

A Comparative Evaluation of Order-Revealing Encryption Schemes and Secure Range-Query Protocols

Dmytro Bogatov
Boston University
Boston, MA 02215
dmytro@bu.edu

George Kollios
Boston University
Boston, MA 02215
gkollios@bu.edu

Leonid Reyzin
Boston University
Boston, MA 02215
reyzin@bu.edu

ABSTRACT

Database query evaluation over encrypted data has received a lot of attention recently. Order Preserving Encryption (OPE) and Order Revealing Encryption (ORE) are two important encryption schemes that have been proposed in this area. These schemes can provide very efficient query execution, but at the same time may leak some information to adversaries. More protocols have been introduced that are based on Searchable Symmetric Encryption (SSE), Oblivious RAM (ORAM) or custom encrypted data structures. In this paper, we present the first comprehensive comparison among a number of important secure range query protocols using a framework that we developed. We evaluate five ORE-based and five generic range query protocols. We analyze and compare them both theoretically and experimentally and measure their performance over database indexing and query evaluation. We report not only execution time but also I/O performance, communication amount, and usage of cryptographic primitive operations. Our comparison reveals some interesting insights concerning the relative security and performance of these approaches in database settings.

PVLDB Reference Format:

Dmytro Bogatov, George Kollios and Leo Reyzin. A Comparative Evaluation of Order-Preserving and Order-Revealing Schemes and Protocols. *PVLDB*, 12(xxx): xxxx-yyyy, 2019.
DOI: <https://doi.org/TBD>

1. INTRODUCTION

Order Preserving Encryption (OPE) was proposed by Agrawal et al. [2] and has received a lot of interest recently. The main idea is to “encrypt” numerical values into ciphertexts that have the same order as the original plaintexts. This is a very useful primitive since it allows a database system to make comparisons between ciphertexts and get the same results as if it had operated on plaintexts. A scheme was proposed in [2] but no security analysis was given.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Articles from this volume were invited to present their results at The 45th International Conference on Very Large Data Bases, August 2019, Los Angeles, California.

Proceedings of the VLDB Endowment, Vol. 12, No. xxx
Copyright 2018 VLDB Endowment 2150-8097/18/10... \$ 10.00.
DOI: <https://doi.org/TBD>

Boldyreva et al. [9] were the first to treat OPE schemes from a cryptographic point of view, providing security models and rigorous analysis. The ideal functionality of such a scheme is to leak only the order of the plaintexts and nothing more. However, it was shown in Boldyreva et al. [9] that the ideal functionality is not achievable if the scheme is *stateless* and *immutable*. Furthermore, they showed that the (stateless) scheme that they proposed leaks at least half of the bits of the plaintext [10]. Since then, a number of OPE schemes have been proposed that provide different performance and security guarantees. In order to achieve the ideal functionality, Popa, Li, and Zeldovich [53] proposed a mutable scheme that maps plaintexts to their ranks and need the full state of the dataset. Notice that, an insertion or a deletion of a value may change the ranks, and therefore the ciphertexts, of multiple values. Kerschbaum [38] proposed an improvement on this scheme that hides also the frequencies of each plaintext (how many times a give value appears in the dataset).

Furthermore, in order to improve the security of these schemes, Boneh et al. [11] proposed the idea of Order Revealing Encryption (ORE). In ORE, ciphertexts have no particular order and look more like typical semantically secure encryptions. The database system has a special comparison function that can be used to compare two ciphertexts. These schemes are more secure than OPE schemes, although still leak some information, and in general are more expensive to compute. Actually, OPE can be seen as a special case of ORE. Since these schemes leak some information, a number of recent works considered attacks on systems that may use these schemes [31, 32, 51, 27, 37, 13, 23, 43, 5, 62]. Most of these attacks assume *auxiliary information* and attack the systems assuming that no other security.

OPE and ORE schemes can be used almost no changes to the underlying database search. However, to provide greater security a number of more complex protocols for securing data in outsourced databases have also been proposed. The most secure of these — so-called Oblivious RAM (ORAM) — provides strong, well-understood, cryptographic privacy guarantees with no information leakage.

