






# Order-Revealing Encryption: File-Injection Attack and Forward Security

Xingchen Wang<sup>1,2,3</sup>  and Yunlei Zhao<sup>1,2,3</sup>  

<sup>1</sup> Shanghai Key Laboratory of Data Science, School of Computer Science,  
Fudan University, Shanghai, China

{xingchenwang16, ylzhaol}@fudan.edu.cn

<sup>2</sup> State Key Laboratory of Integrated Services Networks,  
Xidian University, Xi'an, China

<sup>3</sup> State Key Laboratory of Cryptology,  
Beijing, China

**Abstract.** Order-preserving encryption (OPE) and order-revealing encryption (ORE) are among the core ingredients for encrypted databases (EDBs). In this work, we study the leakage of OPE and ORE and their forward security.

We propose generic yet powerful file-injection attacks (FIAs) on OPE/ORE, aimed at the situations of possessing *order by* and range queries. Our FIAs only exploit the *ideal* leakage of OPE/ORE (in particular, no need of data denseness or frequency). We executed some experiments on real datasets to test the performance, and the results show that our FIAs can cause an extreme hazard on most of the existing OPEs and OREs with high efficiency and 100% recovery rate.

We then formulate forward security of ORE, which is of independent of interest, and propose a practical compilation framework for achieving forward secure ORE in order to resist the perniciousness of FIA. The compilation framework can transform most of the existing OPEs/OREs into forward secure OREs, with the goal of minimizing the extra burden incurred on computation and storage. We also execute some experiments to analyze its performance.

**Keywords:** Order-revealing encryption  
Order-preserving encryption · File-injection attack · Forward security

---

This work is supported in part by National Key Research and Development Program of China under Grant No. 2017YFB0802000, National Natural Science Foundation of China under Grant Nos. 61472084 and U1536205, Shanghai Innovation Action Project under Grant No. 16DZ1100200, Shanghai Science and Technology Development Funds under Grant No. 6JC1400801, and Shandong Provincial Key Research and Development Program of China under Grant Nos. 2017CXG0701 and 2018CXGC0701.

## 1 Introduction

In recent years, many property-preserving encryption (PPE) schemes and property-revealing encryption (PRE) schemes have been proposed with increased efficiency or/and security. This condition promotes the occurrence of encrypted database (EDB) systems. CryptDB [22] has been proposed by Popa et al. as the first practical EDB system for executing data manipulations on encrypted data. Because of its onion encryption model and its proxy architecture, CryptDB supports most of the basic operations on ciphertexts with acceptable efficiency.

As a kind of PPE, order-preserving encryption (OPE) has been gaining more and more attention and studies because of its applications on EDB. OPE was first proposed for numeric data by Agrawal et al. [3], where the order of plaintexts can be obtained by comparing their ciphertexts directly. Later, order-revealing encryption (ORE) was proposed by Boneh et al. [6] as the generalization of OPE, where the ciphertexts reveal their order by a special algorithm rather than comparing themselves directly.

Even though OPE and ORE aim at leaking nothing other than the order of ciphertexts, many attacks have been proposed against OPE and ORE in recent years. Naveed et al. [20] proposed several inference attacks against the deterministic encryption (DTE) and OPE in CryptDB. Durak et al. [10] showed that some ORE schemes, whose security is discussed on uniform inputs, could make the plaintext recovery of some known attacks more accurate on nonuniform data. They also proposed an attack, aiming at multiple encrypted columns of correlated data, which reveals more information than prior attacks against columns individually. Grubbs et al. [13] proposed new leakage-abuse attacks that achieve high-correctness recovery on OPE-encrypted data. They also presented the first attack on frequency-hiding OPE proposed in [16].

### 1.1 Our Contributions

In this paper, we first demonstrate the power of file-injection attacks (FIAs) on OPE/ORE, by developing two categories of FIA schemes (to the best of our knowledge, the first such attacks) against OPE/ORE. (The underlying assumptions and work flows of FIAs are briefly described in Sect. 3.1). Our FIA attacks are generic and powerful, in the sense that they only exploit the *ideal* leakage of OPE/ORE. Specifically, for our FIA attacks to work, the adversary only possesses the plaintext space, some old *order by* or range queries and the corresponding cipher result sets returned from EDB. In particular, the adversary does not need either the ability of comparing ciphertexts with the ORE comparison algorithm, or that of obtaining the ciphertexts outside of the result sets for *order by* and range queries.

In comparison with other attacks against OPE/ORE proposed in recent years, our FIA attacks rely upon less demanding conditions, and are more effective (particularly for attacking systems, like encrypted email systems, with the function of data sharing or transferring). For example, compared with the attacks

against OPE/ORE proposed in [13, 20], our FIA attacks have the following features simultaneously: (1) no need of data denseness or frequency, and (2) generic against any OPE/ORE with *ideal* leakage. Furthermore, as shown in Appendix A, we compare and clarify in detail the advantages of our attacks over the chosen-plaintext attack (CPA) and the inference attack (IA). We present more details and experiments in the extended version [25] of this paper. Moreover, the experiment results show that our FIAs can cause an extreme hazard on most of the existing OPE and ORE schemes with high efficiency and 100% recovery rate.

The strong security property against FIA is forward security, which ensures that the previous data manipulations do not cause any leakage of the newly inserted data. In other words, it is infeasible for the server to produce a correct response when applying an old query to newly inserted ciphertexts encrypted by a forward secure scheme. To the best of our knowledge, no OPE/ORE construction offered the forward security to thwart FIAs up to now.

In this work, we give the formal definition of forward security for OPE/ORE, which might be of independent interest. Then, we propose a compilation framework for achieving forward secure ORE schemes against FIA attacks. Specifically, the compilation framework is applicable to most of the existing OPE/ORE schemes to transform them into forward secure ones. The resultant forward secure schemes leak nothing about newly inserted data that match the previous *order by* or range queries. Moreover, the compilation framework is constructed with the goal of minimizing the extra burden incurred on computation and storage. In particular, the compilation only uses some simple cryptographical tools like pseudo-random function (PRF), keyed hash function and trapdoor permutation (TDP). Finally, we execute some experiments to analyze the additional cost caused when applying our compilation framework to some prominent OPE/ORE schemes developed in recent year.

## 1.2 Related Work

**Order-Preserving Encryption (OPE).** Agrawal et al. [3] first proposed an OPE scheme for numeric data. Afterwards, OPE was formally studied by Boldyreva et al. [4], where, in particular, two leakage profiles were introduced. Boldyreva et al. [5] analyzed the one-wayness security of OPE, and showed that any OPE scheme must have immutable large ciphertexts if the scheme is constructed for leaking only order and frequency information. Popa et al. [21] proposed an OPE scheme in order tree structure, which is the first OPE scheme achieving the security of IND-OCOA (indistinguishability under ordered chosen-plaintext attack). Kerschbaum [16] proposed a frequency-hiding OPE scheme, which supports the security of IND-FA-OCOA (indistinguishability under frequency-analyzing ordered chosen-plaintext attack) for the first time. Later, a partial order preserving encryption (POPE), with a method for frequency-hiding, was developed by Roche et al. [23].

**Order-Revealing Encryption (ORE).** ORE was first generalized from OPE by Boneh et al. [6]. Their ORE scheme is built upon multilinear maps, which

provides better security but at the cost of worse efficiency. Chenette et al. [9] proposed the first practical ORE, which achieves a simulation-based security w.r.t. some leakage functions that precisely quantify what is leaked by the scheme. Recently, Cash et al. [8] presented a general construction of ORE with reduced leakage as compared to [9], but at the cost of using a new type of “property-preserving” hash function based on bilinear maps.

