Probability-p Order-Preserving Encryption

Ce Yang, Weiming Zhang(✉), Jiachen Ding, and Nenghai Yu

CAS Key Laboratory of Electro-magnetic Space Information,
University of Science and Technology of China, Hefei, China
{zhangwm,ynh}@ustc.edu.cn

Abstract. Order-Preserving Encryption (OPE) is an encryption preserving the order relationship of the plaintexts to support efficient range query on ciphertexts. Other than traditional symmetric encryption aiming at absolute security, OPE sacrifices some security for the ability to search on ciphertext. In this paper, we propose a new cryptographic primitive, Probability-p Order-Preserving Encryption (p-OPE), which preserves the order of plaintexts with probability p. When $p = 1$, p-OPE becomes OPE, thus p-OPE is an extension of OPE. We define and analyse the security and precision of the novel primitive, then we propose a construction of p-OPE and conduct experiments to show its performance. As shown in the theoretical analysis and experiment results, p-OPE can improve the security at the cost of some precision sacrifice.

Keywords: Searchable encryption · Order-Preserving Encryption · Range query

1 Introduction

Cloud computing is widely used nowadays. By outsourcing data to the cloud, customers can utilize the computing and storage resources provided by cloud server and reduce maintenance costs. With the rapid development of cloud computing, more and more sensitive information, such as customer information and transaction records, are stored in the cloud. Usually, cloud service can provide better security supportance than individual or small corporation. Nevertheless, the attackers may be willing to spend more resources to intrude the cloud server because of its potential interest. Sometimes, even the administrator of the cloud provider is in collusion with the attackers.

Therefore, sensitive data should be encrypted before outsourcing to the cloud. The mainstream approach used to protect data privacy is cryptography. Data are protected from unauthorized access after encryption. While protecting data privacy, encryption makes data hard to use. For example, traditional plaintext search methods fail on encrypted data. To solve this problem, searchable encryption systems [4,6–9,19–23] have been proposed. Searchable encryption enables searches on ciphertexts without decryption or decryption key, thus the plaintext of sensitive data is protected from attackers, even malicious cloud administrator.

X. Sun et al. (Eds.): ICCCS 2017, Part II, LNCS 10603, pp. 16–28, 2017.
https://doi.org/10.1007/978-3-319-68542-7_2

Order-Preserving Encryption (OPE) [1] is one of the cryptographic primitives enabling search on encrypted data. OPE is a secret-key encryption whose encryption function preserves numerical ordering of the plaintexts. OPE can be used to build an encrypted index supporting range query on numeric data, thus it has been used in encrypted database systems such as CryptDB [16] developed by MIT; Encrypted Bigquery client [10] developed by Google; and other secure retrieval systems [17].

OPE was first proposed by Agrawal et al. [1]. Boldyreva et al. [2] and Popa et al. [15] proved that any practical immutable OPE scheme leaks more than orders. Then, Boldyreva constructed an OPE scheme that uses hypergeometric distribution to lazy-sample a random order-preserving function. Popa et al. [15] designed an order-preserving encoding scheme that leaks at most the order through an interactive protocol. Boneh et al. [3] extended OPE to a more general concept as order-revealing encryption, of which the ciphertexts can be compared by an arbitrary algorithm other than the standard comparison operation as in the case of OPE, and they built a construction that leaks at most the order.

Though OPE provides efficient search ability for range query, it becomes vulnerable in some situation as shown in recent works.

Naveed et al. [14] considered the security of searchable encryption and presented four different attacks that recover the plaintext from property preserving encryption. Two of the attacks, sorting attack and cumulative attack, are applied to OPE and can recover plaintext with high probability. The sorting attack assumed that the plaintext was dense, which is not a typical situation for OPE. The cumulative attack utilized additional information about the distribution of plaintext.

Li et al. [11] developed a differential attack on OPE. Their attack reveals the leakage of distribution by exploiting the difference between ciphertexts. When OPE is used on the inverted index of an encrypted document dataset, experiments shows that the attacker can infer the encrypted keywords using differential attack if the attacker has some background information.