Applications that can benefit of such schemes and protocols include cloud access security broker (CASB) and financial and banking applications. Indeed, a number of commercial CASBs including Skyhigh Networks [57] and CipherCloud [20] have been using some form of OPE or ORE schemes in their systems. In addition a number of financial institutions can encrypt their data using the aforementioned schemes in order to provide another layer of security on their

data assuming that the overhead is small or minimal. For many of these applications either auxiliary information that is needed for the attacks mentioned above may not be available or difficult to get.

Currently, it is a very challenging task for users to choose an appropriate data privacy approach for their outsourced application, because the security and performance tradeoff is not well understood. Both security and performance of every approach need to be thoroughly evaluated. However, characterizing security benefits of different approaches remains an open problem, unlikely to be solved in the immediate future. On the other, it is at least possible to evaluate the performance of each approach, so as to enable better-informed decisions about whether the improved performance of some schemes is worth the uncertainty about the security they achieve.

We emphasize that evaluating performance of these schemes is not a trivial task. Many of the papers presenting the above approaches provide only a theoretical treatment and concentrate more on the security definitions and analysis and less on the performance. Some of these schemes have not been even implemented properly. Furthermore, even-though the main target of these schemes are database applications, most of them have not been evaluated in database settings.

To address this problem, in this paper we design a new framework that allows for systematic and extensive comparison of OPE and ORE schemes and protocols for database applications. We employ these schemes in database indexing techniques (i.e. B+ trees) and query protocols and we report various costs including I/O complexity.

The main contribution of this work is to present an experimental evaluation using both real and synthetic datasets using our new framework that tracks not only time but also primitive usage, I/O complexity, and communication cost. In the process, we present improvements for some of the schemes that make them more efficient and/or more secure. To make understanding of these schemes easier for the reader, we present the main ideas behind these schemes, discuss their security definitions and leakage profiles, and provide an analysis of implementation challenges for each one.

- We discuss the security definitions for a number of important schemes and protocols and we contrast their leakage profiles.
- We present the main ideas behind these schemes and protocols and we provide our analysis and implementation challenges for each one. In addition, we present improvements for some of the schemes that make them more efficient and/or more secure.
- Finally, we give an overall picture of the different methods and we make recommendations to practitioners.

1.1 Related work

A number of OPE schemes have been proposed recently including [2, 52, 9, 10, 59, 39, 63, 35, 38, 36, 65, 64, 47, 24, 46]. Popa, Li, and Zeldovich [53] present a nice analysis of these schemes and they are the first to show that using a stateful scheme you can achieve the ideal security guarantees for OPE. We pick two of these schemes (BCLO [9] and FH-OPE [38]) that are the most representative and outperform other schemes.

In addition, there are a number of ORE schemes [11, 19, 45, 17, 16, 12, 29, 25] that have been proposed. We choose

the most practical and most secure of them [19, 45, 17], to include in the comparison. Also, there are some approaches that assume an outsourced setting where the client may have to communicate with the server during query processing [55, 40, 6, 22]. We choose two of these protocols [55, 40] because they are based on OPE and ORE approaches and therefore have similar security models with these schemes. We would like to point out that there are some other methods that can be used to run range queries on encrypted data that use different types of schemes and techniques. See [6] and [48] for an overview of other methods. In this paper we consider two of the protocols proposed in [22] that use Searchable Symmetric Encryption (SSE). Finally, we would like to stress that the schemes and protocols discussed here should be used with care. The schemes provide specific primitives, security guarantees and leakage profiles, and it is up to the practitioner how to use them.

2. SECURITY PERSPECTIVE

Each scheme and protocol we analyze has its own security definition, which captures leakage from a lot to nothing. We attempt to unify these definitions and analyze them under a common framework. We also attempt to assess relative security of these definitions and analyze their leakages.

In this work we mostly consider the snapshot model, where the attacker can observe all the database contents at different time instants. Note that this excludes timing attacks such as measuring encryption time. All security definitions of the schemes and protocols that we discuss here are based on this model. Also, the snapshot attacker is the most common attacker that we face today [6]. The idea is that a hacker or an insider can steal the entire encrypted database and all its contents at some point in time.

Besides snapshot model a stronger setting allows an adversary to track communication and data manipulation. Most notably, attacker can see and analyze communication volume and access pattern in real time. There are ways to protect against such attacker — ORAM against access pattern and differential privacy against communication volume leakage. Although this model is not a primary target of this paper, our benchmark includes a protocol (Section 4.5.2) that is secure in this setting to show the cost of adding such protection.

