Beyond the Ideal Object: Towards Disclosure-Resilient Order-Preserving Encryption Schemes

Sander Wozniak Michael Rossberg Sascha Grau Ali Alshawish Guenter Schaefer

Technische Universität Ilmenau, Germany

firstname.lastname@tu-ilmenau.de

ABSTRACT

With the emergence of affordable cloud services, users are currently moving data to external services providers. Hence, they implicitly trust providers to not abuse or "lose" sensitive data. To protect this data in the context of cloud computing, the use of Order-Preserving Encryption (OPE) has been suggested to encrypt data while still allowing efficient queries. The reference approach builds on Order-Preserving Functions (OPFs) drawn uniformly at random: the so-called "ideal object". However, recent results question the suitability of this construction, as its security properties turn out to be poor. In this article, we investigate possible alternatives. For this, we introduce two descriptive metrics rating one-wayness-related properties of OPF construction schemes, i.e., the ability of an adversary to estimate the plaintext when given a ciphertext and possible extra information. Furthermore, we propose three novel approaches to draw OPFs and apply the introduced metrics to study their security features in relation to the "ideal object". The results visualize the extent of insecurity caused by using the "ideal object" and qualify the suitability of the alternative schemes under different threat scenarios.

Categories and Subject Descriptors

E.3 [**Data Encryption**]; H.2.7 [**Database Management**]: Database Administration —*Security, integrity, and protection*

General Terms

Design, Security, Algorithms

Keywords

order-preserving encryption; ideal object; one-wayness; disclosure-resilience

CCSW'13, November 08, 2013, Berlin, Germany.

http://dx.doi.org/10.1145/2517488.2517496.

1. INTRODUCTION

While traditional cryptographic operations aim at rendering structured plaintexts into entirely unstructured ciphertexts (and potentially vice versa), recent cryptographic research also considers operations where ciphertexts still reflect some of the properties of the plaintexts. The topic has especially been incited by the emergence of cloud computing services, which provide flexible and cost-efficient access to infrastructure, platforms, and applications. Nevertheless, the outsourced data must be protected from unauthorized access. However, encrypting the data with established algorithms like AES would prevent processing the data, i.e., range querying. Such techniques require sorted index information, that is the order of plaintexts must be kept among the respective ciphertexts, a prerequisite that cannot be achieved with traditional encryption schemes.

In order to address this problem, the use of Order-Preserving Encryption (OPE) has been suggested [2, 5, 9, 1]. A deterministic symmetric OPE scheme is modeled as triple $\mathcal{S} = (\mathcal{K}, \mathcal{E}nc, \mathcal{D}ec)$ [3]. Using a key K obtained from randomized key generation algorithm $\mathcal{K}$, the encryption algorithm $\mathcal{E}nc(K, p) = f(p)$ and its inverse $\mathcal{D}ec(K, c) = f^{-1}(c)$ realize an Order-Preserving Function (OPF) that maps plaintexts from a domain $\mathcal{D}$ to a range $\mathcal{R}$ of possible ciphertexts.

While there is general consent that OPE will never achieve the same level of security as traditional encryption schemes [1, 3, 12, 10], many possible security features are still not fully understood. Existing efforts in the analysis of the security of OPE have primarily been focused on the so-called *"ideal object"* [3, 4, 11, 12]. The "ideal object" is a Random Order-Preserving Function (ROPF), i.e., a strictly monotonically increasing function $f\colon \mathcal{D} \to \mathcal{R}$ that is chosen uniformly at random from the sample space $\mathsf{OPF}_{\mathcal{D},\mathcal{R}} = \{f\colon \mathcal{D} \to \mathcal{R} \mid \forall p_1, p_2 \in \mathcal{D}\colon f(p_1) < f(p_2) \Leftrightarrow p_1 < p_2\}$. It is referred to as "ideal" since it produces every possible OPF with the same probability.

As already indicated, OPE schemes may be used in cloud environments where users are unwilling to fully trust their service provider. While the resistance against known ciphertext attacks is an important property in this scenario, also other aspects must be kept in mind. For example, when considering location-based services, users should be able to compare their whereabouts, e.g, using a cloud service. However, for privacy reasons, it should not be possible to find their accurate positions. Here, in contrast to the commonly referred database scenario, three fundamental differences exist: the attacker may only be able to obtain a limited number of samples of the entire ciphertext space, malicious users

may be able to obtain some plaintext-ciphertext pairs, and might even be aware of the full domain.

Keeping in mind the diverse requirements on OPE, this article investigates randomized construction schemes of order-preserving functions that improve on one-wayness-related properties of the "ideal object". In particular, when considering the functions produced by a specific scheme and an entry $c \in \mathcal{R}$ of the ciphertext range, there is a distribution specifying the probability that entries from the plaintext domain $\mathcal{D}$ are mapped to c. For the "ideal object", this relates to the hypergeometric distribution [3], featuring a dominant peak at the most likely plaintext (m.l.p.), and disqualifying most entries of the domain as possible plaintexts. This enables accurate guesses of the plaintext mapped to a given ciphertext [4]. In this article, we aim for schemes that minimize the importance of the m.l.p., increasing the number of plaintexts that are significant candidates for producing a ciphertext c. Furthermore, this property should even be maintained despite the disclosure of a limited number of ciphertexts or plaintext-ciphertext pairs, as well as in the presence of chosen plaintext attacks. We refer to schemes fulfilling these requirements as *disclosure-resilient*.

Since the "ideal object" is a reference scheme to evaluate the security of OPE approaches, it does not demand that the created functions have a compact representation. In particular, they can be thought of as being stored using $|\mathcal{D}|$ key-value-pairs. In our pursuit of more secure reference schemes, we will follow the same approach.

Within this article, we make the following contributions:

- We propose three novel schemes to draw OPFs from $\mathsf{OPF}_{\mathcal{D},\mathcal{R}}$, that aim to improve the disclosure-resilience of OPF-based OPE schemes.
- In order to evaluate the disclosure-resilience of OPE schemes, we propose the security metrics of the *(average) number of significant plaintexts* and the *expected estimation error* that a maximum-likelihood attacker can achieve given a specific number of ciphertexts or plaintext-ciphertext pairs.
- Finally, we provide an empirical evaluation of the disclosure-resilience of our schemes, comparing them to the "ideal object". From our evaluation, we can see that the proposed approaches significantly improve on the "ideal object" in case of ciphertext-only attacks.

The rest of the article is organized as follows: In Section 2, we provide an overview of existing research on OPE. Section 3 describes our attacker model, as well as our three novel OPF construction schemes. The metrics for evaluating the disclosure-resilience of an OPE scheme are introduced in Section 4. After that, the results of our empirical security evaluation are discussed in Section 5. Finally, in Section 6, we conclude the article with an outlook.

2. RELATED WORK

We now provide an overview of related work. For this, we start with an outline of early ad-hoc realizations and highlight initial efforts in the security analysis of OPE. Then, we review known weaknesses of the "ideal object" and discuss proposed alternatives to OPF-based OPE.

2.1 Early Realizations and the Ideal Object

Several approaches have been suggested to realize OPE. Some of the initially suggested OPF schemes rely on traditional encryption techniques to generate intermediate ciphertexts, assign them to buckets based on the order of their underlying plaintexts, and prepend bucket identifiers to obtain resulting ciphertexts [4, 5, 6]. Thus, with the bucket identifiers, ciphertexts can be compared, allowing efficient range queries. Varying the bucket size, these schemes provide trade-offs between the potential risk of plaintext exposure and the query performance, which may suffer from a large number of false positives in the results of encrypted queries. Nevertheless, while these approaches provide viable solutions for some outsourced database applications, they cannot be securely employed if the domain is known to adversaries. With attackers being aware of the domain, bucket identifiers can already reveal the underlying plaintexts of the given ciphertexts (depending on the bucket size). Therefore, instead of grouping the underlying plaintexts, OPE schemes have to obfuscate the mapping between plain- and ciphertexts in most use cases – for both security and performance.

Apart from bucketing techniques, the summation of random numbers [2], as well as the use of polynomial functions [9] have been suggested for integer domains and ranges.