Durak et al. [5] studied the information leakage of ORE and OPE. They considered two issues: First, they showed that ORE may reveal additional information when multiple columns of correlated data are encrypted using OPE. Second, they discussed the leakage of concrete OPE schemes on non-uniform data.

These researches motivates us to propose a novel cryptographic primitive, which has better security than OPE while preserving the high efficiency of range query. Our contribution includes:

1. We propose a new cryptographic primitive, Probability-p Order-Preserving Encryption (p-OPE), which is an extension of OPE and aims at improving security. We define the security and precision metrics of p-OPE based on realistic situation, and we make a theoretical analysis of it.
2. We propose a construction of p-OPE and conduct experiments to show its performance. The experiment results show that the user can achieve a balance between security and query accuracy by adjusting the order-preserving probability p.

2 Probability-p Order-Preserving Encryption

In this section, we propose the concept of Probability-p Order-Preserving Encryption (p-OPE), which is an extension of OPE scheme. First, we define the novel cryptographic primitive. Then, we study the security and precision of p-OPE.

2.1 Definition

To improve the security of OPE, we propose the concept p-OPE.

Informally, a p-OPE $f : [1, n] \to [1, m]$ is an encryption scheme which preserves the order of the plaintext with probability not less than p when the plaintext follows a uniform distribution, i.e.

$$P(f(x_1) < f(x_2)|x_1 < x_2) \geq p. \tag{1}$$

When $p = 1$, p-OPE becomes OPE.

We give a definition which does not rely on the plaintext distribution. Considering an encryption $f : [1, n] \to [1, m]$, we say a pair of plaintext (x_1, x_2) is an ordered pair if $x_1 < x_2$, and an ordered pair (x_1, x_2) is a reverse pair if $f(x_1) > f(x_2)$. We define reverse number n_r as the number of reverse pairs, i.e. $n_r = |\{(x_1, x_2)|x_1 < x_2 \wedge f(x_1) > f(x_2)\}|$. An encryption f is a p-OPE, if the proportion of reverse pairs in all ordered pairs is smaller than $1 - p$, i.e. f is a p-OPE, if

$$n_r \leq (1 - p)\left(\frac{1}{2}n(n - 1)\right), \tag{2}$$

where n_r is the reverse number, n is the size of plaintext space. We can prove that the formal definition is in accordance with our intuitive definition:

Theorem 1. *If f is a p-OPE, x_1, x_2 are randomly picked from a uniform distribution on X, then*

$$P(f(x_1) < f(x_2)|x_1 < x_2) \geq p. \tag{3}$$

Proof

$$\begin{aligned} P(f(x_1) < f(x_1)|x_1 < x_2) &= \frac{P(f(x_1) < f(x_2) \wedge x_1 < x_2)}{P(x_1 < x_2)} \\ &= 1 - \frac{P(f(x_1) < f(x_2) \wedge x_1 > x_2)}{P(x_1 < x_2)} \\ &= 1 - \frac{|\{(x_1, x_2)|x_1 < x_2 \wedge f(x_1) > f(x_2)\}|}{|\{(x_1, x_2)|x_1 < x_2\}|} \\ &= 1 - \frac{|\{(x_1, x_2)|x_1 < x_2 \wedge f(x_1) > f(x_2)\}|}{\frac{1}{2}n(n - 1)} \\ &\geq 1 - (1 - p) \\ &= p. \end{aligned} \tag{4}$$

Thus the theorem is proved. □

To discuss the property of p-OPE, we start from a special p-OPE named as permutation function, which has a ciphertext space of the same size as the plaintext space. More precisely, a permutation function g is a bijection on integer interval $[1, n]$.

Every general p-OPE f is a composition of a permutation function g and an order-preserving function h. Order-preserving function h can be generated by mapping from i to the ciphertext of the i-th smallest ciphertext, and $g(x) = h_d(g(x))$, where h_d is the corresponding decryption function of the order-preserving function h.

Any p-OPE can do such a decomposition, and any permutation function and any OPE can be combined to build a p-OPE. Based on this, we can analyse the security and precision of p-OPE.