**File-Injection Attack on SSE.** As a kind of query-recovery attack, Islam et al. [15] initiated the study of FIA attack against searchable symmetric encryption (SSE), by showing that a curious service provider can recover most of the keywords-search queries with high accuracy. Cash et al. [7] further improved the power of the attack initiated in [15], by assuming less knowledge about the files of clients even in a larger plaintext space. Except the encrypted email systems like Pmail [2], they also discussed how their active attacks (e.g., query recovery attacks, partial plaintext recovery attacks, FIAs) might be used to break through other systems such as the systems in [14, 17]. Zhang et al. [26] showed that FIA can recover the keywords-search queries with just a few injected files even for SSE of low leakage. Their attacks outperform the attacks proposed in [7, 15] in efficiency and in the prerequisite of adversary’s prior knowledge.

## 2 Preliminaries

In this section we introduce some fundamental knowledge of TDP, ORE and OPE. We use standard notations and conventions below for writing probabilistic algorithms, experiments and protocols. If  $\mathcal{D}$  denotes a domain,  $x \stackrel{\$}{\leftarrow} \mathcal{D}$  is the operation of picking an element uniformly at random from  $\mathcal{D}$ . If  $\mathbf{S}$  is a set, then for any  $k$ ,  $0 \leq k \leq |\mathbf{S}| - 1$ ,  $\mathbf{S}[k]$  denotes the  $(k + 1)$ -th element in  $\mathbf{S}$ . If  $\alpha$  is neither an algorithm nor a set, then  $x \leftarrow \alpha$  is a simple assignment statement. If  $A$  is a probabilistic algorithm, then  $A(x_1, x_2, \dots; r)$  is the result of running  $A$  on inputs  $x_1, x_2, \dots$  and coins  $r$ . We let  $A(x_1, x_2, \dots) \rightarrow y$  denote the experiment of picking  $r$  at random and letting  $y$  be  $A(x_1, x_2, \dots; r)$ . By  $\mathbb{P}[R_1; \dots; R_n : E]$  we denote the probability of event  $E$ , after the ordered execution of random processes  $R_1, \dots, R_n$ .

**Definition 1 (Trapdoor Permutation).** *A tuple of polynomial-time algorithms  $(\text{KeyGen}, \Pi, \text{Inv})$  over a domain  $\mathcal{D}$  is a family of trapdoor permutations (or, sometimes, a trapdoor permutation informally), if it satisfies the following properties:*

- $\text{KeyGen}(1^\lambda) \rightarrow (I, \text{td})$ . On input a secure parameter  $\lambda$ , the parameter generation algorithm outputs a pair of parameters  $(I, \text{td})$ . Each pair of the parameters defines a set  $\mathcal{D}_I = \mathcal{D}_{\text{td}}$  with  $|I| \geq \lambda$ . Informally,  $I$  (resp.,  $\text{td}$ ) is said to be the public key (resp., secret key) of TDP.
- $\text{KeyGen}_1(1^\lambda) \rightarrow I$ . Let  $\text{KeyGen}_1$  be the algorithm that executes  $\text{KeyGen}$  and returns  $I$  as the only result. Then  $(\text{KeyGen}_1, \Pi)$  is a family of one-way permutations.

- $\text{Inv}_{\text{td}}(y) \rightarrow x$ .  $\text{Inv}$  is a deterministic inverting algorithm such that, for every pair of  $(I, \text{td})$  output by  $\text{KeyGen}(1^\lambda)$  and any  $x \in \mathcal{D}_{\text{td}} = \mathcal{D}_I$  and  $y = \Pi_I(x)$ , it holds  $\text{Inv}_{\text{td}}(y) = x$ . For presentation simplicity, we also write the algorithm  $\text{Inv}_{\text{td}}$  as  $\Pi_{\text{td}}^{-1}$ , and denoted by  $\Pi_I^k(x) = \overbrace{\Pi_I(\Pi_I(\dots \Pi_I(x)\dots))}^k \text{ TDPs}$  for some integer  $k \geq 1$ .

## 2.1 Definition of ORE

**Definition 2 (Order-Revealing Encryption).** A secret-key encryption scheme is an order-revealing encryption (ORE), if the scheme can be expressed as a tuple of algorithms  $\text{ORE} = (\text{ORE.Setup}, \text{ORE.Encrypt}, \text{ORE.Compare})$  which is defined over a well-ordered domain  $\mathcal{M}$ .

- $\text{ORE.Setup}(1^\lambda) \rightarrow (pp, sp)$ . On input of a secure parameter  $\lambda$ , the setup algorithm outputs the set of public parameters  $pp$  and the set of secret parameters  $sp$  which includes the secret key for encryption algorithm.
- $\text{ORE.Encrypt}(pp, sp, m, \sigma_1) \rightarrow c$ . On input of  $pp$ ,  $sp$  and a set  $\sigma_1$  of other auxiliary parameters (that are not generated in the setup algorithm), the encryption algorithm encrypts the input plaintext  $m \in \{0, 1\}^*$  to a ciphertext  $c$  that can reveal the correct order with other ciphertexts.
- $\text{ORE.Compare}(pp, sp, c_1, c_2, \sigma_2) \rightarrow b$ . On input of  $pp$ ,  $sp$ , two ciphertexts  $c_1, c_2$ , and the set  $\sigma_2$  of other auxiliary parameters, the comparison algorithm returns a bit  $b \in \{0, 1\}$  as the result of order.

The ORE definition in other literature is simple and only remains the necessary parameters. Our definition above is more complex, and the additional parameters are used for better describing the latter framework. With the above formulation, we aim for a generic and basic definition of ORE, where  $\sigma_1$  and  $\sigma_2$  may depend upon and vary with the concrete implementations of ORE. As a consequence, we do not introduce many details (that may vary with different implementations) and components like clients for interactive queries (as our FIAs are w.r.t. the generic OPE/ORE structure).

**Leakage Profiles.** The *ideal* leakage profile, the *random order-preserving function* profile, the *most significant-differing bit* profile, the *RtM* profile and the *MtR* profile are five leakage profiles that have been proposed in the literature. The first two were described by Boldyreva et al. [4], and the others were described by Chenette et al. [9].

We remark that, in Sect. 3, our FIAs are generic in the sense that they are constructed only with the *ideal* leakage profile. The *ideal* leakage profile just reveals the order and the frequency of the plaintexts. More precisely, only the leakage of order is necessary for our FIAs.

An adversary is said to be adaptive, if it is allowed to adaptively select data to be encrypted by the clients and then stored back to the server. Roughly speaking, an ORE scheme is said to be  $\mathcal{L}$ -adaptively-secure, if any probabilistic polynomial-time (PPT) adaptive adversary cannot learn more than the leakage as described according to the leakage profile  $\mathcal{L}$ .

## 2.2 Definition of OPE

Order-preserving encryption (OPE) is a simplified case of ORE. The ciphertext domain  $\mathcal{C}$  of OPE needs to be well-ordered exactly as the plaintext domain  $\mathcal{M}$ .