While these approaches produce OPFs from specific subclasses of $\mathsf{OPF}_{\mathcal{D},\mathcal{R}}$, Agrawal et al. [1] were the first to uniformly draw OPFs from the complete set of $\mathsf{OPF}_{\mathcal{D},\mathcal{R}}$. In particular, $M = |\mathcal{D}|$ numbers are drawn uniformly at random from $\mathcal{R} = \{1, \ldots, N\}$. Then, these ciphertexts are sorted in ascending order, resulting in a sequence $c_1, \ldots, c_M$. The encryption function $f(p)$ is defined as $f(i) := c_i$ for all $1 \leq i \leq M$. Accordingly, the decryption function is $f^{-1}(c_i) = i$ for all $1 \leq i \leq M$.

Later, this approach was analyzed by [3] and termed as the "ideal object". To evaluate the security features of OPE, researchers have considered existing, as well as novel security notions. We now outline existing security analyses of the "ideal object", as well as weaknesses that have been discovered.

2.2 Weaknesses of the Ideal Object

The initial analysis of the security of the "ideal object" focused on *Indistinguishability under Ordered Chosen-Plaintext Attack (IND-OCPA)*, which is a natural adaptation of the notion of *Indistinguishability under Chosen-Plaintext Attack (IND-CPA)* to OPE [3]. The basic idea behind this notion is that an adversary must not be able to distinguish which one of two chosen plaintexts has been encrypted by the left-right oracle $\mathcal{LR}$. It is easy to see that this notion is not directly applicable to OPE as ciphertexts leak the order of plaintexts. Therefore, Boldyreva et al. introduce the weakened notion of IND-OCPA, where an attacker is allowed to only present pairs $(p_0^1, p_1^1), \ldots, (p_0^q, p_1^q)$ of plaintexts to the $\mathcal{LR}$ oracle such that $p_0^i < p_0^j \iff p_1^i < p_1^j$ for $1 \leq i, j \leq q$. Using a so-called *big-jump attack*, the authors show that any OPE scheme can only achieve IND-OCPA if the size of $\mathcal{R}$ is exponential in the size of $\mathcal{D}$.

While a range size that is exponential in the size of the domain is a necessary condition, it is not sufficient for an approach to provide IND-OCPA. Accordingly, [12] showed that the "ideal object" is not even able to achieve IND-OCPA for $|\mathcal{R}|$ being exponential to $|\mathcal{D}| = 2$. In order to emphasize that the "ideal object" should be considered more carefully in

the analysis of OPE, the authors construct a scheme that is indeed able to provide IND-OCPA in this case. Finally, the authors introduce the notion of *Indistinguishability under Ordered and Local Chosen-Plaintext Attack (IND-OLCPA)* as a further weakened version of IND-OCPA that prevents big-jump attacks by restricting oracle queries to a plaintext interval being at most polylogarithmic in $|\mathcal{R}|$. They show that adversaries can achieve higher advantage against the "ideal object" than against a generalized OPE scheme using *small-jump attacks* (see Section 2.3).

With IND-OCPA or IND-OLCPA not being achieved by the "ideal object", the expected number of bits z_h of a plaintext that remain secret against a known plaintext attack is estimated [11]. For ciphertexts of h chosen plaintexts being disclosed, the authors derive the security bounds of $z_h = \Theta\left(\log \frac{|\mathcal{D}|-h}{h+1}\right)$ for a uniformly-chosen challenge ciphertext based on their analysis. They conclude that for $|\mathcal{R}| \geq |\mathcal{D}|^3$ and $h = o(|\mathcal{D}|^\epsilon)$, where $0 < \epsilon < 1$, a ROPF is able to achieve one-wayness. However, this only applies for classical one-wayness, i.e., aimed to recover the *exact* plaintext of a given ciphertext, not the ability of an adversary to correctly estimate a plaintext *close* to the actually underlying plaintext. Therefore, while considering the disclosure of plaintext-ciphertext pairs, these results are only of limited use for estimating the impact of disclosure on the ability of an adversary to infer information about a given ciphertext.

In order to incorporate the ability of an attacker to estimate the underlying plaintext of a ciphertext in the analysis of the one-wayness of the "ideal object", the advantage of an adversary regarding *Window One-Wayness (WOW)* and *Window Distance One-Wayness (WDOW)* has been considered in [4]. Here, given a challenge set of z uniformly-chosen ciphertexts, the advantage of an attacker is the ability to correctly guess a window of size r in which at least one of the underlying plaintexts respectively at least one of the distances between two plaintexts of the given ciphertexts is within. Considering the advantage of an adversary in determining a relevant plaintext interval instead of an exact plaintext, these notions provide a more generalized version of one-wayness that incorporates the fuzziness of information leakage. In their work, the authors show that, given z challenge ciphertexts and the smallest possible window size of $r = 1$, the "ideal object" is able to provide WOW and WDOW. However, for larger windows of size $r \approx z/\sqrt{|\mathcal{D}|}$, the "ideal object" does not achieve WOW or WDOW. Although, aside from the disclosure of challenge ciphertexts, the authors consider the disclosure of a few plaintext-ciphertext pairs, they only provide a rough suggestion for the size of the range $|\mathcal{R}| \geq 7\,|\mathcal{D}|$ for their analysis to hold in this case. Hence, it remains uncertain how the increasing disclosure of information will affect the accuracy of the estimation that an attacker is able to achieve.

To obtain a concept involving one-wayness, partial indistinguishability, and information disclosure properties, [7] only recently introduced $(\mathcal{X}, \theta, q)$-indistinguishability. Here, an adversary is presented plaintexts m_1^*, m_2^* which satisfy $|m_1^* - m_2^*| \leq \theta$ together with q observed plaintext-ciphertext-pairs whose plaintexts were sampled using distributions $\mathcal{X} = (\mathcal{X}_i)_{i=1..q}$. Furthermore, it is given a ciphertext resulting from the encryption of either m_1^* or m_2^* with probability 1/2. The attacker's advantage measures its ability to correctly guess whether m_1^* or m_2^* was chosen.

The authors propose a novel OPE scheme, that is able to provide indistinguishability of plaintexts m_1^*, m_2^* differing only in their $\lfloor \log \theta \rfloor$ lower-order bits. Additionally, they note that this concept can be seen as a generalization of WOW.

In summary, existing research has shown that the "ideal object" is not the most secure way of realizing OPE. Furthermore, a number of general limitations in the security of OPE schemes have been identified. So, it is a necessary condition that $|\mathcal{R}|$ is exponential in $|\mathcal{D}|$ to achieve IND-OCPA. Higher-order plaintext bits are exposed, while a limited number of lower-order bits can be made indistinguishable when choosing an appropriate scheme. In the following, we will further concentrate on improving one-wayness and disclosure-resilience properties. Therefore, we propose alternative approaches for drawing functions from $\mathsf{OPF}_{\mathcal{D},\mathcal{R}}$ in order to lower the probability of successfully guessing a plaintext (or a set of plaintexts) whose encryption creates (respectively contains) a given ciphertext.

2.3 Alternative Realizations of OPE

Before we describe the details of our schemes in Section 3, we first provide an overview of related work that addressed the weaknesses of the "ideal object" by proposing alternative realizations of OPE.

MOPE. In order to improve the security features of the "ideal object", Boldyreva, Chenette, and O'Neill propose a *Modular Order-Preserving Encryption (MOPE)* scheme [4]. The approach is no longer strictly order-preserving – instead, by prepending a random secret shift to the plaintexts, a modular order is established among the ciphertexts. The modular encryption function for a plaintext p is defined as $\mathcal{E}nc^\star(K, j, p) = \mathcal{E}nc(K, p - j \bmod |\mathcal{D}|)$, where j is the secret random offset. Accordingly, a ciphertext c is decrypted by $\mathcal{D}ec^\star(K, j, c) = \mathcal{D}ec(K, c + j \bmod |\mathcal{D}|)$. This simple modification allows the resulting scheme to achieve optimal WOW security, while WDOW security is equivalent to the security provided by the employed OPE scheme, e.g., the "ideal object". Nevertheless, once a single plaintext-ciphertext pair has been disclosed, the security of MOPE is also reduced to the level of security provided by the underlying OPE scheme. Thus, the security gain does usually not justify the deficits of not being able to use standard database queries.