2.2 Security

Researchers have proposed different security metrics for OPE. Here we use *mean absolute error* (MAE) [13] to measure the security. MAE applies to scenarios that not only the accurate recovery but also a close estimation of plaintext is acceptable, which holds for most application of OPE. For example, if the encrypted data is salary, the adversary usually does not care the difference between 10,000 and 10,100.

We define MAE formally here. For a p-OPE $f : [1, n] \to [1, m]$ with decryption function f_d, the MAE of an adversary is defined as:

$$d_{\text{MAE}}(f_d, f'_d) = \sum_y |f_d(y) - f'_d(y)| P_y, \tag{5}$$

where $f_d(y)$ is the decryption function, $f'_d(y)$ is the approximation of the adversary, and P_y is the probability of ciphertext y. When the order-preserving function is a continuous function, the definition will be an integral. When the plaintext follows a uniform distribution, MAE is equivalent to $d_m = \sum_x |x - f'_d(f(x))|$.

We consider two different attack scenarios.

The first one is the scenario in which the adversary knows the plaintext distribution. Previous work [14] shows that an adversary can recover the plaintext of an OPE with high precision in this scenario. As discussed in previous subsection, f is a composition of a permutation function g and an order-preserving function h. Here we consider the situation that the order-preserving function is insecure, i.e. the adversary knows $h(x)$, thus the security of p-OPE relies on the permutation function.

When the adversary knows $h(x)$ and has no knowledge of $g(x)$, an option for the adversary is to estimate plaintext as: $f'_d(y) = h_d(y)$, where $h_d(y)$ is the decryption function of OPE $h(x)$. We have $f'_d(y) = f'_d(h(g(x))) = h_d(h(g(x))) = g(x)$, i.e. $f_d(y) - f'_d(y) = x - g(x)$. Thus, the security of p-OPE is determined by the permutation function $g(x)$ in our assumption.

Now we analyse the relationship between order-preserving probability p and MAE d_m. We have:

Theorem 2. *For a permutation function $g : [1, n] \to [1, n]$ with reverse number n_r, if the adversary uses the identity function $g'_d(k) = k$ as the estimation and the plaintext follows a uniform distribution, then MAE d_m satisfies*

$$n_r \leq d_m \leq 2n_r. \tag{6}$$

Proof. The reverse pairs can be split to four different parts as:

$$\begin{aligned} &f(x_1) > x_1, f(x_2) > x_2, \\ &f(x_1) > x_1, f(x_2) = x_2, \\ &f(x_1) > x_1, f(x_2) < x_2, \\ &f(x_1) \leq x_1, f(x_2) < x_2. \end{aligned} \tag{7}$$

We abbreviate subscription $1, 2$ and denote them with $n(>, >), n(>, =), n(>, <)$, $n(\leq, <)$.

Consider x_i such that $f(x_i) > x_i$ and $f(x_i) - x_i$ gets maximum. For $x < x_i$, we have $f(x) - x \leq f(x_i) - x_i$, thus

$$f(x) \leq f(x_i) + x - x_i < f(x_i), \tag{8}$$

we know that x, x_i cannot be a reverse pair for $x < x_i$.

For $x > x_i$, (x_i, x) is a reverse pair if $f(x) < f(x_i)$. Because the number of x satisfying $f(x) < f(x_i)$ is $f(x_i) - 1$, and the number of x satisfying $x < x_i$ and $x_i - 1$ is $x_i - 1$. Thus the number of x with $x_i < x$ and $f(x_i) > f(x)$ is $f(x_i) - x_i$, i.e.,

$$|\{x | x > x_i \wedge f(x) < f(x_i)\}| = f(x_i) - x_i. \tag{9}$$

and we have the number of reverse pairs related with x_i is $f(x_i) - x_i$.