**Definition 3 (Order-Preserving Encryption).** *A secret-key encryption scheme is an order-preserving encryption (OPE), if the scheme can be expressed as a tuple of algorithms  $OPE = (OPE.Setup, OPE.Encrypt)$ , which is defined over a well-ordered plaintext domain  $\mathcal{M}$  and a well-ordered ciphertext domain  $\mathcal{C}$ .*

- $OPE.Setup(1^\lambda) \rightarrow (pp, sp)$ . On input of a secure parameter  $\lambda$ , the setup algorithm outputs the set of public parameters  $pp$  and the set of secret parameters  $sp$  which includes the secret key for encryption algorithm.
- $OPE.Encrypt(pp, sp, m, \sigma_1) \rightarrow c$ . On input of  $pp$ ,  $sp$ , and a set  $\sigma_1$  of other auxiliary parameters, the encryption algorithm encrypts the input plaintext  $m$  to a ciphertext  $c$  that preserves the correct order with other ciphertexts.

## 3 File-Injection Attacks on OPE/ORE

### 3.1 Assumptions and Basic Workflow

File injection attack has the following five assumptions: (1) The target system has a dependable component used for data-sharing or data-transmitting; (2) The adversary possesses the plaintext space of the target ciphertexts, and can store correct ciphertexts by sending some forged data to the client without suspicion; (3) The adversary possesses some old encrypted queries and can obtain the correct result sets from the server; (4) The adversary can only get the ciphertexts included in the result sets. (If the plaintext injected by the adversary does not match the queries, the corresponding ciphertext will not be known to it;) (5) The adversary is unable to forge queries or execute any PPE/PRE algorithm.

The basic workflow of FIA is briefly described as following:

- First, the adversary forges some data and sends them to the client from the server. After being encrypted by the client, the resultant ciphertexts of the forged data are sent back to the server for storing.
- Second, the adversary replays some old queries and infers the responses from the database management system (DBMS) with the leakage of newly inserted data.
- Third, the adversary adaptively executes the first two steps repeatedly. And the data will be recovered successfully when the adversary obtains enough leakage.

In some application scenarios like encrypted email system (e.g., Pmail [2]) or the systems in [14, 17], FIA can be easily executed. Assuming that the server has already responded many email-order requests and recorded many encrypted data manipulation statements, the adversary can forge some emails and send to the client. When the new emails are encrypted and sent back to the DBMS, the

adversary can take advantage of the entire set of ciphertexts, as well as the old queries, to collect more leakage and infer the corresponding plaintexts.

Unlike the FIA attacks against SSE, our FIA attacks against OPE/ORE are data-recovery attacks, which are more powerful. Moreover, the forged data are less likely to be detected because of the smaller forged part. Furthermore, by extending the concept of FIA, with our FIA attacks files do not only represent the data elements in NoSQL database, but can also be any kind of data which fit the target system.

**Table 1.** Notations in Sect. 3


Notation	Meaning
$m, c$	Instance variables of plaintext, ciphertext
$\mathcal{M}, \mathcal{C}$	Ordered spaces of plaintexts and ciphertexts
$\mathbf{M}, \mathbf{C}$	Set of plaintexts and set of ciphertexts
$\omega$	Adversary makes at most $\omega$ file-injections
$q, \mathbf{Q}$	Instance variable of query and set of queries
$\phi$	The flag variable which shows whether the plaintext of the target ciphertext has been recovered
$i$	The record of counter that is used for efficiency analysis in our experiments

### 3.2 Notations

Table 1 lists the meaning of some simple notations, which is helpful to comprehend the two FIA algorithms against *ideal*-secure OPEs/OREs presented in Sect. 3. Let  $\mathbf{R}_q^i$  and  $\mathbf{R}_q$  denote the result set of query  $q$  before the  $(i + 1)$ -th file-injection and the current result set of query  $q$ . Let  $c \xleftarrow{\text{file injection}} m$  denote the process in which the adversary sends the forged plaintext  $m$  from the server to the client and the resultant ciphertext  $c$  is sent back (by client) and stored in the EDB. Let  $a$  and  $b$  denote the indices of data which show their locations in their domains or their sets. Let  $\text{mid}(a, b)$  denote an arbitrary scheme for efficient median calculation, regardless of the round-off method. Let  $\mathbf{d}$  and  $\mathbf{d}\text{queue}$  denote a structural body contains two indices  $(a, b)$  and a queue of the structural body. Let  $m_l$  and  $m_r$  (resp.,  $q.c_l$  and  $q.c_r$ ) denote the left plaintext (resp., ciphertext) and right plaintext (resp., ciphertext) boundary values of range condition in a range query  $q$ . We use the composite notation to represent the main part which is related to the additional part. Hence, we let  $m_c$  denote the plaintext of the ciphertext  $c$ , let  $\mathcal{M}_{a,b}$  denote the plaintext space between  $a$  and  $b$ , let  $\mathbf{d}.a$  and  $\mathbf{d}.b$  denote the parameters  $a$  and  $b$  in the structural body  $\mathbf{d}$ . Let  $\mathcal{M}[\text{mid}(a, b)]$  denotes the  $(k + 1)$ -th element in  $\mathcal{M}$  for  $k = \text{mid}(a, b)$ .

### 3.3 Binary Search

The two FIA algorithms presented below are based on a common algorithm – binary search. The difference between the two FIA algorithms lies in the search



**Fig. 1.** Depth first binary search (a) and breadth first binary search (b). (Color figure online)

types they employ: one uses the traditional binary search like the depth first traversal, and another uses the breadth first traversal. The traditional binary search is a kind of (depth-first like) search algorithm, which finds the position of a target value within a sorted array by testing the order of the target value and the median value. In this work, we import the idea of breadth first traversal in the second FIA algorithm, with which we can get the relatively near data (around the target) that does not match the range condition.

We show two types of binary search in Fig. 1, where the colored nodes are the passed nodes with their order marked, and the crosses mark the target nodes. In the second FIA algorithm, our FIA attacker, with the range query determined by  $(m_l, m_r)$ , needs to find a value  $m_1$  matching the range condition, and a pair of relatively near unmatched values  $(m_2, m_3)$  in the file-injected dataset, such that  $m_2 < m_l < m_1 < m_r < m_3$ . The details are presented in Sect. 3.5.

### 3.4 Basic FIA with *order by* Query

Our FIA attacks use two kinds of order queries respectively: *order by* queries and range queries. The *order by* query (e.g., *select \* from table\_1 order by column\_1*, which ensures that the result data are ordered, is one of two Data Manipulation Languages (DMLs) that are based on the order of data. And the other one is the range query with relational operators like “<”, “>” and so on. In Sects. 3.4 and 3.5, we present the attack models and the FIA algorithms, assuming the attacker possesses these two kinds of order queries respectively.

The attack model of basic FIA, with *order by* queries, consists of the adversarial information (i.e., leakage) and the adversarial goal. As to the adversarial information, we limit the power of adversaries in order for more practical attacks in practice. Specifically, the adversary only possesses, as adversarial information, the plaintext space  $\mathcal{M}$ , the set  $\mathbf{Q}$  of old *order by* queries, and the result sets of those queries with forged data. In particular, they do not have any information about the data not in the result sets of the old queries. About the adversarial goal, we partition it into two types: recovering the plaintext of a single ciphertext,



and recovering the plaintexts of all the ciphertexts in the result sets. This partition facilitates the discussion of time complexity as we show later. We formalize the attack model as following:

$$\text{Leakage: } \mathcal{L}(\mathcal{M}, \mathbf{Q}, \mathbf{R}_{\mathbf{Q}} = \{ \bigcup_{q \in \mathbf{Q}, 0 \leq i \leq \omega} \mathbf{R}_{q|\text{ordered}}^i \})$$

$$\text{Goal: } m_c (c \in \mathbf{R}_{q|\text{ordered}}^0, q \in \mathbf{Q}) \text{ or } \mathbf{M}_{\mathbf{C}} (\mathbf{C} = \bigcup_{q \in \mathbf{Q}} \mathbf{R}_q)$$

where  $\mathbf{R}_{q|\text{ordered}}^i$  denotes the *ordered* result set for *order by* query  $q$  before the  $(i+1)$ -th file-injection,  $\mathbf{R}_{q|\text{ordered}}^0$  denotes the original *ordered* result set for *order by* query  $q$ ,  $\mathbf{M}_{\mathbf{C}}$  ( $\mathbf{C} = \bigcup_{q \in \mathbf{Q}} \mathbf{R}_q$ ) denotes the plaintext set  $\mathbf{M}_{\mathbf{C}}$  corresponding to the ciphertext set  $\mathbf{C}$  in the current result sets for all the queries in  $\mathbf{Q}$ . Here,  $\mathbf{M}_{\mathbf{C}}$  can also be expressed as a mapping relation precisely, denoted  $\mathcal{T}_{(\mathbf{M}, \mathbf{C})}$ , between all the ciphertexts in  $\mathbf{C}$  (which includes all the original and forged data) and their corresponding plaintexts in  $\mathbf{M}_{\mathbf{C}}$ .