We wanted to specifically comment on a work of Grubbs et al. [28], which demonstrates a series of attacks against OPE and ORE schemes. The attacks can be very successful, but depend on certain prerequisites. First, all attacks assume the existence of a well-correlated auxiliary dataset. Second, the binomial attack, which works against a “perfectly secure frequency-hiding scheme”, reliably recovers only high-frequency elements. Finally, the attacks are specifically devastating against encrypted strings (e.g. first and last names) as opposed to numerical data, and we also do not recommend using OPE/ORE for strings (see next subsection). One of the conclusions of our work is that security is negatively correlated with performance and it is up to a practitioner to trade off security and performance constraints.

2.1 A note on variable-length inputs

A generic OPE/ORE scheme accepts bit-strings of any length as inputs, and treats them as numbers or processes them bit-by-bit. One might think that supplying raw bytes of variable length (e.g. encoded strings) to OPE/ORE schemes

may naturally work. We warn against it. Such an approach will introduce both performance and security challenges.

From the performance standpoint, OPE/ORE schemes' complexity usually depends on input length at least linearly (see Table 1). 32-bit numbers already introduce a noticeable overhead for some (usually more secure) schemes, and supplying arbitrary-length inputs may worsen performance an order of magnitude.

Security of such a construction will be minimal as most schemes leak some information about a magnitude of a difference, and longer inputs will naturally be treated as larger numbers. Thus, the difference between long and short inputs will be apparent. We refer to the work of Grubbs et al. [28] as they have a practically supported discussion of security consequences of using OPE/ORE with arbitrary strings.

On the other hand, other protocols in our benchmark can usually handle variable-length inputs as long as they fit into a single block for the underlying block cipher.

3. OPE AND ORE SCHEMES

Order-Revealing Encryption scheme is a triple of polynomial-time algorithms KGEN, ENC and CMP. KGEN generates a key of parameterized length (the λ parameter). ENC takes a numerical input (as a bit string) and produces a ciphertext. CMP takes two ciphertexts generated by the scheme and outputs whether the first plaintext was strictly less than the second. Note that being able to check this condition is enough to apply all other comparison operators ($<$, \leq , $=$, \geq , $>$). Also note that an ORE scheme does not include a decryption algorithm, because one can simply append a symmetric encryption of the plaintext to the produced ciphertext and use it for decryption.¹ An Order Preserving Encryption (OPE) scheme is particular case of an ORE scheme where ciphertext is numerical and thus CMP routine is trivial (the numerical order of ciphertexts is the same as underlying plaintexts). OPE may optionally include decryption algorithm, since appending symmetric ciphertext is no longer possible.

Both OPE and ORE schemes by definition allow to totally order the ciphertexts. This is their inherent leakage (by design) and all the OPE/ORE security definitions account for these and possibly additional leakage.

We proceed by describing and analyzing the OPE/ORE schemes we have benchmarked. All plaintexts are assumed to be 32-bit signed integers, or n -bit inputs in complexity analysis. OPE ciphertexts are assumed to be 64-bit signed integers.

From here, we will use the term ORE to refer to both OPE and ORE, unless explicitly stated otherwise. Each scheme has its own subsection where the first part is the construction overview followed by security discussion, and the second part is our theoretical and experimental analysis.

3.1 BCLO OPE

The OPE scheme by Boldyreva et al. [9] was the first OPE scheme that provided formal security guarantees and was used in one of the first database systems that executes queries over encrypted data (CryptDB [54]).

¹ Given the secret key, it is possible to decrypt a ciphertext by doing binary search on the plaintext domain: encrypting known values and comparing them against the target ciphertext, until the target plaintext is found. However, this would require $\mathcal{O}(\log |\mathcal{D}|)$ encryption and comparison operations.

The core principle of their construction is the natural connection between a random order-preserving function and the hypergeometric probability distribution. Authors formalize this principle proving a bijection between the set of all order-preserving functions from a domain of size M to a range of size $N \geq M$ and the set of all possible combinations of M out of N ordered items.

The encryption algorithm works by splitting the domain into two parts according to a value sampled from the hypergeometric distribution (HG) routine and splitting the range in half recursively. When the domain size contains a single element, the corresponding ciphertext is sampled uniformly from the current range.