Index Tagging Schemes. Boldyreva, Chenette, and O'Neill suggest an alternative realization of OPE for static and pre-determined domains, called *Committed Efficient Orderable Encryption (CEOE)* [4]. In order to analyze this scheme, the authors suggest the notion of *Indistinguishability under Committed Chosen-Plaintext Attack (IND-CCPA)*. Here, similar to IND-OCPA, an adversary chooses two challenge vectors of the same size and order before key generation, allowing the key generation algorithm to consider them as input. Then, the advantage of an adversary is defined by the ability to correctly guess whether it is given encryptions of the first or second challenge vector. In their work, the authors propose a combination of traditional encryption and an index tagging scheme that uses a key and the domain as input, and constructs a monotone minimal perfect hash function that maps the i-th largest plaintext of the domain to the tag value i. As indicated, this scheme provides IND-CCPA, but it requires that the domain is not known to the attacker. Otherwise, it is able to infer the underlying plaintext from the index tag, which is part of the ciphertext.

A different tagging scheme is *Mutable Order-Preserving Encoding (mOPE)* [10] that is based on a mutable search tree storing the index information. While the scheme enables the use of variable domains, security relevant restrictions – like the domain being unknown to the attacker – remain.

GOPE. Motivated by the weaknesses of the "ideal object", Xiao and Yen propose a generalized approach [12], called *Generalized Order-Preserving Encryption (GOPE)*. Here, a key is defined as $\{\pi, r_{pp'} \mid 1 \leq p < p' \leq |\mathcal{D}|\}$, where $r_{pp'} \in \mathbb{Z}_3$ is randomly generated for $1 \leq p < p' \leq |\mathcal{D}|$. Furthermore, π is a permutation of the set of all possible comparisons among plaintexts $\{(x, x') \mid 1 \leq x < x' \leq |\mathcal{D}|\}$. For a plaintext p, the corresponding ciphertext c is defined as $\{(\pi(p', p), r_{pp'})) \mid p' < p\} \cup \{(\pi(p, p'), 1 + r_{pp'}) \mid p' > p\}$. Hence, a ciphertext contains $|\mathcal{D}| - 1$ elements which allows to compare $\mathcal{E}nc(p)$ to all other ciphertexts. Accordingly, in order to compare two ciphertexts c and c', they are first tested for equality, returning "=" if $c = c'$. Otherwise, the algorithm retrieves the two distinct elements with matching $i = \pi(p, p')$ from the ciphertexts, i.e., (i, s) from c and (i, s') from c'. Then, it returns "<" if $s - s' = 1$ or ">" if $s - s' = 2$. Finally, in order to decrypt a ciphertext c, two elements (i, s) and (i', s') are retrieved from its set. The underlying plaintext is the element p which appears in $\pi^{-1}(i)$ as well as $\pi^{-1}(i')$. The authors show that GOPE is able to provide both IND-OCPA and IND-OLCPA and can therefore be considered as a secure scheme for OPE.

Compared to OPF-based OPE, this approach has some disadvantages concerning its practical application. First, it does not provide a continuous numeric ciphertext space, which may be a necessary property for an easy adaption of existing database systems and the underlying implementation of index structures to handle ciphertexts produced by GOPE. Second, in GOPE, the comparison operation is more complex since it requires the lookup of a matching pair in the sets of both ciphertexts. This may restrict the applicability of GOPE when extending cloud computing services relying on large-scale database systems with OPE. Finally, GOPE requires $l = \mathcal{O}(\lceil \log_2 |\mathcal{D}| \rceil \cdot (|\mathcal{D}| - 1))$ bits to encode a ciphertext. Accordingly, the practical applicability of GOPE is limited to extremely small domain sizes.

3. OPF CONSTRUCTION SCHEMES

Using the "ideal object", each OPF has the same probability of being chosen for encryption. This, however, leads to a hypergeometric distribution of the underlying plaintexts for each ciphertext, yielding a very dominant peak, i.e., a very probable m.l.p., along the diagonal of domain and range [3, 4]. Therefore, instead of choosing functions uniformly from $\mathsf{OPF}_{\mathcal{D},\mathcal{R}}$, we use construction schemes that have different probabilities of drawing the different OPFs. By preferring functions that are not assigning plaintexts to ciphertexts close to the diagonal, we achieve a distribution of the underlying plaintexts that is closer to uniform.

3.1 Attacker Model and Design Goals

Similar to the notion of one-wayness, we consider an adversary trying to estimate the plaintext responsible for producing an observed ciphertext. For this, it has knowledge of the OPF construction scheme and possibly a set of previously observed ciphertexts or plaintext-ciphertext pairs. Additionally, we study the case where the adversary is allowed to see the ciphertexts of a set of *chosen* plaintexts *before* being presented the challenge ciphertext to decrypt.

For attacks considering the disclosure of information, the general assumption in the literature is that plaintexts are picked uniformly at random from $\mathcal{D}$ and that all ciphertexts and plaintext-ciphertext pairs have the same probability of being disclosed to an attacker [3, 4, 11, 12]. While we follow this assumption for plaintexts, we argue that the probability for a ciphertext or plaintext-ciphertext pair to be disclosed not only depends on the distribution of plaintexts, but also on the probability of a specific ciphertext or plaintext-ciphertext pair of being produced by the chosen OPFs. Therefore, in our attacker model, the probability that a ciphertext or plaintext-ciphertext pair is observed by an attacker depends on the probability that the OPFs chosen by the used construction scheme contain a respective plaintext-ciphertext mapping.

An OPF construction scheme should ideally be able to return OPFs f in such a way, that, given any challenge ciphertext c, all potential plaintexts p have approximately the same probability of satisfying $f(p) = c$ (if c lies in an interval of size $|\mathcal{D}|$ at the edge of $\mathcal{R}$, some plaintexts cannot produce c). In this article, a construction scheme with this property is called *disclosure-resilient*. Considering all functions produced by the scheme, the property enforces that a high number of plaintexts is mapped to any given ciphertext (with significant probability). Hence, it prevents adversaries from accurately guessing respective plaintexts. This condition should also hold for adversaries acquiring additional knowledge before making their guess. Here, we consider the (random) observation of a limited set of ciphertexts or plaintext-ciphertext pairs, as well as the ability to query the ciphertext of a limited number of choosen plaintexts. Although this additional information reduces the number of potentially underlying plaintexts of a challenge ciphertext, their probability of having been assigned to this ciphertext should still be distributed as uniformly as possible over the remaining subspaces.

3.2 Random Offset Addition

Our first approach for constructing OPFs, called *random offset addition*, generalizes the scheme proposed by Xiao, Yen, and Huynh [12]. In their scheme, which is defined for $\mathcal{D} = \{1, 2\}$, the plaintext 1 is encrypted to a random element $r \in [1, N - 1]$ while 2 is always assigned to $r + 1$. Since the authors only use this approach to show that the "ideal object" is not the most secure OPE scheme, we provide a simple extension for an arbitrary $M = |\mathcal{D}|$ here. Accordingly, in our scheme, an OPF is constructed by first choosing a random offset $r \in [1, N - M + 1]$. Then, for each plaintext p, the corresponding ciphertext is given by the encryption function $f(p) = p + r$.

When an OPF is generated by the random offset addition approach, each ciphertext c is produced by all possibly responsible plaintexts $p \in [\max(1, M + c - N), \min(c, M)]$ with equal probability. This property gives random offset addition near-optimal disclosure-resilience in situations when only the challenge ciphertext is known to an adversary. However, as soon as further information is available, its security properties break down. Once a single plaintext-ciphertext pair has been disclosed, an attacker learns about the random offset r and is hence able to decrypt all ciphertexts. Furthermore, with an increasing number of known ciphertexts, an

attacker is able to narrow down the potential range used by the OPF, eventually allowing her to reverse the encryption by inferring unknown ciphertexts between the known ones. Additionally, the use of domains $|\mathcal{D}| > 2$ breaks security under IND-OCPA, since the fixed ciphertext differences allow to distinguish the ciphertexts resulting from oracle query $(1, 1), (2, M)$.

Despite these vulnerabilities against informed adversaries, studying the random offset addition approach provides interesting reference values for the possible disclosure-resilience in uninformed scenarios.

3.3 Random Uniform Sampling

To obtain an OPF construction scheme that is more resilient against attacks based on further knowledge, we propose to study the *random uniform sampling* algorithm shown in Alg. 1. Its concept is similar to that of the NHGD scheme proposed in [3]. However, instead of choosing ciphertexts based on a negative hypergeometric distribution (hence emulating the "ideal object"), it relies on a uniform distribution to prevent the formation of dominant peaks.