---

**Algorithm 1.**

$$m_c \leftarrow \text{FIA\_Orderby}(\mathcal{M}, c \in \mathbf{R}_{q|\text{ordered}}^0, q)$$


---

```

1:  $a \leftarrow -1, \phi \leftarrow 0, b \leftarrow |\mathcal{M}|$ 
2: for  $i \leftarrow 1$  to  $\infty$  do
3:    $c_i \xleftarrow{\text{file injection}} \mathcal{M}[\text{mid}(a, b)]$ 
4:   if  $(\text{Comp}(c_i, c) = 0)$ 
5:      $\phi \leftarrow 1, \text{break}$ 
6:   else if  $(\text{mid}(a, b) = a \text{ or } \text{mid}(a, b) = b)$  break
7:   else if  $(\text{Comp}(c_i, c) = 1)$ 
8:      $b \leftarrow \text{mid}(a, b)$ 
9:   else  $a \leftarrow \text{mid}(a, b)$ 
10:  end if
11: end for
12: if  $(\phi = 1)$  return  $m_c \leftarrow \mathcal{M}[\text{mid}(a, b)]$ 
13: else return  $\perp$ 

```

---

For ease of comprehension, Algorithm 1 describes the elementary FIA based on utilizing a single *order by* query over an entire dataset. The adversary will continually detect the plaintext of the target ciphertext  $c$  with an old query  $q$  by file-injections. We use  $\text{Comp}(c_i, c)$  to express the order result of query  $q$  about the target ciphertext  $c$  and the  $i$ -th injected ciphertext  $c_i$ , where the result expresses as following:

$$\text{Comp}(c_i, c) = \begin{cases} 0 & c_i = c \\ 1 & c_i > c \\ -1 & c_i < c. \end{cases}$$

**Time Complexity.** The time complexity of Algorithm 1 is  $O(\log|\mathcal{M}|)$  obviously in the worst condition for recovering one plaintext. When the adversarial goal is

to recover all the  $N$  nonrepetitive ciphertexts in the entire result set, the time complexity is  $O(N\log|\mathcal{M}| - N\log N)$  in the worst case. This means, in this case, the average time complexity of recovering a single ciphertext becomes smaller because the order of a ciphertext can be used for both sides. In other words, a file-injection for a target will reveal some order information about other target ciphertexts as well.

In Algorithm 1, we only take advantage of the leakage  $\mathcal{L}_1(\mathcal{M}, q, \mathbf{R}'_q)$ , where  $\mathbf{R}'_q = \mathbf{R}_q \setminus \mathbf{R}_q^0$  is the result set after file-injections excluding the original result set. Because the leakage of the original result set  $\mathbf{R}_q^0$  is in the *ideal* leakage profile, we can only get some order information between the target ciphertext  $c_{\text{target}}$  and other ciphertexts. In other words, we can rewrite the original result set as

$$\mathbf{R}_q^0 = \{\mathbf{C}_{\text{ordered}}^-, c_{\text{target}}, \mathbf{C}_{\text{ordered}}^+\}$$

where  $\mathbf{C}_{\text{ordered}}^-$  is the set of ordered ciphertexts which are smaller than the target, and  $\mathbf{C}_{\text{ordered}}^+$  is the set of ordered ciphertexts which are greater than the target. Under the assumption of knowing nothing about the original ciphertexts except their order information, we can only take advantage of  $|\mathbf{C}_{\text{ordered}}^-|$  and  $|\mathbf{C}_{\text{ordered}}^+|$  to curtail the plaintext space. We delete the first  $|\mathbf{C}_{\text{ordered}}^-|$  plaintexts and the last  $|\mathbf{C}_{\text{ordered}}^+|$  plaintexts from the ordered plaintext space  $\mathcal{M}$ , and then we get a smaller new plaintext space  $\mathcal{M}'$  for the target  $c_{\text{target}}$ . Thus, the time complexity of recovering a single ciphertext becomes  $O(\log|\mathcal{M}'|)$  which is even smaller now. In this way, the adversary can adaptively curtail the plaintext space according to the number of ciphertexts on both sides after each file-injection.

Moreover, an improved method with hierarchical idea is presented in the extended version [25] of this paper.

### 3.5 FIA with Range Queries

The attack model of FIA with range queries also consists of the adversarial information and the adversarial goal. As to the adversarial information, the adversary just has the plaintext space  $\mathcal{M}$ , the old range queries in  $\mathbf{Q}$ , and the result sets of those queries without inner order. In this condition, the leakage is less than that with *order by* queries, because the adversary only knows the result set matching the range conditions without knowing the inner order. As to the adversarial goal, the adversary needs to recover the boundary plaintexts of the range conditions as well as all the plaintexts matching the range conditions. We formalize the attack model as following:

$$\text{Leakage: } \mathcal{L}(\mathcal{M}, \mathbf{Q}, \mathbf{R}_{\mathbf{Q}} = \{ \bigcup_{q \in \mathbf{Q}, 0 \leq i \leq \omega} \mathbf{R}_q^i \})$$

$$\text{Goal: } \mathbf{M}_l, \mathbf{M}_r, \mathbf{M}_{\mathbf{C}} \ (\mathbf{C} = \{c \mid q.c_l < c < q.c_r, q \in \mathbf{Q}\})$$

where  $\mathbf{M}_{\mathbf{C}}$  can be expressed as a mapping relation precisely, denoted  $\mathcal{T}_{(\mathbf{M}, \mathbf{C})}$ , between all the ciphertexts in  $\mathbf{C}$  (which includes all the original and forged data) and their plaintexts in  $\mathbf{M}_{\mathbf{C}}$ ,  $\mathbf{R}_q^i$  is not ordered,  $\mathbf{M}_l$  and  $\mathbf{M}_r$  contain all the

boundary plaintexts of the range queries in  $\mathbf{Q}$ . In our construction, we design 3 steps to achieve the goal as following:

- First, the adversary must find a plaintext matching the range condition, whether its ciphertext is in the original EDB or not.
- Second, the adversary recovers the boundary plaintexts using Algorithm 1.
- Third, the adversary recovers all the plaintexts of the ciphertexts matching the range condition by several file-injections.

Here, to describe the FIA scheme briefly, Algorithm 2 is based on utilizing a single range query without any *order by* operation. In the following descriptions,  $q$  denotes the range query with the boundary ciphertexts denoted  $q.c_l$  and  $q.c_r$  respectively.  $\mathbf{M}_{\mathbf{R}_q}$  denotes the plaintext set corresponding to the cipher result set  $\mathbf{R}_q$  for query  $q$ .

