq \delta$ it creates:

$$U_i = (g_p^b)^{t_i} R_{u,i}, H_i = (g_p^b)^{-t_i L_{\beta,i}} \Gamma^{y_i} R_{h,i}, W_i = \Gamma^{x_i} R_{w,i}$$

We observe that the parameters are distributed identically to the real scheme.

Query 1 The adversary will make private key queries to the simulator. The way they are handled depends upon the type of query. Suppose the adversary queries for a key \mathcal{I} , where the set of indexes that are non-wild cards is denoted as S .

Type 1 If $\delta \notin S$ then the simulator first chooses random $r_{i,1}, r_{i,2} \in \mathbb{Z}_n \forall i \in S$. Next it creates:

$$K_0 = g_p^\alpha \prod_{i \in S} g_p^{r_{i,1} (\mathcal{I}_i - L_{\beta,i}) t_i} (g_p^a)^{r_{i,1} y_i} g_p^{r_{i,2} x_i}$$

Additionally, it creates:

$$\forall i \in S : K_{i,1} = (g^a)^{r_{i,1}}, K_{i,2} = g_p^{r_{i,2}}$$

Type 2 Suppose $\delta \in S$, but $\mathcal{I}_\delta \neq L_{0,\delta}$ and $\mathcal{I}_\delta \neq L_{1,\delta}$. The simulator first chooses random $r_{i,1}, r_{i,2} \in \mathbb{Z}_n \forall i \in S$. Next it creates:

$$K_0 = \left(\prod_{i \in S/\{\delta\}} g_p^{r_{i,1}(\mathcal{I}_i - L_{\beta,i})t_i} (g_p^a)^{r_{i,1}y_i} g_p^{r_{i,2}x_i} \right) g_p^{r_{\delta,1}(\mathcal{I}_\delta - L_{\beta,\delta})t_\delta x_\delta}$$

Additionally, it creates:

$$\forall i \in S/\{\delta\} : K_{i,1} = (g^a)^{r_{i,1}}, K_{i,2} = g_p^{r_{i,2}}$$

Finally, it creates:

$$K_{\delta,1} = (g_p^a)^{x_\delta r_{\delta,1}} g_p^{x_\delta r_{\delta,2}}, K_{\delta,2} = (g_p^a)^{-y_\delta r_{\delta,1}} (g_p^b)^{t_\delta (\mathcal{I}_\delta - L_{\beta,\delta})} g_p^{-r_{\delta,2}}$$

The keys are distributed as if the randomness for the δ component was:

$$\begin{aligned} \tilde{r}_{\delta,1} &= r_{\delta,1}x_\delta/b + r_{\delta,2}x_\delta/(ab) \pmod{p} \\ \tilde{r}_{\delta,2} &= -y_\delta r_{\delta,1}/b - (y_\delta/ab + t_\delta(\mathcal{I}_\delta - L_{\beta,\delta})/a)r_{\delta,2} \pmod{p} \end{aligned}$$

Since, $\tilde{r}_{\delta,1}, \tilde{r}_{\delta,2}$ are independent the keys generated from the simulation are identical to that of the real scheme.

Type 3 Suppose $\delta \in S$ and $\mathcal{I}_\delta = L_{\beta,\delta}$, but there exists an (arbitrary) index $\gamma \in S$ such that $\mathcal{I}_\gamma \neq L_{\beta,\gamma}$.

The simulator first chooses random $r_{i,1}, r_{i,2} \in \mathbb{Z}_n \forall i \in S$. Next it creates:

$$K_0 = \left(\prod_{i \in S/\{\delta\}} g_p^{r_{i,1}(\mathcal{I}_i - L_{\beta,i})t_i} (g_p^a)^{r_{i,1}y_i} g_p^{r_{i,2}x_i} \right) g_p^{r_{\delta,1}y_\delta y_\gamma}$$