Algorithm 1 Random Uniform Sampling

```
1: function RAND-UNIF-SAMPLE(M, N)
2:     f ← {}
3:     SAMPLE(f, 1, M, 1, N)
4:     return f
5: end function

6: procedure SAMPLE(f, d_min, d_max, r_min, r_max)
7:     p ←$ [d_min, d_max]   ▷ select random splitting element
8:     m_S ← p − d_min       ▷ number of plaintexts p' < p
9:     m_L ← d_max − p       ▷ number of plaintexts p' > p
10:    c ←$ [r_min + m_S, r_max − m_L]   ▷ randomly select c
11:    f ← f ∪ {(p, c)}      ▷ add (p, c) to OPF f

12:    ▷ recursively sample lower subspace
13:    if p > d_min then
14:        SAMPLE(f, d_min, p − 1, r_min, c − 1)
15:    end if

16:    ▷ recursively sample upper subspace
17:    if p < d_max then
18:        SAMPLE(f, p + 1, d_max, c + 1, r_max)
19:    end if
20: end procedure
```

In this approach, the RAND-UNIF-SAMPLE procedure initializes the OPF f with {} and invokes the initial call to SAMPLE with f, the minimum and maximum element of $\mathcal{D}$ ($d_{min} = 1$ and $d_{max} = M$), as well as the minimum and maximum element of $\mathcal{R}$ ($r_{min} = 1$ and $r_{max} = N$). The SAMPLE function picks a plaintext p from $[d_{min}, d_{max}]$ as *splitting element* (line 7). Here, we distinguish between two splitting strategies, namely choosing p uniformly at random as shown in Alg. 1, or using the median, i.e., the middle element of $[d_{min}, d_{max}]$ (or one of the two middle elements chosen uniformly at random in case of an even number of plaintexts). Having selected p, a respective ciphertext c is chosen uniformly at random from $[r_{min} + m_S, r_{max} - m_L]$, where $m_S = p - d_{min}$ and $m_L = d_{max} - p$ are the numbers of plaintexts smaller and larger than p. The resulting pair (p, c) is added to f (line 11), dividing both the domain and the range into subspaces. In particular, the lower subspace has domain $[d_{min}, p-1]$ and range $[r_{min}, c-1]$, whereas the upper subspace has domain $[p+1, d_{max}]$ and range $[c+1, r_{max}]$. Each subspace with non-empty domain is then recursively sampled using the SAMPLE function (lines 14 and 18). Once all calls to SAMPLE are finished and therefore all plaintexts have been assigned to a ciphertext, RAND-UNIF-SAMPLE returns f.

Note that using an argument as applied in [12], it is possible to show that for $|\mathcal{D}| = 2$, the random uniform sampling scheme does not achieve the same level of IND-OCPA as the random offset addition approach.

3.4 Random Subrange Selection

Our third approach, referred to as *random subrange selection*, builds on drawing an OPF from a randomly chosen subrange N' of the original range N. For the actual sampling of an OPF from this subrange, an alternative OPF construction scheme is used. Although the OPFs constructed with a specific subrange N' may still feature specific most likely plaintexts for each ciphertext, the randomization step spreads the most likely plaintexts of the subranges over the full domain. Hence, this reduces the overall probability of the most likely plaintexts of each ciphertext. Accordingly, the subrange selection scheme performs the following steps to construct an OPF:

1. Randomly decide whether to choose a lower bound or an upper bound first.
2. According to the previous choice, select the upper and lower bounds uniformly at random as follows:
   a) If a lower bound is to be chosen first, draw the lower bound $r_{min} \in [1, N-M+1]$. Then, draw the upper bound $r_{max} \in [r_{min} + M - 1, N]$.
   b) Otherwise, draw the upper bound $r_{max} \in [M, N]$ and the lower bound $r_{min} \in [1, r_{max} - M + 1]$.
3. Sample an OPF from the domain $[1, M]$ and the range $[1, r_{max} - r_{min} + 1]$ using an alternative construction scheme. In this work, we consider two such schemes, namely the "ideal object" and the random uniform sampling scheme proposed in Section 3.3.
4. Finally, adjust the range of the obtained OPF by adding $r_{min} - 1$ to all ciphertexts and return the resulting function.

Note, that the expected size of the used subrange remains linear in N. Furthermore, the expected subrange interval lies symmetrically around $\frac{N+1}{2}$. For $|\mathcal{D}| = 2$, this reduces the IND-OCPA properties of the random subrange selection approach to that of the approach used in the subrange. In particular, it is again possible to adopt the argument given in [12], to show that *for this specific domain size* random offset addition provides a higher level of IND-OCPA than random subrange selection in combination with the "ideal object" or a random uniform sampling approach.

4. SECURITY EVALUATION

In order to evaluate the security features of our proposed OPF construction schemes and to compare them to the

"ideal object", we first introduce two novel metrics that are able to assess the attacker's ability of accurately guessing the underlying plaintexts of ciphertexts.

4.1 Evaluation Metrics

Before we consider the proposed metrics in detail, we first define the following prerequisites.

Let $\mathcal{S}$ be a randomized construction scheme that produces each order-preserving function $f \in \mathsf{OPF}_{\mathcal{D},\mathcal{R}}$ with probability $\Pr(f)$. Then, the probability of a range element $c \in \mathcal{R}$ being a value of the OPF f produced by $\mathcal{S}$ is given by:

$$\Pr(c \in f(\mathcal{D})) \;=\; \sum_{\substack{f \in \mathsf{OPF}_{\mathcal{D},\mathcal{R}} \\ c \in f(\mathcal{D})}} \Pr(f)$$

Accordingly, the probability of a pair $(p, c) \in \mathcal{D} \times \mathcal{R}$ being a plaintext-ciphertext pair produced by f is defined as follows:

$$\Pr(f(p)=c) \;=\; \sum_{\substack{f \in \mathsf{OPF}_{\mathcal{D},\mathcal{R}} \\ f(p)=c}} \Pr(f)$$

4.1.1 Number of Significant Plaintexts

For each $\alpha \in [0, 1]$ and $c \in \mathcal{R}$, we define the *number of significant plaintexts for ciphertext c and threshold α* as a random variable of value

$$M_\alpha^{\mathcal{S}}(c) = \min\{|Q| \mid Q \subseteq \mathcal{D} \wedge \alpha \leq \sum_{p \in Q} \Pr(f(p)=c \mid c \in f(\mathcal{D}))\}$$

if $\Pr(c \in f(\mathcal{D})) > 0$ and of value 0 otherwise. In the first case, this metric measures the cardinality of the smallest set of plaintexts which has at least probability α of containing the plaintext that is mapped to c by a function f randomly constructed using scheme $\mathcal{S}$. Generally speaking, the higher the value of $M_\alpha^{\mathcal{S}}(c)$, the higher is the disclosure-resilience of $\mathcal{S}$ when considering only ciphertext c.

To obtain a metric summarizing over all ciphertexts, we use $M_\alpha^{\mathcal{S}}(c)$ to define *the average number of significant plaintexts for threshold α* as a random variable $M_\alpha^{\mathcal{S}}$:

$$M_\alpha^{\mathcal{S}} := \frac{1}{|\mathcal{D}|} \sum_{c \in \mathcal{R}} \Pr(c \in f(\mathcal{D})) \cdot M_\alpha^{\mathcal{S}}(c)$$

As mentioned earlier we will empirically estimate $M_\alpha^{\mathcal{S}}(c)$ and $M_\alpha^{\mathcal{S}}$ by conducting simulations.

Note, that there is a connection to the concept of r, z-Window One-Wayness introduced in [4]. Here, an adversary has to determine a domain interval of size r containing the plaintext being mapped to one of z randomly chosen ciphertexts. The advantage of the adversary equals its probability of success. For $z = 1$ and presented ciphertext c, the adversary has to return an interval of at least size $M_\alpha^{\mathcal{S}}(c)$ to achieve an advantage of value α. For higher values of z a similar relationship depends on the question, whether the presented ciphertexts have disjoint sets of significant plaintexts. Furthermore, note that the WOW interval size r and $M_\alpha^{\mathcal{S}}(c)$ may significantly differ for multimodal plaintext distributions with maxima lying far apart. $M_\alpha^{\mathcal{S}}(c)$ will better reflect the behavior of an optimal attacker in this case.