We can remove x_i from points of interested, and repeat it. Thus for all x with $f(x) > x$, the sum of number of reverse pairs related with x is

$$N(f(x) > x) = \sum_{f(x) > x} f(x) - x. \tag{10}$$

i.e.

$$n(>, >) + n(>, =) + n(>, <) = \sum_{f(x) > x} f(x) - x. \tag{11}$$

For x satisfying $f(x) < x$, we have a similar conclusion,

$$N(f(x) < x) = \sum_{f(x) < x} x - f(x). \tag{12}$$

i.e.

$$n(>,<) + n(\leq,<) = \sum_{f(x)<x} x - f(x). \tag{13}$$

Thus

$$\begin{aligned} n_r &= n(>,>) + n(>,=) + n(>,<) + n(\leq,<) \\ &\leq n(>,>) + n(>,=) + n(>,<) + n(>,<) + n(\leq,<) \\ &= \sum_x |f(x) - x| \\ &= d_m, \end{aligned} \tag{14}$$

and similarly,

$$2n_r \geq d_m. \tag{15}$$

In conclusion, the theorem is proved. □

Because order-preserving probability p is determined by n_r, this theorem shows the relationship between p and security. The security increases with the reverse number, and $d = 0$ when $p = 1$.

The second one is the adversary knows several plaintext-ciphertext pairs. In this scenario, the adversary can use interpolation to estimate the decryption function. We show the performance of p-OPE and compare it with OPE constructions using experiments in the latter.

2.3 Precision

Precision is related with applications. We consider the precision of range query. Here we adopt the false positive and false negative from the evaluation of plaintext information retrieval [12] to illustrate the precision of ciphertext query. For a plaintext query Q_X and the corresponding ciphertext query Q_Y, a false negative is a plaintext x appear in the plaintext query but the corresponding ciphertext not in the ciphertext query, i.e. $x \in Q_X$ and $f(x) \notin Q_Y$; a false positive is the opposite, i.e. $x \notin Q_X$ and $f(x) \in Q_Y$.

We measure the precision using the ratio of right results and errors. For a range query, if the query result contains n_e related results and n_{fp} false positives, and the number of false negatives is n_{fn}, then we define the precision as:

$$P_e = \frac{n_e}{n_e + n_{fp} + n_{fn}}. \tag{16}$$

The precision of a p-OPE is the average of precision of all possible range queries.

In term of range query, different query distribution will lead to different precision. A simple way to generate query is randomly pick two plaintext uniformly and make them the endpoint of query interval. However, interval generated such way is somewhat not uniform. Assume the plaintext chosen is x_1, x_2, then the probability of plaintext x in the interval $[x_1, x_2]$ or $[x_2, x_1]$ is

$$\begin{aligned} P &= P(x_1 \le x \wedge x_2 \ge x) + P(x_1 > x \wedge x_2 \le x) \\ &= P(x_1 \le x)P(x_2 \ge x) + P(x_1 > x)P(x_2 \le x) \\ &= \frac{x}{N}\frac{N-x+1}{N} + \frac{N-x}{N}\frac{x}{N}. \end{aligned} \tag{17}$$

Obviously, different plaintext will be queried with different probability, and the midpoint of plaintext space will be queried in half of the queries. Because we assume the plaintext follows a uniform distribution, we also hope that each plaintext will be queried equiprobablely. We generate intervals with the same length and different centers as queries. The queries is generated as $[x, x+k]$, where x obeys a uniform distribution, and k is a fixed number.

For each plaintext interval $[x_1, x_2]$, the user can traverse different ciphertext intervals to find the most suitable ciphertext query $[y_1', y_2']$ to minimize the error. Here we propose a more simple way to response the query. For plaintext interval $[x_1, x_2]$, we use $[x_1 - e, x_2 + e]$ as ciphertext query, where e is a constant chosen by user.

Now we analyse the precision of our ciphertext query. We have

Theorem 3. *Assume p-OPE $g(x) : [1, n] \to [1, n]$ is a permutation function. If range query is generated as $[x, x+k]$, where x follows a uniform distribution on $[1-k, n]$, then the precision of g satisfies*

$$P_e \ge \frac{k+1}{(n+k)(k+2e+1)}\left(n - \left\lceil \frac{d_m}{k+e+1} \right\rceil\right). \tag{18}$$

Proof. We say a plaintext x matches an plaintext interval query $[x_1, x_2]$, if the plaintext x is in the plaintext interval $[x_1, x_2]$, and the corresponding ciphertext $g(x)$ is in the corresponding ciphertext query $[x_1 - e, x_2 + e]$.