All pseudo-random decisions are made by an internal PRG (TAPEGEN in [9]). This way not only they ensure that the algorithm is deterministic, but also allow for decryption. The decryption procedure takes the same "path" of splitting domain and range, and when the domain size reaches one, the only left value is the original plaintext.

Security. This scheme is POPF-CCA secure [9], meaning that it is as secure as the underlying ideal object — randomly sampled order-preserving function from a certain domain to a certain range. For practical values of the parameters, Boldyreva, Chenette, and O'Neill [10] showed that the distance between the plaintexts can be approximated to an accuracy of about the square root of the domain size. In other words, approximately, half of the bits (the most significant) are leaked. Grubbs et al. [28] showed that this leakage allows to almost entirely decrypt the ciphertexts (given auxiliary data with a similar distribution) and encrypting strings with this scheme is especially dangerous (see Section 2.1).

Analysis and implementation challenges

Efficient sampling from hypergeometric distribution is a challenge by itself. Authors suggest using a randomized yet exact (not approximate) Fortran algorithm by Kachitvichyanukul and Schmeiser [34]. It should be noted that the algorithm relies on infinite precision floating-point numbers, which most regular frameworks do not have. The security consequences of finite precision computations is actually an open question. The complexity of this randomized algorithm is hard to analyze; however, we empirically verified that its running time is no worse than linear in the input bit length. The authors also suggest a different algorithm for smaller inputs [61].

On average, encryption and decryption algorithms make n calls to HG, which in turn consumes entropy generated by the internal PRG. The entropy, and thus the number of calls to PRG, needed for one HG run is hard to analyze theoretically. However, we derived this number experimentally (see Section 5).

BCLO has been implemented in numerous languages and has been deployed in a number of secure systems. We add C# implementation to the list, and recommend using a library that supports infinite precision floating-point numbers when building the hypergeometric sampler.

3.2 CLWW ORE

The ORE scheme by Chenette et al. [19], which authors call "Practical ORE", is the first efficient ORE implementation based on PRFs.

On encryption, the plaintext is split into n values in the following way. For each bit, a value is this bit concatenated

Scheme	Primitive usage		Cipher size, or state size	Leakage (In addition to inherent total order)
	Encryption	Comparison		
BCLO [9]	n HG	none	$2n$	\approx Top half of the bits
CLWW [19]	n PRF	none	$2n$	Most-significant differing bit
Lewi-Wu [45]	$\frac{2n}{d}$ PRP $\frac{n}{d} (2^d + 1)$ PRF $\frac{n}{d} 2^d$ Hash	$\frac{n}{2d}$ Hash	$\frac{n}{d} (\lambda + n + 2^{d+1}) + \lambda$	Most-significant differing block
CLOZ [16]	n PRF n PPH 1 PRP	n^2 PPH	$n \cdot h$	Equality pattern of most-significant differing bit
FH-OPE [38]	1 Traversal	3 Traversals	$3 \cdot n \cdot N$	Insertion order

Table 1: Primitive usage by OPE / ORE schemes. Ordered by security rank — most secure below. n is the input length in bits, d is a block size for Lewi-Wu scheme, λ is a PRF output size, N is a total data size, **HG** is hyper-geometric distribution sampler, **PPH** is property-preserving hash with h -bit outputs built with bilinear maps and **bolded** are weak points of the schemes.

with all less significant bits. This value is given to a keyed PRF and the result is added to the next more significant bit. This resulting list of n elements is the ciphertext.

The comparison is trivial. The algorithm compares two lists traversing them in-order looking for the case when one value is greater than the other by exactly one. This would mean that the first differing bit is found. If no such index exists, the plaintexts are equal.

Security. A generic ORE security definition was introduced along with the scheme [19]. ORE leakage is more clearly quantified than in OPE. The definition says that the scheme is secure with a leakage $\mathcal{L}(\cdot)$ if there exists an algorithm (simulator) that has access to the leakage function and can generate output indistinguishable from the one generated by the real scheme. This scheme satisfies ORE security definition with the leakage $\mathcal{L}(\cdot)$ of the location and value of the first differing bit of every pair of plaintexts. Note that the most significant differing bit also leaks the approximate distance between two values. For example, if the most significant differing bit is the last bit, then the plaintexts’ difference is one.