4.1.2 Expected Estimation Error

For most applications of order-preserving functions (e.g., geographic locations or account balances), information is disclosed not only by the exact decryption of a given ciphertext c, but also if it is possible to estimate a plaintext lying in a narrow interval around the actual plaintext producing c. Therefore, we are interested in the expected estimation error of a maximum-likelihood attacker, estimating the plaintext mapped to a given ciphertext c by choosing uniformly at random one element from the set $\mathrm{mlp}^{\mathcal{S}}(c)$:

$$\mathrm{mlp}^{\mathcal{S}}(c) = \arg\max_{p \in \mathcal{D}} \Pr(f(p) = c)$$

For the "ideal object", it is known, that this estimator has maximum error of $O(\sqrt{|\mathcal{D}|})$ with probability arbitrarily close to one [4]. The *expected estimation error* the adversary achieves for ciphertext c and scheme $\mathcal{S}$ can be written as:

$$E^{\mathcal{S}}(c) = \sum_{p \in \mathcal{D}} \Pr(f(p) = c) \cdot \frac{\sum_{mlp \in \mathrm{mlp}^{\mathcal{S}}(c)} |mlp - p|}{|\mathrm{mlp}^{\mathcal{S}}(c)|}$$

Weighting this value with the probabilities of actually observing ciphertext c, we obtain the *expected estimation error of the maximum-likelihood attacker considering scheme $\mathcal{S}$*:

$$E^{\mathcal{S}} = \frac{1}{|\mathcal{D}|} \sum_{c \in \mathcal{R}} \Pr(c \in f(\mathcal{D})) \cdot E^{\mathcal{S}}(c)$$

Again, when comparing different OPF construction schemes, a higher value indicates increased disclosure-resilience. In our evaluation, we will empirically estimate $E^{\mathcal{S}}(c)$ and $E^{\mathcal{S}}$ for the considered schemes. Please note that, while the expected estimation error may provide insight into the estimation accuracy that an adversary may achieve, a maximum-likelihood estimator does not necessarily present an optimal attacker model for some cases. For example, given a bimodal distribution of underlying plaintexts with two maxima, a maximum-likelihood attacker would have to decide for one maximum, resulting in a larger estimation error in case of a poor choice.

4.1.3 Adversaries with Additional Knowledge

As outlined in Section 3.1, we are also interested in the disclosure-resilience properties of the studied schemes when the adversary already possesses a limited set of previously observed ciphertexts or plaintext-ciphertext pairs. Accordingly, given $c_1, \ldots, c_z \in \mathcal{R}$ or $(p_1, c_1), \ldots, (p_z, c_z) \in \mathcal{D} \times \mathcal{R}$ together with a metric $\phi \in \{M_\alpha^{\mathcal{S}}(c), M_\alpha^{\mathcal{S}}, E^{\mathcal{S}}(c), E^{\mathcal{S}}\}$, we write $\phi_{|c_1,\ldots,c_z}$ or $\phi_{|(p_1,c_1),\ldots,(p_z,c_z)}$ to denote the version of ϕ where only OPFs f satisfying $c_1, \ldots, c_z \in f(\mathcal{D})$ or $f(p_1) = c_1, \ldots, f(p_z) = c_z$ are considered, respectively. In particular, all involved probabilities are subject to these conditions.

Then, we define the metric ϕ *under the condition of z known ciphertexts* by computing a weighted average over all possible combinations:

$$\phi_{|z \cdot c} = \frac{1}{\left|\binom{\mathcal{D}}{z}\right|} \sum_{\{c_1,\ldots,c_z\} \in \binom{\mathcal{R}}{z}} \Pr(c_1, \ldots, c_z \in f(\mathcal{D})) \cdot \phi_{|c_1,\ldots,c_z}$$

Note, that the individual probabilities $\Pr(c_1, \ldots, c_z \in f(\mathcal{D}))$ sum up to a value of $\left|\binom{\mathcal{D}}{z}\right|$, which is corrected by a normalizing term.

Accordingly, the metric ϕ *under the condition of* z *known plaintext-ciphertext pairs* is defined as

$$\phi_{|z\cdot(p,c)} = \frac{1}{\left|\binom{\mathcal{D}}{z}\right|} \sum_{P \in \binom{\mathcal{D}\times\mathcal{R}}{z}} \Pr\left(\bigwedge_{(p,c)\in P} f(p) = c\right) \cdot \phi_{|P}$$

4.1.4 The Case of Chosen Plaintexts

Next, we consider the scenario where an adversary is allowed to see the ciphertexts of z chosen plaintexts before being presented a challenge ciphertext. Again, let $\phi \in \{M_\alpha^{\mathcal{S}}, E^{\mathcal{S}}\}$ be one of the metrics defined in Sections 4.1.1 and 4.1.2.

Depending on the OPF construction scheme $\mathcal{S}$, the chosen plaintexts $p_1, \ldots, p_z$ are mapped to ciphertexts $c'_1, \ldots, c'_z$ with different probability. Considering the consequences of all possible mappings, we compute the *expected value of* $\phi(c)$ *for chosen plaintexts* $p_1, \ldots, p_z$ *and observed challenge* c as:

$$\overline{\phi(c)}_{|p_1,..,p_z} :=$$

$$\sum_{(c'_1,\ldots,c'_z)\in\mathcal{R}^z} \Pr\left(\bigwedge_{i\in[1,z]} f(p_i) = c'_i \,\middle|\, c \in f(\mathcal{D})\right) \cdot \phi(c)_{|(p_1,c'_1)..(p_z,c'_z)}$$

To obtain a metric summarizing over all possible ciphertexts, we compute the *average expected value of* ϕ *for chosen plaintexts* $p_1, \ldots, p_z$ as:

$$\overline{\phi}_{|p_1,\ldots,p_z} := \frac{1}{|\mathcal{D}|} \sum_{c\in\mathcal{R}} \Pr(c \in f(\mathcal{D})) \cdot \overline{\phi(c)}_{|p_1,\ldots,p_z}$$

Since the adversary cannot predict the challenge ciphertext when choosing plaintexts, she will be tempted to query a combination $p_1, \ldots, p_z$ which leads to worst-case (i.e., minimum) global disclosure-resilience. Therefore, we can adopt its value as a metric named *expected* ϕ *under* z *chosen plaintexts*:

$$\overline{\phi}_{|z\cdot p} := \min_{(p_1,\ldots,p_z)\in\mathcal{D}^z} \overline{\phi}_{|p_1,\ldots,p_z}$$

Note that for all OPF construction schemes considered in this paper, the worst-case occurs for the choice of (approximately) equally spaced plaintexts, partitioning the domain into (approximately) equal parts.

4.2 Simulation Setup

We implemented the "ideal object", the random offset addition, the random subrange selection, and both variants of the random uniform sampling approach in C++ using the *Boost.Random* library[1] and its *Mersenne Twister* implementation [8] for pseudo-random number generation.

First, we studied the disclosure-resilience of these schemes against adversaries without the knowledge of further ciphertexts or plaintext-ciphertext pairs. For each construction scheme, we generated 10^8 OPFs using a domain size of $M = 500$ and a range size of $N = 5000$. While generating the OPFs, we recorded the frequencies of plaintext-ciphertext pairs occurring among these functions.

Then, in order to empirically estimate the average number of significant plaintexts and the expected estimation error, we generated 10^8 OPFs for each approach and computed an estimation of both metrics assuming the prior knowledge or choice of up to $z = 2$ ciphertexts, plaintext-ciphertext pairs, and plaintexts, respectively. We varied the domain size $M \in \{10, 20, 30\}$ and chose a range size of $N = M^2$. Due to memory constraints and the complexity of computing $\phi_{|z\cdot c}$, $\phi_{|z\cdot(p,c)}$, and $\overline{\phi}_{|z\cdot p}$, which requires the collection of the frequency distribution of all occurring ciphertext and plaintext-ciphertext pair combinations, we were limited in our evaluation to these rather small domain and range sizes, as well as to $z = 2$.