Additionally, it creates:

$$\forall i \in S/(\{\delta\} \cup \{\gamma\}) : K_{i,1} = (g_p^a)^{r_{i,1}}, K_{i,2} = g_p^{r_{i,2}}$$

$$K_{\delta,1} = g_p^{-r_{\delta,2}x_\delta} (g_p^b)^{-r_{\delta,1}(\mathcal{I}_\gamma - L_{\beta,\gamma})t_\gamma}, K_{\delta,2} = g_p^{r_{\delta,2}y_\delta}$$

$$K_{\gamma,1} = (g_p^a)^{r_{\gamma,1}} g_p^{y_\delta r_{\delta,1}}, K_{\gamma,2} = g_p^{r_{\gamma,2}}$$

The keys are distributed as if the randomness for the δ, γ components was:

$$\tilde{r}_{\delta,1} = -r_{\delta,2}x_\delta/(ab) - r_{\delta,1}t_\gamma(\mathcal{I}_\gamma - L_{\beta,\gamma})/a \pmod{p}$$

$$\tilde{r}_{\delta,2} = r_{\delta,2}y_\delta/(ab) \pmod{p}$$

$$\tilde{r}_{\gamma,1} = r_{\gamma,1}/b + y_\delta r_{\delta,1}/(ab) \pmod{p}$$

Since, $\tilde{r}_{\delta,1}, \tilde{r}_{\delta,2}, \tilde{r}_{\gamma,1}$ are independent the keys generated from the simulation are identical to that of the real scheme.

Challenge The adversary first gives the simulator messages M_0, M_1 . Let \bar{X}_j be the first j indexes in \bar{X} . The simulator chooses random $Z \in \mathbb{G}_q$ ($Z_{1,1}, Z_{1,2}, \dots, (Z_{\ell,1}, Z_{\ell,2}) \in \mathbb{G}_q^2$ (this can be done since the simulator has g_q). and outputs the challenge as follows:

$$C_0 = YZ, C_{\delta,1} = T^{y_\delta} Z_{\delta,1}, C_{\delta,2} = T^{x_\delta} Z_{\delta,2}, \quad \forall i \text{ s.t. } i \neq \delta \text{ and } i \notin \bar{H}_j : C_{i,1} = Y^{y_i} Z_{i,1}, C_{i,2} = Y^{x_i} Z_{i,2}.$$

For all $i \in \bar{H}_j$ the simulator chooses random elements in \mathbb{G} for $C_{i,1}, C_{i,2}$. If $M_0 = M_1$ the simulator creates C' as $C' = e(Y, g_p)^{\alpha'} M_0$, otherwise it chooses a random group element for C' .

If T' is forms a tuple, then the simulator is playing game H_j , otherwise it is playing game H'_j .

Query Phase 2 Same as Query Phase 1.

Guess The adversary outputs a guess β' . If $\beta = \beta'$ output 0 otherwise output 1. By our assumption the probability that the adversary guesses β correctly in game G_j has a non-negligible ϵ difference from that of it guessing it correctly in game G'_j . However, it is in game G'_j if and only if the challenger gave the simulator R instead of T . Therefore the simulator has advantage ϵ in the Composite 3-Party Diffie-Hellman game. □

B.4 Proof of Lemma 5.6

We prove our lemma by supposing that a poly-time adversary \mathcal{A} has non-negligible difference ϵ between its advantage in game G'_j and its advantage in game G_{j+1} for some index j . We build a simulator that plays the Composite 3-Party Diffie-Hellman game with advantage ϵ .

The challenger first creates a 3-Party challenge as:

$$\begin{aligned} (p, q, \mathbb{G}, \mathbb{G}_T, e) &\stackrel{R}{\leftarrow} \mathcal{G}(\lambda), \quad n \leftarrow pq, \quad g_p \stackrel{R}{\leftarrow} \mathbb{G}_p, \quad g_q \stackrel{R}{\leftarrow} \mathbb{G}_q \\ R_1, R_2, R_3 &\stackrel{R}{\leftarrow} \mathbb{G}_q \\ a, b, c &\stackrel{R}{\leftarrow} \mathbb{Z}_n \\ \bar{Z} &\leftarrow ((n, \mathbb{G}, \mathbb{G}_T, e), g_q, g_p, g_p^a, g_p^b, \Gamma = g_p^{ab} \cdot R_1, Y = g_p^{abc} \cdot R_2) \\ T &\leftarrow g_p^c \cdot R_3 \end{aligned}$$

It then randomly decides whether to give $(\bar{Z}, T' = T)$ or $(\bar{Z}, T' = R)$ where R is a random element in \mathbb{G} .