Analysis and implementation challenges

On encryption the algorithm makes n calls to PRF and the comparison procedure does not use any cryptographic primitives. Ciphertext is a list of length n , where each element is an output of a PRF modulo 3. The authors claim that the ciphertext’s size is $n \log_2 3$, just 1.6 times larger than the plaintext’s size. While this may be true for large enough n if ternary encoding is used, we found that in practice the ciphertext size is still $2n$. $1.6n$ for 32-bit words is 51.2 bits, which will have to occupy one 64-bit word, or two 32-bit words, therefore resulting in $2n$ anyway.

3.3 Lewi-Wu ORE

Lewi and Wu [45] introduced an improved version of CLWW [19] which leaks strictly less.

The novel idea was to use the “left / right framework” in which two ciphertexts get generated — left and right. The right ciphertexts are semantically secure, so an adversary will learn nothing from them. Comparison is only defined

between the left ciphertext of one plaintext and the right ciphertext of another plaintext.

The approach is to split the plaintext into the blocks of certain number of bits. Next, apply CLWW [19] procedure to blocks. Within one block, the algorithm computes all possible permutations of values, hashes them and adds the result of the comparison between the value and the non-permuted block value. This way the location of the differing bit inside the block is hidden, but the location of the first differing block is revealed.

Comparison recomputes the hashes for each block and verifies that any of them have the property that one is greater than the other by one.

Security. This scheme satisfies the ORE security definition introduced by Chenette et al. [19] with the leakage $\mathcal{L}(\cdot)$ of the location of the first differing *block*. This property allows a practitioner to set performance-security tradeoff by tuning the block size. Left / right framework is particularly useful in a B+ tree since it is possible to store only one (semantically secure) ciphertext in the structure (see Section 4.1).

Analysis and implementation challenges

Let n be the size of input in bits (e.g. 32) and d be the number of bits per block (e.g. 2).

Left encryption loops $\frac{n}{d}$ times making one PRP call and two PRF calls each iteration. Right encryption loops $\frac{n}{d} 2^d$ times making one PRP call, one hash call and two PRF calls each iteration. Comparison makes $\frac{n}{d}$ calls to hash at worst and half of that number on average. Please note that the complexity of right encryption is exponential in the block size. In the Table 1 the PRP usage is linear due to our improvement.

Assume PRF output size is λ . Left ciphertext size is therefore $\frac{n}{d} (\lambda + n)$. Right ciphertext size is $\lambda + \frac{n}{d} 2^{d+1}$. The total ciphertext size is then $\frac{n}{d} (\lambda + n + 2^{d+1}) + \lambda$.

The implementation details of this approach has an interesting security question. Although the authors suggest using 3-rounds Feistel networks [56] for PRP and use it in their implementation, it may not be secure for small input sizes. Feistel networks security depends on the input size [30] — exponential in the inputs size. However, the typical input

for PRP in their scheme is 2–8 bits, thus even exponential number is small.

We have considered multiple PRP implementations to use instead of the Feistel networks. We have found that Knuth shuffle algorithm [42] fits particularly well. It is secure for any input size, and its performance, although degrades with input size, is acceptable for small inputs. Another important aspect of the implementation is that for each block we need to compute the permutation of all the values inside the block. This operation applied many times can be expensive. To address this, we propose to generate a PRP table once for the whole block and then use this table when you need to compute the location of an element of permutation. This can reduce the PRP usage (indeed, we observe a reduction from 80 to 32 calls in our case.) We evaluate this improved approach in our experimental section.

3.4 CLOZ ORE

Cash et al. [16] introduced a new ORE scheme that provably leaks less than any previous scheme. Their construction uses bilinear maps for a new primitive they have defined, which allows to hide the first differing bit.

The idea is to use Chenette et al. [19] construction, but permute the list of PRF outputs. It is not necessary to know the original order of those outputs, as one can simply find a pair where one element is greater than the other by one. This is not enough to reduce leakage, however, since an adversary can count how many elements two ciphertexts have in common.

To address this problem, the authors define a new primitive they call a *property-preserving hash* (PPH).

A PPH as defined in [16] should be instantiated with a predicate (property) on two elements of the domain. PPH TEST will then output 1 if the predicate is true for the original elements, and false otherwise. For the purposes of this scheme, the predicate we are interested in is $y = x + 1$, or testing if the first element is greater than the second by 1.