[1] `http://www.boost.org/libs/random/`

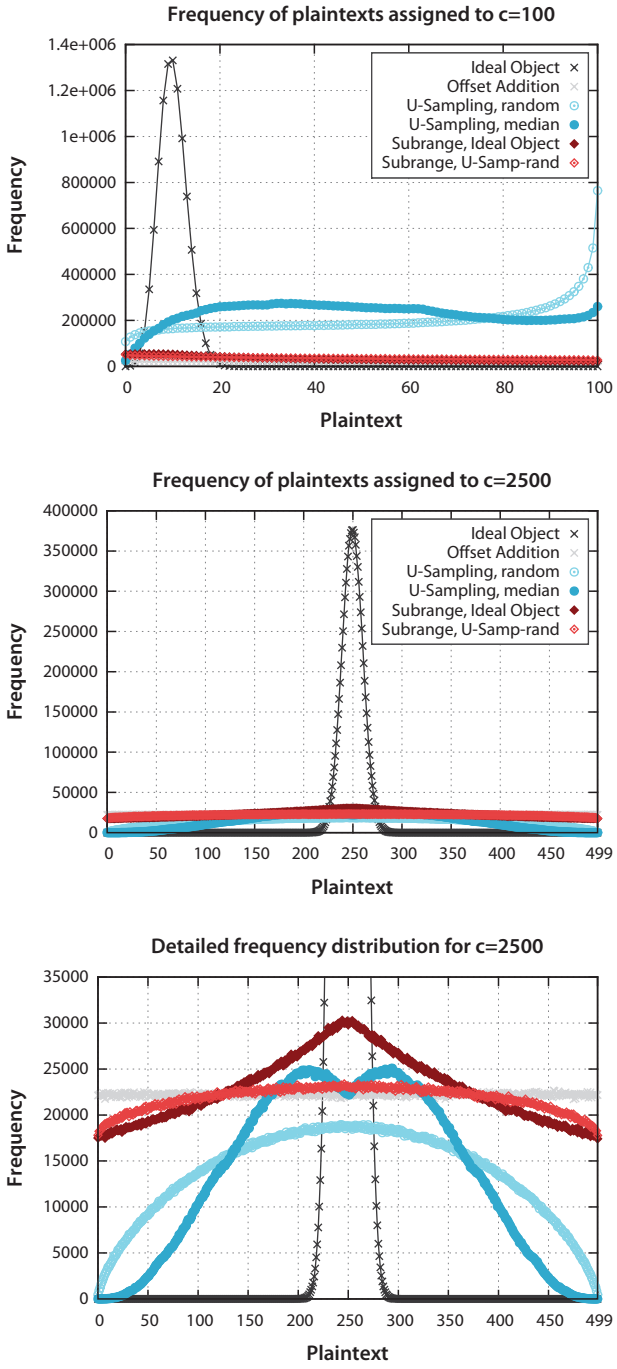

Figure 1: Empirically measured frequency distributions for specific ciphertexts (10^8 OPFs).

5. RESULTS & DISCUSSION

We now provide an overview and a discussion of the simulation results.

5.1 Plaintext-Ciphertext Assignments

For the first experiment, we studied the frequency distribution of plaintexts that had been assigned a specific ciphertext. Fig. 1 shows the measured frequency distributions for the ciphertexts $c = 100$ and $c = 2500$. We chose these two ciphertexts in order to compare the approaches at the edge of the domain, where a range element c can only be assigned to domain elements $\{p \in \mathcal{D} \mid p \leq c \leq |\mathcal{R}| - (|\mathcal{D}| - p)\}$, as well as the middle of the range, where these restriction do not play a role. Accordingly, for $c = 100$, the distribution does not cover the full domain as plaintexts $p > 100$ cannot be assigned to $c = 100$. Note furthermore, that, since the compared approaches may use different ciphertexts with different probabilities, in this plot, the sum of the frequencies does not have to be equal for different schemes.

As we can see from Fig. 1, the "ideal object" closely follows a hypergeometric distribution according to [3, 4], yielding a frequency of over $1.3 \cdot 10^6$ assignments of $c = 100$ to $p = 10$ and over $3.7 \cdot 10^5$ assignments of $c = 2500$ to $p = 250$. Hence, although the absolute frequency reduces for the ciphertext in the middle of the range, it still shows a prominent peak compared to our proposed schemes in both situations.

In contrast, the random offset approach has a constant frequency of about $2 \cdot 10^4$ assignments over the (assigned) domain both for $c = 100$ and $c = 2500$. Both random subrange selection schemes only slightly vary between $2 \cdot 10^4$ to $5 \cdot 10^4$ for $c = 100$ and about $1.7 \cdot 10^4$ to $3 \cdot 10^4$ assignments for $c = 2500$. Comparing the approaches, using random uniform sampling instead of the "ideal object" to draw OPFs from the subrange yields a more even distribution of ciphertexts.

Finally, the direct application of both random uniform sampling approaches shows a slightly higher maximum frequency compared to the offset and subrange selection schemes for $c = 100$. Despite the presence of a peak at $p = 100$, the uniform sampling schemes still show a roughly uniform distribution over the largest fraction of the usable part of the domain. For $c = 2500$, both feature a non-uniform frequency distribution decreasing towards the edges of the domain.

In addition to the plaintext distributions for two specific ciphertexts, for the first experiment, we plotted the measured plaintext-ciphertext assignments over the full domain and range. Fig. 2 shows the corresponding heatmaps obtained from the frequency of plaintext-ciphertext assignments occurring among the generated functions for each approach. Here, according to the m.l.p. deduced in [4], we see that the "ideal object" yields a dominant peak of plaintext-ciphertext assignments along the diagonal of the plot. Furthermore, for the offset addition approach, we can see an almost perfectly uniform distribution over the domain and range. While the random subrange selection approaches still show a slightly higher frequency around the diagonal, they drastically reduce the height of the peak when compared to the "ideal object". For the random uniform sampling approaches, we see that both schemes yield the tendency to encrypt small plaintexts to small ciphertexts and large plaintexts to large ciphertexts. The reason for this may be the fact that, as soon as a splitting element p is assigned a ciphertext near the lower (upper) edge of $\mathcal{R}$, the subrange available to unassigned domain elements $p' < p$ ($p' > p$) becomes so small that the their ciphertexts have to be very close to each other. Due to the continued repetition of the splitting random experiment, it is probable that this case eventually occurs. Accordingly, for plaintexts from the beginning of $\mathcal{D}$, it is more likely to be a assigned a ciphertext from the beginning of $\mathcal{R}$, while ciphertexts from the end of $\mathcal{R}$ tend to be assigned to plaintexts from the end of $\mathcal{D}$.

In summary, these results give indication that the proposed schemes are able to improve the security properties of OPF-based OPE by reducing the probability of the most likely plaintexts and spreading the distribution of possible plaintexts over the whole domain.

5.2 Quantifying the Disclosure-Resilience

In order to quantify the improvement of the security properties of the suggested schemes and to compare them to the "ideal object", we now discuss our results regarding the average number of significant plaintexts and the expected estimation error considering ciphertext-only as well as known and chosen plaintext attacks.

5.2.1 Number of Significant Plaintexts

Fig. 3 shows the average number of significant plaintexts $M^{\mathcal{S}}_{\alpha=0.5}$ for the "ideal object" and our proposed schemes for different domain and range sizes. We chose $\alpha = 0.5$ since, in this case, $M^{\mathcal{S}}_{\alpha}$ corresponds to the number of plaintexts that an attacker has to consider for them to contain the underlying plaintext of a challenge ciphertext with a probability of at least 50 %. For $M = 10$, according to our expectations, we can see that in case that no additional information (e.g., ciphertexts or plaintext-ciphertext pairs) has been disclosed, the random offset addition approach has the highest number of plaintexts that have to be considered by an adversary. We can also see that, since $M^{\text{offset}}_{\alpha=0.5} \approx 4.9$, which roughly corresponds to 50 % of M, the offset addition approach provides a nearly uniform distribution of plaintexts that have been assigned to each ciphertext.

Note, that $M^{\text{offset}}_{\alpha=0.5}$ is not exactly 5.0 since ciphertexts at the beginning and the end of the range have a smaller number of potentially underlying plaintexts. Therefore, for these ciphertexts, the number of significant plaintexts is smaller as well, slightly decreasing the overall result.

The subrange selection approaches ($M^{\text{sr-u-samp}}_{\alpha=0.5} \approx 4.7$ and $M^{\text{sr-ideal}}_{\alpha=0.5} \approx 4.4$) and the random uniform sampling schemes ($M^{\text{u-samp-rand}}_{\alpha=0.5} \approx 3.5$ and $M^{\text{u-samp-med}}_{\alpha=0.5} \approx 2.9$) perform increasingly worse. However, all proposed schemes clearly outperform the "ideal object" with $M^{\text{ideal}}_{\alpha=0.5} \approx 1.9$. For $M = 20$ and $M = 30$, we can see similar results confirming the observations for $M = 10$.