In Algorithm 2, we adopt the breadth first search, because under the assumption of FIA the adversary does not know the order between file-injected data and the boundary ciphertexts in case the file-injected data do not match the range condition. With this limitation, the breadth first search is beneficial to find a plaintext matching the condition, and to get the relatively near unmatching plaintexts that are necessary for recovering the boundary plaintexts. Then, the boundary plaintexts  $m_l$  and  $m_r$  are recovered by calling Algorithm 1. Finally, the plaintext set  $\mathbf{M}_{\mathbf{R}_q}$  is recovered by several file-injections over the entire plaintext set matching the condition.

---

**Algorithm 2.**

$m_l, m_r, \mathbf{M}_{\mathbf{R}_q} \leftarrow \text{FIA\_Rangequery}(\mathcal{M}, q)$

---

- 1:  $a \leftarrow -1, b \leftarrow |\mathcal{M}|, \mathbf{d} \leftarrow (a, b)$
  - 2: insert  $\mathbf{d}$  into the queue  $\mathbf{dqueue}$
  - 3: **while**  $\mathbf{dqueue} \neq \emptyset$
  - 4:     take out the first  $\mathbf{d}$  in  $\mathbf{dqueue}, a \leftarrow \mathbf{d}.a, b \leftarrow \mathbf{d}.b$
  - 5:      $c \xleftarrow{\text{file injection}} \mathcal{M}[\text{mid}(a, b)]$
  - 6:     **if**  $|\mathbf{R}_q^0| \neq |\mathbf{R}_q|$  **break**
  - 7:     **if**  $(\text{mid}(a, \text{mid}(a, b)) \neq a$  **and**  $\text{mid}(a, \text{mid}(a, b)) \neq \text{mid}(a, b))$
  - 8:          $\mathbf{d} \leftarrow (a, \text{mid}(a, b)),$  insert  $\mathbf{d}$  into  $\mathbf{dqueue}$
  - 9:     **end if**
  - 10:    **if**  $(\text{mid}(\text{mid}(a, b), b) \neq b$  **and**  $\text{mid}(\text{mid}(a, b), b) \neq \text{mid}(a, b))$
  - 11:         $\mathbf{d} \leftarrow (\text{mid}(a, b), b),$  insert  $\mathbf{d}$  into  $\mathbf{dqueue}$
  - 12:    **end if**
  - 13:  $m_l \leftarrow \text{FIA\_Orderby}(\mathcal{M}_{a, \text{mid}(a, b)}, q.c_l, q)$
  - 14:  $m_r \leftarrow \text{FIA\_Orderby}(\mathcal{M}_{\text{mid}(a, b), b}, q.c_r, q)$
  - 15:  $\mathbf{M}_{\mathbf{R}_q} \leftarrow$  do file-injections from  $m_l$  to  $m_r$  and get their mapping table or corresponding plaintext set briefly
  - 16: **return**  $m_l, m_r, \mathbf{M}_{\mathbf{R}_q}$
- 

Most of the boundary values are very special in practice. For instance, the numbers, which are the multiple of  $10^\gamma$  ( $\gamma = 0, 1, 2, \dots$ ), are frequently used for

range query over numerical data; and the 26 letters are used for the same purpose over string data usually. Based on the different frequency of the plaintexts which are between every two adjacent common boundary plaintexts, the adversary may recover them more rapidly by several file-injections instead of the first step.

For space limitation, the analysis of time complexity, the discussions on FIA with both *order by* Queries and Range Queries, the description of our experiments and the FIA against Frequency-Hiding OPE are presented in the extended version [25] of this paper.

## 4 Formulating Forward Secure ORE

Forward security is a strong property of the dynamic SSE leakage profile. For a dynamic SSE scheme, its forward security means that: the previous data manipulations do not cause any leakage of the newly inserted data. Stefanov et al. [24] proposed this notion informally. Stefanov et al. [24] also proposed the concept of backward security, which ensures that the previous data manipulations do not leak any information about the newly deleted data. In this work, we extend this concept from SSE to OPE/ORE. Specifically, we give the definitions of forward security and backward security informally, as following:

**Definition 4 (Forward/Backward Security).** *An  $\mathcal{L}$ -adaptively-secure ORE scheme is forward (resp., backward) secure if the leakage profile, denoted  $\mathcal{L}_{\text{update}}$ , of update operation for  $\text{update} = \text{add}$  (resp.,  $\text{update} = \text{delete}$ ) can be described as following:*

$$\mathcal{L}_{\text{update}}(\text{update}, \mathbf{W}_{\text{update}}) = (\text{update}, \mathbf{IND}_{\text{update}})$$

where  $\text{add}$  (resp.,  $\text{delete}$ ) denotes the addition (resp., deletion) of data.  $\mathbf{W}_{\text{update}}$  is the data set of the update operations, in which the data have their own data storage structure, indices, and constraints according to the database.  $\mathbf{IND}_{\text{update}}$  is a set that only describes the modified column (in SQL database) or the document (in NoSQL database) and the indices of updated data.

Informally, a forward secure ORE ensures that the previous data order manipulations do not leak any information about the newly inserted data. Meanwhile, the new data order manipulations can be executed normally, and can correctly leak the order information about the newly inserted data. And in a forward secure ORE scheme,  $\mathbf{W}_{\text{update}}$  of a simple insertion can be briefly described as  $\mathbf{W}_{\text{add}} = (m, s)$ , where  $s$  denotes the order space of the related data on which order queries may be executed. For SQL databases,  $s$  can represent a column of a table. And for NoSQL databases,  $s$  can represent a set of documents.  $\mathbf{IND}_{\text{update}}$  of a simple insertion can be briefly described as  $\mathbf{IND}_{\text{add}} = (j, s)$ , where the incremental timestamp  $j$  is initially set to be 0 and is shared by all the manipulations.

Let  $e$  denote the intermediate ciphertext without forward security. Let  $\mathbf{op}(s)$  denote the order pattern of an order space  $s$ , which lists all the timestamps of

the order queries.  $\mathbf{Hist}(s)$  contains all the data-updating histories of  $s$  as well as the index  $index_s$  of  $s$ . Here, we only use it to list all the data-addition histories over the time. More formally, they can be defined as:

$$\mathbf{op}(s) = \{j : (j, s, \text{order}) \in \mathbf{List}_{\mathbf{SQL}}\}$$

$$\mathbf{Hist}(s) = \{index_s, (j, \text{add}, e) : (j, s, \text{add}, e) \in \mathbf{List}_{\mathbf{SQL}}\}$$

where  $\mathbf{add}$  denotes the addition manipulation,  $\mathbf{order}$  denotes the ordering manipulation,  $\mathbf{List}_{\mathbf{SQL}}$  denotes the list of data-manipulations. And we give the formal definition of forward secure ORE below.