Denote the set of plaintexts matching a interval $[x_1, x_2]$ as $Q_e([x_1, x_2])$, i.e. if $x \in [x_1, x_2]$ and $g(x) \in [x_1 - e, x_2 + e]$, then $x \in Q_e([x_1, x_2])$, and vice versa.

For each plaintext x, we can calculate $n_q(x)$, the number of intervals it matches. The definition of $n_q(x)$ is

$$n_q(x) = |\{(x_1, x_2) | x_2 - x_1 = s \wedge x \in Q_e([x_1, x_2])\}|. \tag{19}$$

For plaintext x with $|x - f(x)| \le e$, x matches each interval $[x_1, x_2]$ it in. For plaintext x with $e < |x - f(x)| \le k + e$, x matches $s + 1 + d - |x - f(x)|$ different intervals. For plaintext x with $k + e < |x - f(x)|$, x matches none interval. Thus, if we define function $f_r(k)$ as

$$f_r(i) = \begin{cases} k+1, & i \le e \\ k+1+e-i, & e \le i \le e+k+1 \\ 0, & e+k+1 \le k \end{cases} \tag{20}$$

then $n_q(x) = f_r(|x - f(x)|)$.

According to the definition of precision, the precision of g is

$$\begin{aligned} P_e &= \sum_{x=1-k}^{n} P_e(x)\frac{1}{n+k} \\ &= \frac{1}{n+k}\sum_{i=1-k}^{n}\frac{n_e}{k+2e+1} \\ &= \frac{1}{n+k}\sum_{i=1-k}^{n}\frac{|\{x|x\in Q_e([i,i+k])\}|}{k+2e+1} \\ &= \frac{1}{(n+k)(k+2e+1)}\sum_{x=1}^{n} n_q(x). \end{aligned} \tag{21}$$

Now we discuss the relationship between $\sum n_q$ and MAE d_m. When the MAE of two sequences is a fixed num, we estimate the precision between them.

For two numbers k_1, k_2 with fixed sum, we calculate $f_r(k_1) + f_r(k_2)$. It is easy to know, if $k_1 + k_2 > k + e + 1$, then $f_r(k_1) + f_r(k_2)$ get minimum when $k_1 = k + e + 1$; if $k_1 + k_2 < k + e + 1$, then $f_r(k_1) + f_r(k_2)$ get minimum when $k_1 = 0$.

Thus for n different numbers, the sum of their precision function will greater than the sequence of $(k + e + 1, k + e + 1, \ldots, k + e + 1, s_r, 0, \ldots, 0)$, where $0 \le s_r \le k + e + 1$. If $(k + e + 1)(n_l - 1) < d_m \le (k + e + 1)n_l$, we have $\sum n_r(x) > (n - n_l)(k + 1)$, i.e.

$$\sum n_r(x) \ge \left(n - \left\lceil \frac{d_m}{k+e+1} \right\rceil\right)(k+1). \tag{22}$$

Thus, we have

$$\begin{aligned} P_e &= \frac{\sum n_q(x)}{(n+k)(k+2e+1)} \\ &\ge \frac{k+1}{(n+k)(k+2e+1)}\left(n - \left\lceil \frac{d_m}{k+e+1} \right\rceil\right). \end{aligned} \tag{23}$$

□

Theorem 3 shows the relationship between precision and security. The lower bound ensures the precision in the worst case. When the size of plaintext space is large enough, the probability of precision reaching lower bound is small. Thus, in most situation, actual precision is better than the lower bound.

3 Construction

In this section we give a construction of p-OPE based on the framework in previous section. First, we give the algorithm of our construction, then we show its actual security and precision performance.

3.1 Construction

As discussed in previous section, a p-OPE can be constructed from a permutation function and an order-preserving function. The OPE scheme has been discussed by many previous researchers, and ready-made schemes such as random-partition scheme or OPE based on hypergeometric distribution can be adopted.