Informally, PPH is correct if PPH TEST routine is correct, and PPH is secure if an adversary cannot distinguish a real PPH from a fake one without the secret key. Please refer to the original paper [16] for formal correctness and security definitions.

Equipped with the PPH primitive, the authors “hash” the elements of the ciphertexts before outputting them. Due to security of PPH, the adversary would not be able to count how many elements two ciphertexts have in common, thus, would not be able to tell the location of differing bit.

Security. The strong side of the scheme is its security. The scheme leaks $\mathcal{L}(\cdot)$ an *equality pattern* of the most-significant differing bits (satisfying Chenette et al. [19] definition). As defined in [16], the intuition behind equality pattern is that for any triple of plaintexts m_1, m_2, m_3 , it leaks whether m_2 differs from m_1 before m_3 does. We do not know of any attacks against this construction (partially because no implementation exists yet, see next subsection), but it is inherently vulnerable to frequency attacks that apply to all frequency-revealing ORE schemes (see Section 2).

Analysis and implementation challenges

On encryption, the scheme makes n calls to PRF, n calls to PPH HASH and one call to PRP. Comparison is more expensive, however, as the scheme makes n^2 calls to PPH

TEST. Ciphertext size depends on PPH implementation and is equal to n times the PPH hash size.

The scheme has two limitations that make it impractical. The first one is the square number of calls to PPH, which is around 1024 for a single comparison.

The second problem is the PPH itself. Authors suggest a construction based on bilinear maps. Hash of an argument is an element of a group, and the test algorithm is computing a pairing. This operation is very expensive — order of magnitude more expensive than any other primitive we have implemented for other schemes.

We have implemented the scheme in C++ using PBC library [49] to empirically assess schemes’s performance and on our machine (see Section 5), a single comparison takes 1.9 seconds on average. Although we have produced the first (correct and secure) real implementation of this scheme in C++, it is infeasible to use it in the benchmark (it will take years to complete a single run). Therefore, for the purposes of our benchmark, we implemented a “fake” version of PPH — correct, but insecure, which does not use pairings. Consequently, in our analysis we did not benchmark the speed of the scheme, but measured all other data.

3.5 FH-OPE

Frequency-hiding OPE by Kerschbaum [38] is a stateful scheme that hides the frequency of the plaintexts — adversary would not be able to construct a frequency histogram.

This scheme is stateful, which means that the client needs to keep a data structure and update it with every encryption and decryption. The data structure is a binary search tree where the node’s value is the plaintext and node’s position in a tree is the ciphertext. For example, consider the range $[1, 128]$. The first *any* plaintext, let it be 6, will be the root, and thus the cipher is 64. Then any plaintext smaller than the root 6, say 3, will become the left child of the root, and will produce the ciphertext 32.

To encrypt a value, the algorithm traverses the tree until it finds a spot for the new plaintext, or finds the same plaintext. If the same plaintext is found, in order to hide the frequency, the algorithm tosses a coin and goes left or right depending on the outcome up to the leaf. This way, the invariant of the tree — intervals of the same plaintexts do not overlap — is maintained. The ciphertext generated from the new node’s position is returned.

The property that every duplicate plaintext will have a new pseudo-random ciphertext makes the scheme randomized. Therefore, the comparison algorithm is more complicated than in the regular deterministic OPE.

To properly compare ciphertexts, the algorithm needs to know the boundaries — the minimum and maximum ciphertexts for a particular plaintext. The client is responsible for traversing the tree to find the plaintext for the ciphertext and then minimum and maximum ciphertext values. Having these values, the comparison is trivial — equality is a check that the value is within the boundaries, and other comparison operators are similar.

Authors have designed a number of heuristics to minimize the state size, however, these are mostly about compacting the tree and the result highly depends on the tree content. In our analysis, we consider the worst case performance without the use of heuristics. In our experimental evaluation, however, we did implement compaction.