**Definition 5 (Forward Secure ORE).** *Let the algorithm tuple*

$$\Gamma = (\text{ORE\_Setup}, \text{ORE\_Encrypt}, \text{ORE\_Compare})$$

be an ORE scheme. Let  $\mathcal{A}$  denote a PPT adaptive adversary. Define a real security game  $\mathbf{FS-ORE-R}_A^\Gamma(\lambda)$ , in which  $\mathcal{A}$  gets the public parameters output by  $\text{ORE\_Setup}(\lambda)$  and gets access to the encryption oracle and the comparison oracle adaptively. Based on the given public parameters and all the answers received from the oracles,  $\mathcal{A}$  outputs a bit as the result of the game. Define an ideal security game  $\mathbf{FS-ORE-I}_{A,S,\mathcal{L}_\Gamma}^\Gamma(\lambda)$ , in which a PPT simulator  $\mathcal{S}$  only takes the leakage profile  $\mathcal{L}_\Gamma$  as input.  $\mathcal{L}_\Gamma$  has two parts as following:

$$\mathcal{L}_{\text{update}}(\mathbf{add}, (m, s)) = (\mathbf{add}, (j, s))$$

$$\mathcal{L}_{\text{compare}}(c_1, c_2, s) = (\mathbf{op}(s), \mathbf{Hist}(s))$$

The simulator  $\mathcal{S}$  will output a bit as the result of the ideal game. The scheme  $\Gamma$  is said to be forward secure, if the following equation holds for any sufficient large  $\lambda$ :

$$|\mathbb{P}[\mathbf{FS-ORE-R}_A^\Gamma(\lambda) = 1] - \mathbb{P}[\mathbf{FS-ORE-I}_{A,S,\mathcal{L}_\Gamma}^\Gamma(\lambda) = 1]| \leq \text{negl}(\lambda)$$

where  $\text{negl}(\lambda)$  denotes a negligible function.

## 5 A Compilation Framework for Forward Secure ORE

To the best of our knowledge, all the existing OPE and ORE schemes in the literature do not have forward security *precisely*. Here, we use “*precisely*” with only the special case of POPE [23]. In [23], there is not any statement about whether the interactive processes need a client authorization or not. For the common application scenarios of OPE/ORE in practice, there is not any client authorization for querying. However, if the client authorization is mandated, POPE has forward security.

In the general case, the ciphertexts in EDB do not cover the entire ciphertext space. In other words, the ciphertexts in EDB are not dense usually. Thus, it is difficult to recover all the stored ciphertexts correctly with the limited leakage of

OPE/ORE. However, according to our FIA constructions and experiments, FIA schemes are powerful and effective in recovering data encrypted by OPE/ORE without forward security in practice. Though forward security can be achieved with oblivious RAM (ORAM) [11, 12] in general, it incurs massive overburden of performance [19] (large bandwidth consumption, multiple data round-trips, and/or large client storage complexity). Thus, it is desirable to have practical forward secure OPE/ORE schemes.

**Table 2.** Notations in Sect. 5

Notation	Meaning
$e$	Instance of intermediate ciphertext output by the original ORE or OPE scheme in EDB
$\Pi, \text{KeyGen}, sk, pk$	A TDP scheme, its key-generation algorithm and its secret key, public key
$\text{PRF}, H$	A pseudo-random function and a keyed hash function
$\text{OT}, \mathbf{OT}, \mathcal{OT}$	Instance of order token, map of order token stored on the client and domain of order token
$i$	The counter of order tokens, which is equal to the number of order tokens minus one
$s$	Instance of order space, which ensures that the data in different order spaces cannot be ordered

In this section, we present a practical compilation framework that transforms most of the existing OPE/ORE schemes into forward secure ones. To ease the understanding of the framework, we first give the meaning of some notations in Table 2.

## 5.1 Basic Ideas

With forward security, the *add* operation should leak nothing to server. In other words, the server should not distinguish between the ciphertexts output by a forward secure ORE and the ciphertexts encrypted by a perfect encryption scheme, when they are just inserted to the database before undergoing any search operation. In order to realize this goal, the ciphertext  $e$  generated by original OPE/ORE should be salted in our compilation framework. And we use TDP to link the salts to reduce the bandwidth consumption.

The salt is a hash value of an order token  $\text{OT}$  in our construction. To insert a new datum to EDB (say, the  $(i + 1)$ -th insertion,  $i \geq 0$ ), the client generates an order token  $\text{OT}_i$  based on the TDP scheme  $\Pi$ , its secret key  $sk$ , and the last order token  $\text{OT}_{i-1}$ . If  $\text{OT}_i$  ( $i = 0$ ) is the first order token in the order space, it will be randomly selected from the domain of order token  $\mathcal{OT}$ . In order to reduce the client storage, the client only stores the latest order token  $\text{OT}_i$  and the corresponding counter  $i$  in our basic construction. When an order query needs

to be executed, the client sends the current order token  $\text{OT}_i$  and the counter  $i$  to the server. The server can then calculate all the order tokens with the public key  $pk$ , and gets the original OPE/ORE ciphertexts by desalting operations. At last, the client will receive the correct comparison result which is calculated with the comparison algorithm of the original OPE/ORE by the server.

## 5.2 The Compilation Framework

Given any OPE or ORE scheme, denoted  $\Gamma = (\text{ORE\_Setup}, \text{ORE\_Encrypt}, \text{ORE\_Compare})$ , the compiled ORE scheme is described in Algorithm 3, which is denoted by  $\Gamma_{fp} = (\text{Setup}, \text{Encrypt}, \text{Compare})$ . In Algorithm 3, the parts of the original OPE/ORE are only briefly described.

---

### Algorithm 3. $\Gamma_{fp}$

---

#### Setup( $1^\lambda$ )

- 1:  $(pp, sp) \leftarrow \text{ORE\_Setup}(1^\lambda)$
  - 2:  $\mathbf{OT} \leftarrow$  empty map
  - 3:  $(sk, pk) \leftarrow \text{KeyGen}(1^\lambda)$
  - 4:  $k_0 \xleftarrow{\$} \{0, 1\}^\lambda$
  - 5: **return**  $((pk, pp), (sk, sp, k_0, \mathbf{OT}))$
- 

#### Encrypt( $pp, (sk, sp, k_0, \mathbf{OT}), m, (\text{add}, s, \sigma_1)$ )

*Client :*

- 1:  $k_s \leftarrow \text{PRF}_{k_0}[s]$
- 2:  $(\text{OT}_i, i) \leftarrow \mathbf{OT}[s]$
- 3: **if**  $(\text{OT}_i, i) = \perp$   $\{i \leftarrow -1, \text{OT}_{i+1} \xleftarrow{\$} \mathcal{OT}\}$
- 4: **else**  $\text{OT}_{i+1} \leftarrow \Pi_{sk}^{-1}(\text{OT}_i)$
- 5: **end if**
- 6:  $\mathbf{OT}[s] \leftarrow (\text{OT}_{i+1}, i + 1)$
- 7:  $c_{i+1} \leftarrow \text{ORE\_Encrypt}(pp, sp, m, (\text{add}, s, \sigma_1)) \oplus \text{H}(k_s, \text{OT}_{i+1})$
- 8: Send  $c_{i+1}$  to server.

*Server :*

- 9: Insert  $c_{i+1}$  into EDB.
- 

#### Compare( $(pk, pp), (sp, k_0, \mathbf{OT}), c_{s_\alpha}, c_{s_\beta}, (s, \sigma_2)$ )

*Client :*

- 1:  $k_s \leftarrow \text{PRF}_{k_0}[s]$
- 2:  $(\text{OT}_i, i) \leftarrow \mathbf{OT}[s]$
- 3: **if**  $(\text{OT}_i, i) = \perp$  **OR**  $i = 0$  **return**  $\emptyset$
- 4: Send  $(\text{OT}_i, i, k_s)$  to server.