Thus, we focus on the permutation scheme. If the size of plaintext space is n, the targeted order-preserving probability is p, then we use a random-shift algorithm to generate a permutation function.

The algorithm runs in a loop. Based on the size of p, the algorithm execute multiple rounds of interval-shift operations to adjust the order relationship of the sequence. A right interval-shift on an sequence $(x_1, x_2, \ldots, x_n)$ is to move the first element to the last position, and move other elements to the left position next to it, i.e. $(x_1, x_2, \ldots, x_n) \rightarrow (x_2, x_3, \ldots, x_n, x_1)$. A left interval-shift is a similar operation in the opposite direction, i.e. $(x_1, x_2, \ldots, x_n) \rightarrow (x_n, x_1, x_2, \ldots, x_{n-1})$. In each round of the algorithm, the algorithm randomly choose a continuous subsequence of x^n, and evaluate the reverse number after interval-shift. If the reverse number decreases with the interval-shift, the operation will not be executed but discarded. The same operation is repeated until the actual order-preserving probability is close to p enough.

The detailed algorithm is shown in Algorithm 1.

3.2 Experiments

In this subsection, we conduct experiments to show the performance of our constructions. First, we verify the theoretical analysis.

The security and precision of proposed construction is shown in Fig. 1. As we expect, MAE d_m increases with the increase of reverse number n_r, and precision decreases with the decrease of order-preserving probability. Figure 1 also shows that a linear function of n_r is a close enough approximation of d_m when the size of plaintext space n is large enough.

Second, we compare the security of p-OPE with three different OPE constructions, random order-preserving function proposed by Boldyreva [2], random uniform sampling proposed by Wozniak [18], and order-revealing encryption proposed by Boneh [3]. We consider the scenario that the attacker knows k plaintext-ciphertext pairs and he uses linear interpolation to estimate the plaintext. The results of proposed method, Boldyreva, and Wozniak are shown in Fig. 2. A point at rate r and MAE d means the number of ciphertext which has a MAE smaller than d is nr when the experiment repeats n times. The proposed method improves the security significantly. With the increase of k, the security of OPE decreases.

Because the ciphertext of ORE is not number, we map the ciphertext of ORE to integers according to the orders before we use interpolation. For example, when the ciphertext sequence is (y_1, y_2, y_3) and $y_2 < y_3 < y_1$, we map the ciphertexts to $(3, 1, 2)$. We compare proposed method of different order-preserving probability with Boneh. The experiment result is shown in Fig. 3. As

Algorithm 1. Random-Shift algorithm

Input the size of plaintext space n, expected order-preserving probability p.
Output permutated sequence x^n.
$n_r = \lfloor \frac{1}{2} n(n-1) * p \rfloor$.
Initialize $x^n = (1, 2, \ldots, n)$.
while $n_r > 3$ **do**
  Choose l from a uniform distribution on $[1, n]$.
  Choose k from a uniform distribution on $[1, n_r]$.
  Randomly choose d from $\{-1, 1\}$.
  if $d < 0$ **then**
    flip the sequence x^n
  end if
  if $x + k > n$ **then**
    $k = n - x$.
  end if
  $s = 0$
  for i in $[l + 1, l + k]$ **do**
    if $x_l < x_i$ **then**
      $s = s + d$
    else
      $s = s - d$
    end if
  end for
  if $s > 0$ **then**
    $x = x_l$
    for i in $[l + 1, l + k]$ **do**
      $x_{i-1} = x_i$
    end for
    $x_{l+k} = x$
  end if
  $n_r = n_r - s$
  if $d < 0$ **then**
    flip the sequence x^n
  end if
end while