Protocol	I/O requests		Leakage	Communication (result excluded)	
	Construction	Query		Construction	Query
B+ tree with ORE	$\log_B \frac{N}{B}$	$\log_B \frac{N}{B} + \frac{r}{B}$	Same as ORE	1	1
Kerschbaum [40]	$\frac{N}{B}$	$\log_2 \frac{N}{B} + \frac{r}{B}$	Total order	$\log_2 N$	$\log_2 N$
POPE [55] cold	1	N/B	Fully hiding	1	N
POPE [55] warm		$\log_L \frac{N}{B} + \frac{r}{B}$	Partial order		$\log_L N$
Logarithmic-BRC [21]	—	r	Same as SSE	—	$\log_2 N$
ORAM	$\log^2 \frac{N}{B}$	$\log_2 \frac{N}{B} \left(\log_B \frac{N}{B} + \frac{r}{B} \right)$	Fully hiding (access pattern)	$\log^2 \frac{N}{B}$	$\log^2 \frac{N}{B}$

Table 2: Performance of protocols. Ordered by security rank — most secure below. N is a total data size, B is an I/O page size, L is a POPE tree branching factor, r is the result size in records and **bolded** are weak points of the protocols. All values are in \mathcal{O} notation.

Security. The security of the scheme relies on the large range size to domain size ratio. Authors recommend at least 6 times longer ciphertexts than the plaintexts in bit-length, which means ciphertexts should be 192-bit numbers that are not commonly supported. It is possible to operate over arbitrary-length numbers, but the performance overhead would be substantial. We did a quick micro-benchmark in C# and the overhead of using `BigInteger` is 15–20 times for basic arithmetic operations.

This scheme satisfies IND-FAOCPA definition (introduced along with the scheme [38]), meaning that it does not leak the equality pattern or relative distance between the plaintexts. This definition has been criticized in [50], who claim that the definition is imprecise and propose the enhanced definition along with a small change to construction to satisfy this new definition. Both schemes leak the insertion order, because it affects the tree structure. We do not know of any attacks against this leakage, but it does not mean they cannot exist. Grubbs et al. [28] describe an attack against this scheme (binomial attack), but it applies to any perfectly secure (leaking only total order) frequency-hiding OPE.

Analysis and implementation challenges

If the binary tree grows in only one direction, at some point it will be impossible to generate another ciphertext. In this case, the tree has to be re-balanced. This procedure will invalidate all ciphertexts already generated. However, the client can still generate new ones by manipulating its state. This property makes the scheme difficult to use in some protocols since they usually rely on the ciphertexts on the server being always valid. The authors explicitly mention that the scheme works under the assumption of uniform input. However, the re-balancing will be caused by insertion of just $r+1$ sequential data points for r being the range size in bits, which means that 65 consecutive input elements is enough for 64-bit integer range.

The scheme makes one tree traversal on encryption and decryption. Comparison is trickier as it requires one traversal to get the plaintext, and two traversals for minimum and maximum ciphertexts. We understand that it is possible to get these values in fewer than three traversals, but we did not optimize the scheme for the analysis and evaluation.

For practitioners we note that the stateful nature of the scheme implies that the client storage is no longer negligible

as the state grows proportionally to the number of encryptions. We also note that implementing compaction extensions will affect code complexity and performance. Finally, we stress again that some inputs — namely all non-uniform inputs — can break the scheme by causing all ciphertexts to be invalid. It is up to the users of the scheme to ensure uniformity of the input, which poses serious restrictions on the scheme usage.

4. SECURE RANGE QUERY PROTOCOLS

We proceed by describing and analyzing the range query protocols we have chosen. For the purposes of this paper, a secure search protocol is defined as a client-server communication involving construction and search stages. Communication occurs between a client, who owns some sensitive data, and an honest server, who securely stores it. In construction stage, a client sends the server the encrypted data points (index-value tuples) and the server stores them in some internal data structure. In search stage, a client asks the server for a range (usually specifying it with encrypted endpoints) and the server returns a set of encrypted records matching the query. Note that the server may interact with the client during both stages (e.g. ask the client to sort a small list of ciphertexts). Also note that we do not allow batch insertions as it would limit the use cases (e.g. client may require interactive one-by-one insertions).

The first protocol is a family of constructions where a data structure (B+ tree in this case) uses ORE schemes internally. Following are the alternative solutions with varying performance and security profiles, not relying on ORE. Final subsection introduces two baseline solutions we will use in the benchmark — best performance and maximal security.

4.1 Search protocol from ORE

So far we have analyzed OPE and ORE schemes without much context. One of the best uses of an ORE is within a secure protocol. In this section we provide a construction of a search protocol built with a B+ tree working on top of an ORE scheme and analyze its security and performance.