We create the following simulation:

Init The attacker gives the simulator two identities L_0, L_1 . The challenger then flips the coin β internally.

Setup Let δ be the $j + 1$ -th index in \bar{X} . The simulator first chooses random

$$R \in \mathbb{G}_q, (R_{u,1}, R_{h,1}, R_{w,1}), \dots, (R_{u,\ell}, R_{h,\ell}, R_{w,\ell}) \in \mathbb{G}_q^3,$$

$\nu \in \mathbb{Z}_n$ and random $t_1, x_1, y_1, \dots, t_\ell, x_\ell, y_\ell$.

The simulator first publishes the group description and $g_q, V = g_p R$. It picks a random $\alpha' \in \mathbb{Z}_n$ and lets $A = e(g_p, g_p)^{\alpha'}$. It next creates

$$U_\delta = (g_p^b)^{t_\delta} R_{u,\delta}, H_\delta = (g_p^b)^{-t_\delta L_{\beta,\delta}} g_p^{y_\delta} R_{h,\delta}, W_\delta = \Gamma R_{w,\delta}$$

Finally, for all $i \neq \delta$ it creates:

$$U_i = (g_p^b)^{t_i} R_{u,i}, H_i = (g_p^b)^{-t_i L_{\beta,i}} g_p^{y_i} R_{h,i}, W_i = g_p^{x_i} R_{w,i}$$

We observe that the parameters are distributed identically to the real scheme.

Query 1 The adversary will make private key queries to the simulator. The way they are handled depends upon the type of query. Suppose the adversary queries for a key L , where the set of indexes that are non-wild cards is denoted as S .

Type 1 If $\delta \notin S$ then the simulator first chooses random $r_{i,1}, r_{i,2} \in \mathbb{Z}_n \forall i \in S$. Next it creates:

$$K_0 = \prod_{i \in S} (g_p^b)^{r_{i,1}(\mathcal{I}_i - L_{\beta,i})t_i} g_p^{r_{i,1}y_i} g_p^{r_{i,2}x_i}$$

Additionally, it creates:

$$\forall i \in S : K_{i,1} = (g_p)^{r_{i,1}}, K_{i,2} = g_p^{r_{i,2}}$$

Type 2 Suppose $\delta \in S$, but $\mathcal{I}_\delta \neq L_{0,\delta}$ and $\mathcal{I}_\delta \neq L_{1,\delta}$. The simulator first chooses random $r_{i,1}, r_{i,2} \in \mathbb{Z}_n \forall i \in S$.

Next it creates:

$$K_0 = \left(\prod_{i \in S/\{\delta\}} (g_p^b)^{r_{i,1}(\mathcal{I}_i - L_{\beta,i})t_i} g_p^{r_{i,1}y_i} g_p^{r_{i,2}x_i} \right) (g_p^a)^{r_{\delta,1}y_\delta} (g_p^b)^{r_{\delta,2}(\mathcal{I}_\delta - L_{\beta,\delta})t_\delta} g_p^{r_{\delta,2}y_\delta}$$

Additionally, it creates:

$$\forall i \in S/\{\delta\} : K_{i,1} = (g_p)^{r_{i,1}}, K_{i,2} = g_p^{r_{i,2}}$$

Finally it creates:

$$K_{\delta,1} = (g_p^a)^{r_{\delta,1}} g_p^{r_{\delta,2}}, K_{\delta,2} = g_p^{-t_\delta(\mathcal{I}_\delta - L_{\beta,\delta})r_{\delta,1}}$$

The keys are distributed as if the randomness for the δ component was:

$$\begin{aligned} \tilde{r}_{\delta,1} &= ar_{\delta,1} + r_{\delta,2} \pmod{p} \\ \tilde{r}_{\delta,2} &= -t_\delta(\mathcal{I}_\delta - L_{\beta,\delta})r_{\delta,1} \pmod{p} \end{aligned}$$

Since, $\tilde{r}_{\delta,1}, \tilde{r}_{\delta,2}$ are independent the keys generated from the simulation are identical to that of the real scheme.