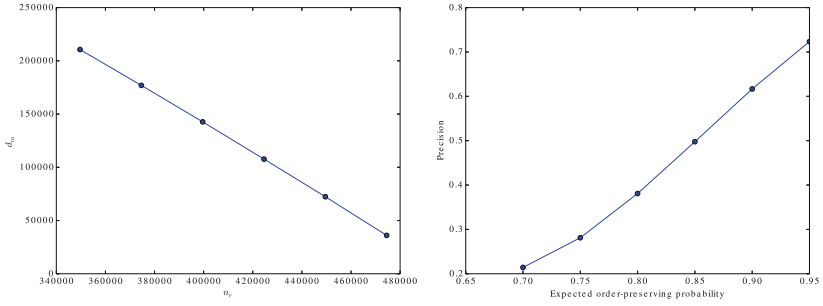

Fig. 1. The security and precision of p-OPE. The size of plaintext space is 1000. Left figure shows the relationship between reverse number n_r and MAE d_m, and right figure shows the relationship between order-preserving probability and query precision.

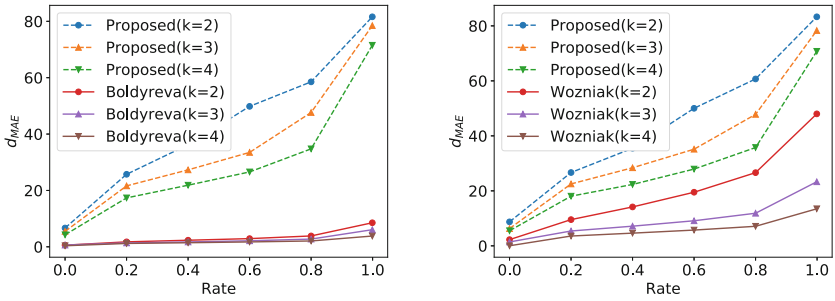

Fig. 2. The security of proposed method (p = 0.80) and Boldyreva, Wozniak.

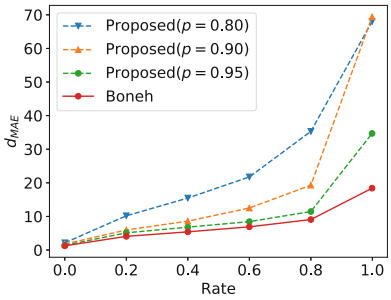

Fig. 3. The security of proposed method and Boneh.

we expected, the attack accuracy of proposed method is lower than Boneh, which leaks only orders, and the security improves with the decrease of order-preserving probability.

4 Conclusion

In this paper, we propose a new cryptographic primitive p-OPE, aiming at improving the security of OPE while preserving searching efficiency. We analyse the security and precision of p-OPE, which is an extension of OPE. We also propose a construction of p-OPE and conduct experiment to verify its performance. The experiment results are in accordance with the theoretical analysis.

Acknowledgements. This work was supported in part by the Natural Science Foundation of China under Grant U1636201 and 61572452.

References

1. Agrawal, R., Kiernan, J., Srikant, R., Xu, Y.: Order preserving encryption for numeric data. In: Proceedings of the 2004 ACM SIGMOD International Conference on Management of Data, SIGMOD 2004, NY, USA, pp. 563–574. ACM, New York (2004)