In case of the adversary's knowledge of $z \in \{1, 2\}$ additional ciphertexts, we see in Fig. 3 that, for $M \in \{10, 20, 30\}$, the average number of significant plaintexts reduces for all proposed schemes, while the decrease is less considerable for the "ideal object". Nevertheless, even despite the reduction of $M^{\mathcal{S}}_{\alpha=0.5|z\cdot c}$, our schemes still improve on the "ideal object".

Considering the disclosure of plaintext-ciphertext pairs and the chosen plaintext scenario, Fig. 3 shows the decrease of both $M^{\mathcal{S}}_{\alpha=0.5|z\cdot(p,c)}$ and $\overline{M^{\mathcal{S}}_{\alpha=0.5|z\cdot p}}$ for all schemes, including the "ideal object". The experiment confirms the assumption that the offset addition approach is unable to provide *any* level of security as soon as a single plaintext-ciphertext pair has been disclosed. While for $z = 1$, the proposed

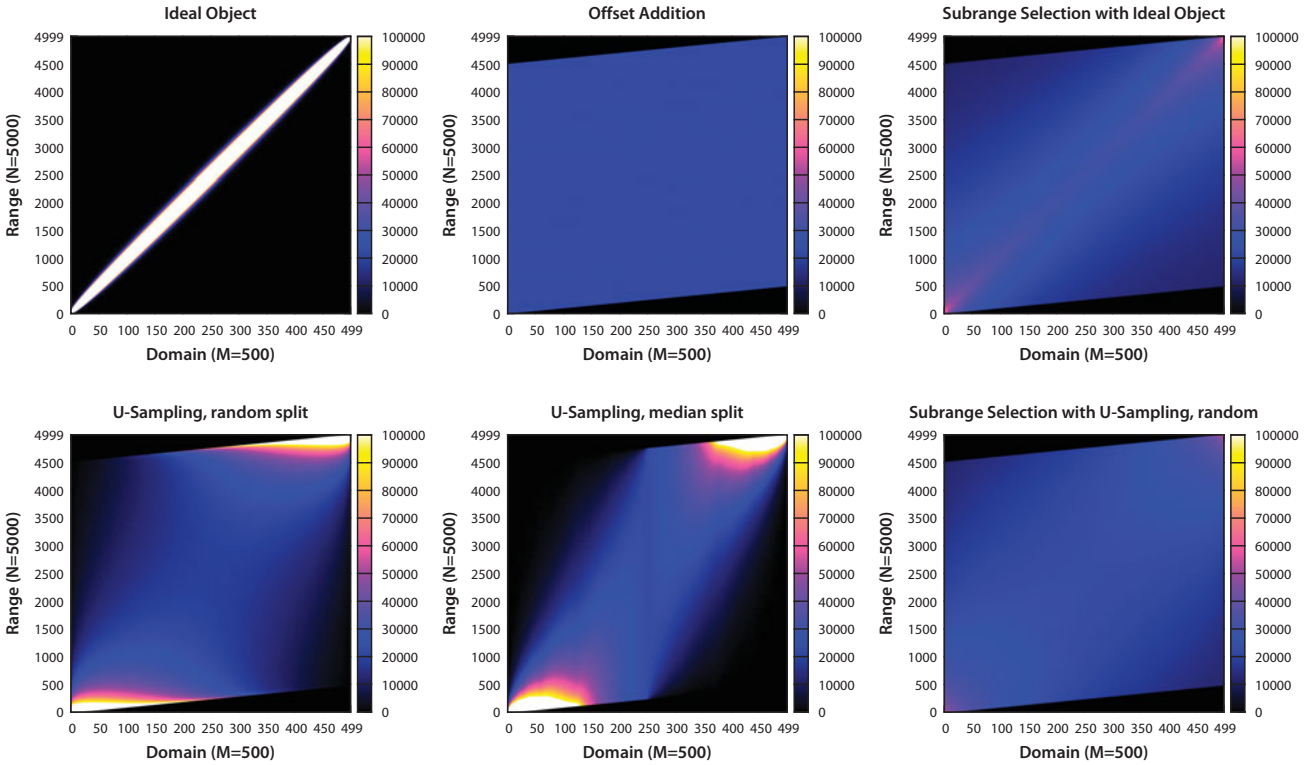

Figure 2: Empirically measured frequency distributions of plaintext-ciphertext assignments (10^8 OPFs).

schemes still clearly outperform the "ideal object" in the larger domains $M \in \{20, 30\}$, they differ only marginally for the case $z = 2$. Although the uniform sampling approaches perform slightly worse compared to the subrange selection approaches if no information or only ciphertexts are available, they are competitively resilient when considering the disclosure of plaintexts.

Generally, when using the uniform sampling approach, the random splitting strategy seems preferable over selecting the median as splitting element. Considering the studied subrange selection schemes, both variants yield approximately the same number of significant plaintexts.

5.2.2 Estimation Error

Finally, we compared the proposed schemes and the "ideal object" in terms of expected estimation error. Fig. 4 shows the values of $E^{\mathcal{S}}$ for different domain and range sizes.

Over all scenarios, using the "ideal object" leads to very low expected estimation error results, with values between $5 - 8\%$ of the respective domain size. Similar to $M^{\text{ideal}}_{\alpha=0.5|z\cdot c}$, knowledge of additional ciphertexts does not have a noticeable impact on $E^{\text{ideal}}{}_{|z\cdot c}$. However, the expected error decreases in the presence of known or chosen plaintext attacks.

The random offset addition scheme proves second-best when only the challenge ciphertext is provided to an adversary. However, it is again not able to uphold *any* level of disclosure-resilience against known plaintext-ciphertext pairs, resulting in $E^{\text{offset}}{}_{|1\cdot(p,c)} = 0$ and $\overline{E^{\text{offset}}}{}_{|1\cdot p} = 0$.

The subrange selection schemes feature high expected estimation errors given that only the challenge ciphertext is provided to an adversary (the column *'no disclosure'*). However, they turn out to be comparatively sensitive to adversaries observing additional ciphertexts or plaintext-ciphertext pairs. In this case, their expected estimation error roughly drops to the same level as that of the uniform sampling approaches. Considering the variant based on the "ideal object", the results for plaintext-ciphertext disclosure can be explained by the linear character of its m.l.p.s inside the subrange. As soon as two such pairs are known, it allows to estimate the subrange limits with high accuracy.

5.2.3 Summary

Recapitulating the results, all proposed OPF construction schemes show higher disclosure-resilience (in terms of the number of significant plaintexts and estimation error) than the "ideal object" in cases where only the challenge ciphertext or a set of ciphertexts are known to the adversary. Furthermore, the subrange selection and the random uniform sampling schemes also outperform the "ideal object" when considering known or one chosen plaintext.

However, we also note that the disclosure-resilience of all proposed schemes heavily decreases with a growing amount of additional information available to the adversary. Especially the simulations considering known or chosen plaintext show disclosure-resilience properties close to (or even worse than) the "ideal object".

Here, it is important to carefully interprete the already achieved values with respect to the possible optima. Due to the inherent structure of OPFs, the availability of just a small portion of extra information to the adversary must necessarily lead to considerable decreases of the maximum possible values of our disclosure-resilience metrics. This is

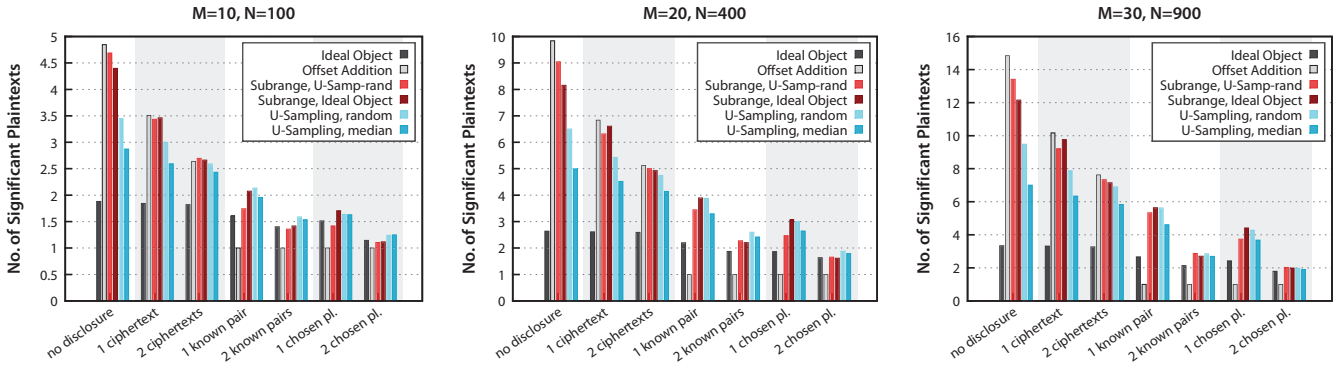

Figure 3: Number of significant plaintexts with $\alpha = 0.5$ for $M \in \{10, 20, 30\}$ and $N = M^2$.