Type 3 Suppose $\delta \in S$ and $\mathcal{I}_\delta = L_{\beta,\delta}$, but there exists an (arbitrary) index $\gamma \in S$ such that $\mathcal{I}_\gamma \neq L_{\beta,\gamma}$.

The simulator first chooses random $r_{i,1}, r_{i,2} \in \mathbb{Z}_n \forall i \in S$.

Next it creates:

$$K_0 = \left(\prod_{i \in S/\{\delta\}} (g_p^b)^{r_{i,1}(\mathcal{I}_i - L_{\beta,i})t_i} g_p^{r_{i,1}y_i} g_p^{r_{i,2}x_i} \right) (g_p^b)^{r_{\delta,2}(\mathcal{I}_\delta - L_{\beta,\delta})t_\delta} g_p^{r_{\delta,2}y_\delta}$$

Additionally, it creates:

$$\forall i \in S/(\{\gamma\} \cup \{\delta\}) : K_{i,1} = (g_p)^{r_{i,1}}, K_{i,2} = g_p^{r_{i,2}}$$

Finally it creates:

$$K_{\delta,1} = g_p^{r_{\delta,2}} (g_p^a)^{r_{\gamma,1}y_\delta/y_\gamma}, K_{\delta,2} = g_p^{-t_\gamma(\mathcal{I}_\gamma - L_{\beta,\gamma})r_{\delta,1}}$$

$$K_{\gamma,1} = (g_p)^{r_{\gamma,1}} (g_p^a)^{r_{\delta,1}}, K_{\gamma,2} = g_p^{r_{\gamma,2}}$$

The keys are distributed as if the randomness for the δ, γ components was:

$$\tilde{r}_{\delta,1} = r_{\delta,2} + ar_{\delta,1}y_\gamma/y_\gamma \pmod{p}$$

$$\tilde{r}_{\delta,2} = -t_\gamma r_{\delta,1}(\mathcal{I}_\gamma - L_{\beta,\gamma}) \pmod{p}$$

$$\tilde{r}_{\gamma,1} = r_{\gamma,1} + ar_{\delta,1} \pmod{p}$$

Since, $\tilde{r}_{\delta,1}, \tilde{r}_{\delta,2}, \tilde{r}_{\gamma,1}$ are independent the keys generated from the simulation are identical to that of the real scheme.

Challenge The adversary first gives the simulator messages M_0, M_1 . Let \bar{X}_j be the first j indexes in \bar{X} . The simulator chooses random $Z \in \mathbb{G}_q$ $(Z_{1,1}, Z_{1,2}), \dots, (Z_{\ell,1}, Z_{\ell,2}) \in \mathbb{G}_q^2$ It also chooses random $s' \in \mathbb{Z}_n$. It outputs the challenge as follows:

$$C_0 = g_p^{s'} Z, C_{\delta,1} = T^{y_\delta} Z_{\delta,1}, C_{\delta,2} = Y^{x_\delta} Z_{\delta,2}, \quad \forall i \text{ s.t. } i \neq \delta \text{ and } i \notin \bar{H}_j : C_{i,1} = g_p^{s'y_i} Z_{i,1}, C_{i,2} = g_p^{s'x_i} Z_{i,2}.$$

For all $i \in \bar{H}_j$ the simulator chooses random elements in \mathbb{G} for $C_{i,1}, C_{i,2}$. If $M_0 = M_1$ the simulator creates C' as $C' = e(g_p, g_p)^{s'\alpha'} M_0$, otherwise it chooses a random group element for C' .