2. 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). doi:10.1007/978-3-642-01001-9_13
3. 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). doi:10.1007/978-3-662-46803-6_19
4. Chen, C., Zhu, X., Shen, P., Hu, J., Guo, S., Tari, Z., Zomaya, A.Y.: An efficient privacy-preserving ranked keyword search method. IEEE Trans. Parallel Distrib. Syst. **27**(4), 951–963 (2016)
5. 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, CCS 2016, NY, USA, pp. 1155–1166 (2016). http://doi.acm.org/10.1145/2976749.2978379
6. Fu, Z., Huang, F., Sun, X., Vasilakos, A., Yang, C.N.: Enabling semantic search based on conceptual graphs over encrypted outsourced data. IEEE Trans. Serv. Comput. **PP**(99), 1 (2016)
7. Fu, Z., Ren, K., Shu, J., Sun, X., Huang, F.: Enabling personalized search over encrypted outsourced data with efficiency improvement. IEEE Trans. Parallel Distrib. Syst. **PP**(99), 1 (2015)
8. Fu, Z., Sun, X., Ji, S., Xie, G.: Towards efficient content-aware search over encrypted outsourced data in cloud. In: IEEE INFOCOM 2016 - The 35th Annual IEEE International Conference on Computer Communications, pp. 1–9, April 2016
9. Fu, Z., Wu, X., Guan, C., Sun, X., Ren, K.: Toward efficient multi-keyword fuzzy search over encrypted outsourced data with accuracy improvement. IEEE Trans. Inf. Forensics Secur. **11**(12), 2706–2716 (2016)
10. Google: The encrypted bigquery client. https://github.com/google/encrypted-bigquery-client
11. Li, K., Zhang, W., Yang, C., Yu, N.: Security analysis on one-to-many order preserving encryption-based cloud data search. IEEE Trans. Inf. Forensics Secur. **10**(9), 1918–1926 (2015)
12. Manning, C.D., Raghavan, P., Schütze, H., et al.: Introduction to Information Retrieval, vol. 1. Cambridge University Press, Cambridge (2008)
13. Martınez, S., Miret, J.M., Tomas, R., Valls, M.: Security analysis of order preserving symmetric cryptography. Appl. Math. Inf. Sci. (AMIS) **7**(4), 1285–1295 (2013)
14. 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, CCS 2015, NY, USA, pp. 644–655. ACM, New York (2015)
15. Popa, R., Li, F., Zeldovich, N.: An ideal-security protocol for order-preserving encoding. In: 2013 IEEE Symposium on Security and Privacy (SP), pp. 463–477, May 2013
16. Popa, R.A., Redfield, C.M.S., Zeldovich, N., Balakrishnan, H.: Cryptdb: protecting confidentiality with encrypted query processing. In: Proceedings of the Twenty-Third ACM Symposium on Operating Systems Principles, SOSP 2011, NY, USA, pp. 85–100. ACM, New York (2011)
17. Wang, C., Cao, N., Ren, K., Lou, W.: Enabling secure and efficient ranked keyword search over outsourced cloud data. IEEE Trans. Parallel Distrib. Syst. **23**(8), 1467–1479 (2012)

18. Wozniak, S., Rossberg, M., Grau, S., Alshawish, A., Schaefer, G.: Beyond the ideal object: towards disclosure-resilient order-preserving encryption schemes. In: Proceedings of the 2013 ACM Workshop on Cloud Computing Security Workshop, CCSW 2013, NY, USA, pp. 89–100. ACM, New York (2013)
19. Xia, Z., Wang, X., Sun, X., Wang, Q.: A secure and dynamic multi-keyword ranked search scheme over encrypted cloud data. IEEE Trans. Parallel Distrib. Syst. **27**(2), 340–352 (2016)
20. Xia, Z., Wang, X., Zhang, L., Qin, Z., Sun, X., Ren, K.: A privacy-preserving and copy-deterrence content-based image retrieval scheme in cloud computing. IEEE Trans. Inf. Forensics Secur. **11**(11), 2594–2608 (2016)
21. Xia, Z., Zhu, Y., Sun, X., Qin, Z., Ren, K.: Towards privacy-preserving content-based image retrieval in cloud computing. IEEE Trans. Cloud Comput. **PP**(99), 1 (2015)
22. Xia, Z., Xiong, N.N., Vasilakos, A.V., Sun, X.: EPCBIR: an efficient and privacy-preserving content-based image retrieval scheme in cloud computing. Inf. Sci. **387**, 195–204 (2017). http://www.sciencedirect.com/science/article/pii/S0020025516321971
23. Zhangjie, F., Xingming, S., Qi, L., Lu, Z., Jiangang, S.: Achieving efficient cloud search services: multi-keyword ranked search over encrypted cloud data supporting parallel computing. IEICE Trans. Commun. **98**(1), 190–200 (2015)

http://www.springer.com/978-3-319-68541-0

Cloud Computing and Security
Third International Conference, ICCCS 2017, Nanjing, China, June 16-18, 2017, Revised Selected Papers, Part II
Sun, X.; Chao, H.-C.; You, X.; Bertino, E. (Eds.)
2017, XXII, 879 p. 281 illus., Softcover
ISBN: 978-3-319-68541-0