*Server :*

- 5:  $e_{s_\alpha} \leftarrow c_{s_\alpha} \oplus \text{H}(k_s, \Pi_{pk}^{i-\alpha}(\text{OT}_i))$
  - 6:  $e_{s_\beta} \leftarrow c_{s_\beta} \oplus \text{H}(k_s, \Pi_{pk}^{i-\beta}(\text{OT}_i))$
  - 7:  $b \leftarrow \text{ORE\_Compare}(pp, sp, e_{s_\alpha}, e_{s_\beta}, (s, \sigma_2))$
  - 8: Send the result  $b$  to client
-

In our construction, we let  $\lambda$  denote the secure parameter. Let  $k_0$  denote the main key of our compilation framework. For each order space  $s$ , the key  $k_s$  of keyed hash function  $H$  is calculated by pseudo-random function  $\text{PRF}_{k_0}[s]$ . Let  $\text{add}$  denote the addition/insertion of data. The order tokens are calculated with TDP one by one in sequence, and the hash values of these tokens will xor the original OPE/ORE ciphertexts to generate the final ciphertexts without extra storage consumption at the server side. The salt of the final ciphertext is of  $\lambda$  bits, and will be desalted in the comparison algorithm. In the comparison algorithm, we let  $c_{s_\alpha}$  and  $c_{s_\beta}$  denote two ciphertexts to be compared in the order space  $s$  with their indices  $\alpha$  and  $\beta$  respectively. We let  $e_{s_\alpha}$  and  $e_{s_\beta}$  denote their intermediate ciphertexts output by the original OPE/ORE respectively.

For space limitation, the methods of data deletion and batch encryption is postponed to Appendix B.

### 5.3 Analysis of Forward Security

In our framework, the ciphertexts output by the original OPE/ORE xor the one-way generated salts. Hence, the newly inserted data leak nothing to the server if they have not been queried. Once the data have been queried and desalted, the ciphertexts turn into the security level of the original OPE/ORE scheme for the adversary with continuous monitoring. Hence, the security of the composite forward secure ORE cannot be weaker than that of the original OPE/ORE. On the other hand, our compilation framework is powerful against FIAs, because the forged data will not leak any information with the old queries. The data need a new credible order query from the client to desalt.

For space limitation, the formal proof of forward security is presented in the extended version [25] of this paper. Moreover, the description of our experiments is postponed to Appendix C.

## 6 Conclusion and Future Work

In this work, we study the leakage of OPE and ORE. We propose generic yet devastating FIA attacks which only exploit the *ideal* leakage of OPE/ORE. We also propose various improved methods to further boost the efficiency. Compared with existing attacks against OPE/ORE, our FIA attacks rely upon less demanding conditions, and can be more effective. We executed some experiments on real datasets to test the performance, and the results show that our FIA attacks can cause an extreme hazard on most of the existing OPE and ORE schemes with high efficiency and 100% recovery rate.

We then formulate forward-secure ORE, which may be of independent interest. In order to resist the disastrous effectiveness of FIA, we propose a practical compilation framework for transforming most existing OPE/ORE schemes into forward-secure ones. Finally, we execute experiments on some prominent OPE/ORE schemes developed in recent years, and the results show that our compilation framework is practical and useful for most of the systems.



Our compilation framework does not fit the OPE/ORE schemes which store the inserted data in order trees. Achieving forward security for these OPE/ORE schemes is an interesting direction for future research.

**Acknowledgement.** We thank the anonymous reviewers for their insightful comments. We are grateful to Yuan Li and Hongbing Wang for many helpful discussions.

## A Comparison Among FIA and Other Generic Attacks Against OPE/ORE

Known- and chosen-plaintext attacks (CPAs) have been considered in many OPE/ORE works. To the best of our knowledge, the latest discussion of CPA is in [13]. Besides these attacks against OPE/ORE, inference attack (IA) is also a kind of powerful generic attack which has been described in [20] detailedly. In this subsection, we make brief comparisons among FIA, CPA and IA.

About the adversarial prerequisite, these three attacks all need an unbroken auxiliary dataset as the plaintext space, but only IA needs the data-frequency statistics. About the source of leakage information, CPA only utilizes the data which are chosen by the adversary and encrypted by the encryption scheme in the system; IA only utilizes the original ciphertexts in the EDB; but FIA can utilize both the main forged data (which are chosen by the adversary and encrypted by the system) and the secondary original ciphertexts in the EDB. Specially, the leakage information in FIA is obtained through old queries. In other words, the ciphertexts, which are not included in the result sets of the old queries, are not required in the adversarial prerequisite. As for the comparison algorithm, CPA and IA must use it; FIA only calls it normally through the old queries. About the performance, 100% accuracy can be achieved easily with either CPA or FIA, but it is difficult to achieve with IA. Additionally, CPA and FIA can attack the frequency-hiding OPE, but IA cannot do this without the decline of accuracy and applicability.

Overall, FIA and CPA require less auxiliary information, while IA needs more auxiliary information. CPA needs more adversarial abilities, but FIA/IA need less adversarial abilities. Because of the utilization rate of leakage, FIA is more efficient than CPA, while the efficiency of IA depends on the size of plaintext space overly.

## B The Methods of Data Deletion and Batch Encryption

For data deletion, the first method is to store the deleted data in another EDB. Then a checking procedure should be added into the comparison algorithm to ensure that the ordered data have not been deleted. When the computing and bandwidth resource are sufficient and the server does not receive any query, the system can execute a **refresh** operation, which deletes all the deleted data from both EDB and recalculate all the order tokens and ciphertexts in sequence for curtailing storage and lifting efficiency. In this case, the scheme also achieves

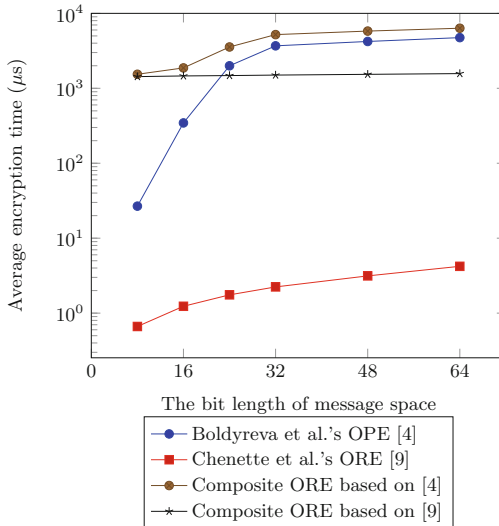
backward security. The second method is to insert the sequence numbers into EDB when inserting data. Then, for the situation where some data have been deleted, the server can calculate the salts of the remaining data by executing TDP exact times according to the interval leaked by sequence numbers. For space limitation, the analysis of storage and computational complexity is presented in the extended version [25] of this paper.

For batch encryptions, we can simply arrange all the elements in random order, and run the `Encrypt` algorithm in sequence. If batch encryptions are common in the system, we can add an extra batch index in the database for each datum, and use the same calculated order token for salting all the intermediate ciphertexts in a batch. This solution reduces the average computational complexity at the expense of leaking some information (e.g. equality) of the elements in the same batch.

## C Applicability and Experiments

Our compilation framework can be applied to all the OPE/ORE schemes except the OPE schemes (like the *ideal*-secure schemes proposed in [16, 21]) that store the ciphertexts in order trees. These OPE schemes [16, 21] leak all the ciphertext order from the tree structure. Hence, the salting of our framework is useless in this case.

In the experiments, we combined the OPE/ORE schemes in [4] and [9] with our forward secure framework. The experiments are implemented in C/C++,



**Fig. 2.** The comparison of average encryption time between two existing OPE/ORE schemes and the composite forward secure ORE schemes with our compilation framework.

and our experiments were performed using a single core on a machine with an Intel Pentium G2020 2.9 GHz CPU and 4 GB available RAM. We operate at 128-bits of security. We use HMAC as the PRF and the keyed hash function, and we use the RSA implementation (with 2048 bits RSA keys) in OpenSSL's BigNum library as the TDP. We use Blake2b as the underlying hash function. For our basic implementation of Boldyreva et al.'s OPE scheme, we use the C++ implementation from CryptDB [22], and for the implementation of Chenette et al.'s ORE scheme, we use the C-implemented FastORE mentioned in [18]. We use the California public employee payroll data from 2014 [1] as the experimental plaintext sets.