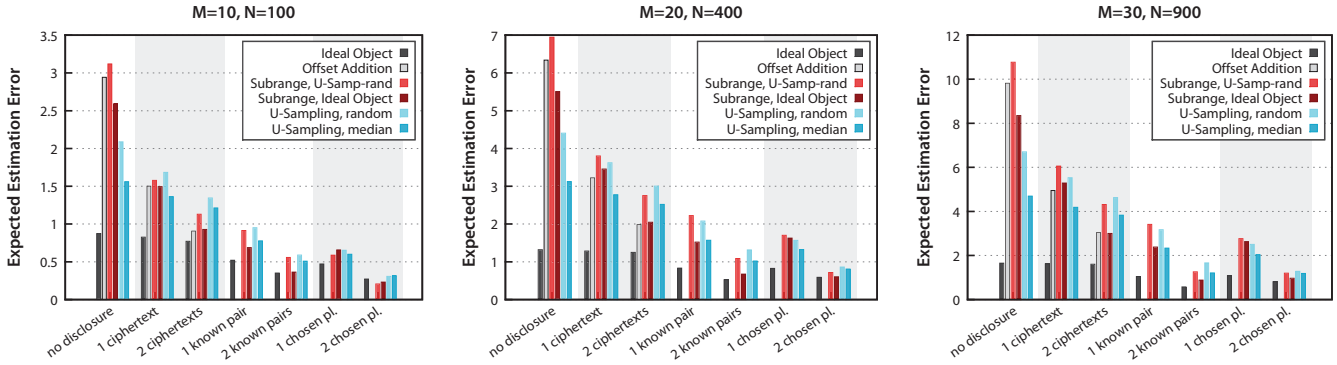

Figure 4: Expected estimation error for $M \in \{10, 20, 30\}$ and $N = M^2$.

best visualized for the chosen plaintext scenarios. Here, the most successful approach of an attacker was to choose z plaintexts splitting the domain into $z+1$ equally sized parts. Consequently, the possible maxima of the number of significant plaintexts and estimation errors drop to *less than* a $\frac{1}{z+1}$ fraction of their original value.

However, even in this light, the results obtained for the studied schemes are unsatisfactory. This conclusion clearly conveys the need to further investigate OPF construction schemes with high disclosure-resilience in the presence of well-informed and powerful attackers.

6. CONCLUSION

With OPE allowing users of cloud applications to protect their sensitive information while still enabling service providers to perform efficient query operations on the encrypted data, several approaches and security notions have been proposed for the analysis of OPE. Since it has been shown that providing indistinguishability either demands for exponential (hence inefficient) range sizes or an extensive weakening of the security notion, we studied how to enhance the one-wayness properties of OPE schemes. We argue that these notions are most important in practical applications where encrypted data may be observed during communication or due to database access. We proposed two novel metrics to enable a descriptive evaluation and comparison of the security features of OPE schemes under these conditions. Furthermore, building upon the knowledge of the weaknesses of the "ideal object", we suggested three novel schemes for OPF construction that show improved one-wayness and disclosure-resilience properties.

In our evaluation, we were able to show that all proposed schemes outperform the "ideal object" when considering ciphertext-only attacks. Moreover, in case of known or chosen plaintexts, while the random offset addition cannot provide any security features anymore, the proposed random subrange selection and random uniform sampling schemes mostly improve on the "ideal object".

However, the obtained results clearly suggest that further improvement is possible in different directions.

On the one hand, the proposed approaches should be fine-tuned. Especially, the random uniform sampling schemes show a tendency to concentrate ciphertexts at the edges of $\mathcal{R}$. To lower the probability of the responsible unbalanced recursions, the usage of alternative distributions to choose splitting elements and ciphertexts should be investigated.

On the other hand, the observed sensitivity of *all* studied OPFs construction schemes against known and chosen plaintext attacks, motivates the further investigation of alternative schemes achieving higher disclosure-resilience under such adverse conditions. Given the current results, only applications guaranteeing an at most marginal disclosure of plaintext information will profit from the proposed schemes.

Several further issues remain to be studied in future work. First of all, the mathematical modeling and analysis can be improved. Particular next steps should include a more thorough mathematical analysis of the suggested OPF construction schemes regarding the proposed and previously es-

tablished metrics. Furthermore, it would be interesting to study extensions of the proposed disclosure-resilience metrics that allow to consider different plaintext distributions. Finally, since our empirical evaluation was limited to domain sizes $M \in \{10, 20, 30\}$ and $N = M^2$, we plan to investigate the impact of increasing domain and range sizes on the disclosure-resilience of the proposed schemes.

7. ACKNOWLEDGMENTS

The authors would like to thank Amir Herzberg, as well as the anonymous reviewers for making valuable suggestions and comments and pointing us to [7].

Sander Wozniak is supported by the German Research Foundation (DFG GRK 1487: Selbstorganisierende Mobilkommunikationssysteme für Katastrophenszenarien).

8. REFERENCES

[1] R. Agrawal, J. Kiernan, R. Srikant, and Y. Xu. Order Preserving Encryption for Numeric Data. In *ACM SIGMOD*, pages 563–574, 2004.

[2] G. Bebek. Anti-Tamper Database Research: Interference Control Techniques. Technical report, EECS 433 Final Report, Case Western Reserve University, 2002.

[3] A. Boldyreva, N. Chenette, Y. Lee, and A. O'Neill. Order-Preserving Symmetric Encryption. In *Advances in Cryptology – EUROCRYPT 2009*, pages 224–241. 2009.

[4] A. Boldyreva, N. Chenette, and A. O'Neill. Order-Preserving Encryption Revisited: Improved Security Analysis and Alternative Solutions. In *Advances in Cryptology – CRYPTO 2011*, pages 578–595, 2011.

[5] H. Hacigümüş, B. Iyer, C. Li, and S. Mehrotra. Executing SQL over Encrypted Data in the Database-Service-Provider Model. In *ACM SIGMOD*, pages 216–227, 2002.

[6] B. Hore, S. Mehrotra, M. Canim, and M. Kantarcioglu. Secure Multidimensional Range Queries over Outsourced Data. *The VLDB Journal*, 21(3):333–358, 2012.

[7] T. Malkin, I. Teranishi, and M. Yung. Order-Preserving Encryption Secure Beyond One-Wayness. Cryptology ePrint Archive, Report 2013/409, 2013.

[8] M. Matsumoto and T. Nishimura. Mersenne Twister: A 623-dimensionally Equidistributed Uniform Pseudo-Random Number Generator. *ACM Transactions on Modeling and Computer Simulation (TOMACS)*, 8(1):3–30, 1998.

[9] G. Ozsoyoglu, D. Singer, and S. S. Chung. Anti-Tamper Databases: Querying Encrypted Databases. In *IFIP WG 11.3 Working Conference on Database and Applications Security*, volume 11, pages 4–6, 2003.

[10] R. A. Popa, F. H. Li, and N. Zeldovich. An Ideal-Security Protocol for Order-Preserving Encoding. In *IEEE Symposium on Security and Privacy*, pages 463–477, 2013.

[11] L. Xiao and I.-L. Yen. Security Analysis for Order Preserving Encryption Schemes. In *Information Sciences and Systems (CISS)*, pages 1–6, 2012.

[12] L. Xiao, I.-L. Yen, and D. Huynh. A Note for the Ideal Order-Preserving Encryption Object and Generalized Order-Preserving Encryption. Cryptology ePrint Archive, Report 2012/350, 2012.