If T is forms a tuple, then the simulator is playing game H'_j , otherwise it is playing game H_{j+1} .

Query Phase 2 Same as Query Phase 1.

Guess The adversary outputs a guess β' . If $\beta = \beta'$ output 0 otherwise output 1. By our assumption the probability that the adversary guesses β correctly in game G'_j has a non-negligible ϵ difference from that of it guessing it correctly in game G_{j+1} . However, it is in game G_{j+1} if and only if the challenger gave the simulator R instead of T . Therefore the simulator has advantage ϵ in the Composite 3-Party Diffie-Hellman game.

C Comparison queries with \sqrt{n} size ciphertext

In this section we focus on the comparison searching problem discussed in Section 3.1 for the special case $w = 1$, namely the case considered in Figure 1. We let n denote the domain size. Recall that the trivial system in this case achieves ciphertext size $O(n)$ as does the system based on Hidden Vector Encryption.

Here, we briefly describe a construction that achieves ciphertext size of \sqrt{n} . Boneh, Sahai, and Waters [11] recently described a tracing traitors system where ciphertext size is \sqrt{n} where n is the number of users in the system. Their construction is based on a general primitive called PLBE (Private Linear Broadcast Encryption). Boneh and Waters [12] recently generalized the construction to obtain a trace and revoke system with ciphertexts having the same size. Their generalization is based on a construction for Augmented Broadcast Encryption (ABE). Setting the recipient set S to $S = \{1, \dots, n\}$ in an ABE system results in a *public* variant of PLBE which we call public-PLBE. The definition of a public-PLBE is implicit in [12]. For completeness, we give the complete definition in Appendix D here. The main result in [12] is an ABE system with the following parameters:

$$\text{CT-size} = \text{Key-size} = \text{PK-size} = O(\sqrt{n})$$

This gives a public-PLBE with similar parameters (by setting $S = \{1, \dots, n\}$). We denote the algorithms in the BW public-PLBE by $(Setup_{PKLBE}, Encrypt_{PKLBE}, Decrypt_{PKLBE})$. We also note that the PLBE of [11] can be easily extended as in [12] to obtain a public-PLBE with parameters

$$\text{Key-size} = O(1), \quad \text{CT-size} = \text{PK-size} = O(\sqrt{n})$$

In Section 3.1 we defined the set of comparison predicates $\Phi_{n,w}$. We show that for $w = 1$, any secure public-PLBE gives a $\Phi_{n,1}$ -searchable encryption as follows:

Setup(λ) Run $Setup_{PKLBE}(n, \lambda)$ to obtain a public key PK and n secret keys (SK_1, \dots, SK_n) .
Output PK and $SK := (SK_1, \dots, SK_n)$.

Encrypt(PK, s, M) where $s \in \{1, \dots, n\}$. Output $C \xleftarrow{R} Encrypt_{PKLBE}(PK, s, M)$.

GenToken(SK, $\langle P \rangle$) A predicate $P \in \Phi_{n,1}$ is a number $i \in \{1, \dots, n\}$. Output $TK \leftarrow (i, SK_i)$.

Query(TK, C) Let $TK = (i, SK_i)$. Run $Decrypt_{PKLBE}(i, SK_i, C)$.

Using a public-PLBE we thus obtain a $\Phi_{n,1}$ -searchable public key encryption where ciphertext size is \sqrt{n} . Security follows easily from the properties of public-PLBE.

Theorem C.1. *The $\Phi_{n,1}$ -searchable encryption system is secure assuming the underlying public-PLBE is secure.*

D Definition of public-PLBE

Boneh and Waters [12] define a primitive called Augmented Broadcast Encryption (ABE) which they use to build a trace and revoke system. Setting the recipient set S to $S = \{1, \dots, n\}$ in an ABE results in a concept we call public-PLBE. For completeness, we include the full definition here.