Figure 2 shows the comparison results of the average encryption time between the original schemes and the composite schemes. We respectively used 100000 data for testing each of the points in Fig. 2 and calculated the average results.

About the additional average encryption time, the composite forward secure ORE schemes demand about 1.5 ms for each datum encryption. Moreover, as to the additional average comparison time, the composite forward secure ORE schemes demand about 47  $\mu$ s. Because we chose two of the most practical OPE/ORE schemes as the contrasts, the composite forward secure ORE schemes seem slower. However, the additional comparison time is still at the microsecond level. Hence, our scheme is still practical and useful for the most common systems.

## References

1. California Public Employee Payroll Data (2014). <http://transparentcalifornia.com/downloads/>
2. Pmail (2014). <https://github.com/tonypr/Pmail>
3. Agrawal, R., Kiernan, J., Srikant, R., Xu, Y.: Order preserving encryption for numeric data. In: ACM SIGMOD International Conference on Management of Data, pp. 563–574. ACM, Paris (2004)
4. Boldyreva, A., Chenette, N., Lee, Y., O'Neill, A.: Order-preserving symmetric encryption. In: Joux, A. (ed.) EUROCRYPT 2009. LNCS, vol. 5479, pp. 224–241. Springer, Heidelberg (2009). [https://doi.org/10.1007/978-3-642-01001-9\\_13](https://doi.org/10.1007/978-3-642-01001-9_13)
5. Boldyreva, A., Chenette, N., O'Neill, A.: Order-preserving encryption revisited: improved security analysis and alternative solutions. In: Rogaway, P. (ed.) CRYPTO 2011. LNCS, vol. 6841, pp. 578–595. Springer, Heidelberg (2011). [https://doi.org/10.1007/978-3-642-22792-9\\_33](https://doi.org/10.1007/978-3-642-22792-9_33)
6. Boneh, D., Lewi, K., Raykova, M., Sahai, A., Zhandry, M., Zimmerman, J.: Semantically secure order-revealing encryption: multi-input functional encryption without obfuscation. In: Oswald, E., Fischlin, M. (eds.) EUROCRYPT 2015. LNCS, vol. 9057, pp. 563–594. Springer, Heidelberg (2015). [https://doi.org/10.1007/978-3-662-46803-6\\_19](https://doi.org/10.1007/978-3-662-46803-6_19)
7. Cash, D., Grubbs, P., Perry, J., Ristenpart, T.: Leakage-abuse attacks against searchable encryption. In: Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security, pp. 668–679. ACM, Denver (2015)
8. Cash, D., Liu, F.H., O'Neill, A., Zhang, C.: Reducing the leakage in practical order-revealing encryption. Cryptology ePrint Archive, Report 2016/661 (2016). <http://eprint.iacr.org/2016/661>

9. Chenette, N., Lewi, K., Weis, S.A., Wu, D.J.: Practical order-revealing encryption with limited leakage. In: Peyrin, T. (ed.) FSE 2016. LNCS, vol. 9783, pp. 474–493. Springer, Heidelberg (2016). [https://doi.org/10.1007/978-3-662-52993-5\\_24](https://doi.org/10.1007/978-3-662-52993-5_24)
10. Durak, F.B., DuBuisson, T.M., Cash, D.: What else is revealed by order-revealing encryption? In: Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security, pp. 1155–1166. ACM, Vienna (2016)
11. Garg, S., Mohassel, P., Papamanthou, C.: TWORAM: Efficient Oblivious RAM in Two Rounds with Applications to Searchable Encryption. In: Robshaw, M., Katz, J. (eds.) CRYPTO 2016. LNCS, vol. 9816, pp. 563–592. Springer, Heidelberg (2016). [https://doi.org/10.1007/978-3-662-53015-3\\_20](https://doi.org/10.1007/978-3-662-53015-3_20)
12. Goldreich, O., Ostrovsky, R.: Software protection and simulation on oblivious RAMs. *J. ACM (JACM)* **43**(3), 431–473 (1996)
13. Grubbs, P., Sekniqi, K., Bindschaedler, V., Naveed, M., Ristenpart, T.: Leakage-abuse attacks against order-revealing encryption. In: 2017 IEEE Symposium on Security and Privacy (SP), pp. 655–672. IEEE (2017)
14. He, W., Akhawe, D., Jain, S., Shi, E., Song, D.: Shadowcrypt: encrypted web applications for everyone. In: Proceedings of the 2014 ACM SIGSAC Conference on Computer and Communications Security, pp. 1028–1039. ACM, Scottsdale (2014)
15. Islam, M.S., Kuzu, M., Kantarcioglu, M.: Access pattern disclosure on searchable encryption: ramification, attack and mitigation. In: 19th Annual Network and Distributed System Security Symposium, vol. 20, p. 12. The Internet Society, San Diego (2012)
16. Kerschbaum, F.: Frequency-hiding order-preserving encryption. In: Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security, pp. 656–667. ACM, Denver (2015)
17. Lau, B., Chung, S.P., Song, C., Jang, Y., Lee, W., Boldyreva, A.: Mimesis aegis: a mimicry privacy shield—a system’s approach to data privacy on public cloud. In: Proceeding of the 23rd USENIX conference on Security Symposium, pp. 33–48. USENIX Association, San Diego (2014)
18. Lewi, K., Wu, D.J.: Order-revealing encryption: new constructions, applications, and lower bounds. In: Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security, pp. 1167–1178. ACM, Vienna (2016)
19. Naveed, M.: The fallacy of composition of oblivious ram and searchable encryption. Cryptology ePrint Archive, Report 2015/668 (2015). <http://eprint.iacr.org/2015/668>
20. Naveed, M., Kamara, S., Wright, C.V.: Inference attacks on property-preserving encrypted databases. In: Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security, pp. 644–655. ACM, Denver (2015)
21. Popa, R.A., Li, F.H., Zeldovich, N.: An ideal-security protocol for order-preserving encoding. In: 2013 IEEE Symposium on Security and Privacy (SP), pp. 463–477. IEEE, San Francisco (2013)
22. Popa, R.A., Redfield, C., Zeldovich, N., Balakrishnan, H.: Cryptdb: protecting confidentiality with encrypted query processing. In: Proceedings of the 23rd ACM Symposium on Operating Systems Principles, pp. 85–100. ACM, Cascais (2011)
23. Roche, D.S., Apon, D., Choi, S.G., Yerukhimovich, A.: Pope: partial order preserving encoding. In: Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security, pp. 1131–1142. ACM, Vienne (2016)
24. Stefanov, E., Papamanthou, C., Shi, E.: Practical dynamic searchable encryption with small leakage. In: 21st Annual Network and Distributed System Security Symposium, vol. 71, pp. 72–75. The Internet Society, San Diego (2014)

25. Wang, X., Zhao, Y.: Order-revealing encryption: file-injection attack and forward security. Cryptology ePrint Archive, Report 2017/1086 (2017). <http://eprint.iacr.org/2017/1086>
26. Zhang, Y., Katz, J., Papamanthou, C.: All your queries are belong to us: the power of file-injection attacks on searchable encryption. In: Proceeding of the 25th USENIX conference on Security Symposium, pp. 707–720. USENIX Association, Austin (2016)