A public-PLBE is a restricted broadcast system comprising of the following algorithms:

Setup_{PKLBE}(N, λ) A probabilistic algorithm that takes as input N , the number of users in the system, and a security parameter λ . The algorithm runs in polynomial time in λ and outputs a public key PK and private keys SK_1, \dots, SK_N , where SK_u is given to user u .

Encrypt_{PKLBE}(PK, i, M) Takes as input a public key PK, an integer i satisfying $1 \leq i \leq N+1$, and a message M . It outputs a ciphertext C . This ciphertext is intended for users $\{i, i+1, \dots, N\}$.

Decrypt_{PKLBE}(j, SK_j, C) Takes as input the private key SK_j for user j and a ciphertext C . The algorithm outputs a message M or \perp .

The system must satisfy the following **correctness property**: for all $i, j \in \{1, \dots, N+1\}$ (where $j \leq N$), and all messages M :

Let $(PK, (SK_1, \dots, SK_N)) \stackrel{R}{\leftarrow} \text{Setup}_{PKLBE}(N, \lambda)$ and $C \stackrel{R}{\leftarrow} \text{Encrypt}_{PKLBE}(PK, i, M)$.

If $j \geq i$ then $\text{Decrypt}_{PKLBE}(j, SK_j, C) = M$.

Security. We define security of an PKLBE system using two games. The first game is a **message hiding game** and says that a ciphertext created using index $i = N+1$ is unreadable by anyone. The second game is an **index hiding game** and captures the intuition that a broadcast ciphertext created using index i reveals no non-trivial information about i . We will consider all these games for a fixed number of users, N .

Game 1. The first game, called the **Message Hiding Game** says that an adversary cannot break semantic security when encrypting using index $i = N+1$. The game proceeds as follows:

- **Setup** The challenger runs the Setup_{PKLBE} algorithm and gives the adversary PK and all secret keys $\{SK_1, \dots, SK_N\}$.
- **Challenge** The adversary outputs two equal length messages M_0, M_1 . The challenger flips a coin $\beta \in \{0, 1\}$ and sets $C \stackrel{R}{\leftarrow} \text{Encrypt}_{PKLBE}(PK, N+1, M_\beta)$. The challenger gives C to the adversary.
- **Guess** The adversary returns a guess $\beta' \in \{0, 1\}$ of β .

We define the advantage of adversary \mathcal{A} in winning the game as $\text{MH Adv}_{\mathcal{A}} = |\Pr[\beta' = \beta] - 1/2|$.

Game 2. The second game, called the **Index Hiding Game** says that an adversary cannot distinguish between an encryption to index i and one to index $i+1$ without the key SK_i . The game takes as input a parameter $i \in \{1, \dots, N\}$ which is given to both the challenger and the adversary. The game proceeds as follows:

- **Setup** The challenger runs the Setup_{PKLBE} algorithm and gives the adversary PK and the set of private keys $\{SK_j \text{ s.t. } j \neq i\}$.
- **Challenge** The adversary outputs a message M . The challenger flips a coin $\beta \in \{0, 1\}$ and computes $C \stackrel{R}{\leftarrow} \text{Encrypt}_{PKLBE}(PK, i + \beta, M)$. The challenger returns C to the adversary.

- **Guess** The adversary returns a guess $\beta' \in \{0, 1\}$ of β .

We define the advantage of adversary \mathcal{A} as the quantity $\text{IH Adv}_{\mathcal{A}}[i] = |\Pr[\beta' = \beta] - 1/2|$. In words, the game captures the fact that even if all users other than i collude they cannot distinguish whether i or $i + 1$ was used to create a ciphertext C .

With this games we define a secure PKLBE as follows.

Definition D.1. *We say that an N -user public-PLBE system is secure if for all polynomial time adversaries \mathcal{A} we have that $\text{MH Adv}_{\mathcal{A}}$ and $\text{IH Adv}_{\mathcal{A}}[i]$ for $i = 1, \dots, N$, are negligible functions of λ .*