The general idea is to consider some data structure that is optimized for range queries, and to modify it to change all comparison operators to ORE scheme’s `CMP` calls. This way the data structure can operate only on ciphertexts. Performance overhead would be that of using the ORE scheme’s

CMP routine instead of a plain comparison. Space overhead would be that of storing ciphertexts instead of plaintexts.

In this paper, we have implemented B+ tree [4] (with a proper deletion algorithm [33]) as a data structure.

B+ tree is a good choice for underlying data structure because all its operations — insertions, updates, searches and deletions — require nothing beyond ordering of its elements. Moreover, these operations require comparisons of arguments to already inserted elements, thus enabling the use of “left / right framework”.

For protocols, we also analyze the I/O performance and the communication cost. In particular, we are interested in the expected number of I/O requests the server would have made to the secondary storage, and the number and size of messages parties would have exchanged.

The relative performance of the B+ tree depends only on the page capacity (ciphertexts are larger than the plaintexts and therefore the B+ tree will have smaller branching factor). Therefore, the query complexity is:

$$\mathcal{O}(\log_B(N/B) + r/B)$$

where B is the number of records (ciphertexts) in a block, N is the number of records (ciphertexts) in the tree and r is the number of records (ciphertexts) in the result (none for insertions).

Communication volume of the protocol is relatively small. For insertions, the client transfers one ciphertext in one message. For queries, the client transfers two ciphertexts in one message, and gets one message of result size back.

Security. The leakage of this protocol consists of leakage of the underlying ORE scheme plus whatever information about insertion order is available in the B+ tree. Please note that Lewi-Wu [45] ORE is particularly well-suited in this construction with its left / right framework. In this case, the ORE leakage becomes only the total order and the security of the protocol is comparable with other non-ORE constructions.

4.2 Kerschbaum-Tueno

Kerschbaum and Tueno [40] proposed a new data structure, which satisfies their own definitions of security (IND-CPA-DS) and efficiency (search operation has poly-logarithmic running time and linear space complexity).

In short, the idea is to maintain a (circular) array of symmetrically encrypted ciphertexts in order. On insertion, the array is rotated around a uniformly sampled offset to hide the location of the smallest element. Client interactively performs a binary search requesting an element, decrypting it and deciding which way to go.

Security. Authors prove that this construction is IND-CPA-DS secure (definition introduced in the same paper [40]). The definition assumes an array data structure and therefore serves specifically this construction (as opposed to being generic). It provably hides the frequency due to semantic encryption and hides the location of the first element due to random rotations. Leakage-wise this construction is strictly better than B+ tree with ORE — they both leak total order, but [40] hides distance information and smallest / largest elements. Specifically, for all pairs of consecutive elements e_i and e_{i+1} it is revealed that $e_{i+1} \geq e_i$ except for one pair of smallest and largest element in the set.

4.2.1 Analysis and implementation challenges

Insertions are I/O-heavy because they involve rotation of the whole data structure. All records will be read and written, thus the complexity is $\mathcal{O}(N/B)$. Searches are faster since they involve logarithmic number of blocks. The first few blocks can be cached and the last substantial number of requests during the binary search will target a small number of blocks. The complexity is then $\mathcal{O}(\log_2 N/B)$.

Communication volume is small as well. Insertion requires $\log_2 N$ messages from each side. Searches require double that number because separate protocol is run for both endpoints. Inherently, the response is sent in a single message.

The data structure is linear in size, and the client storage is always small. Sizes of messages are also small as only a single ciphertext is usually transferred.

For practitioners we have a few points. The construction in the original paper [40] contains a typo as m and m' must be swapped in the insertion algorithm. Also, we have found some rare edge cases; when duplicate elements span over the modulo, the algorithm may not return the correct answer. Both inconsistencies can be fixed however. This protocol is not optimized for I/O operations for insertions, and thus would be better suited for main memory datasets.

4.3 POPE

Roche et al. [55] presented a protocol, direct improvement over mOPE [53], which is particularly suitable for large number of insertions and small number of queries. The construction is heavily based on buffer trees [3] to support fast insertion and lazy sorting.

The idea is to maintain a POPE tree on the server and have the client manipulate that tree. POPE tree is similar to B-tree, in that the nodes have multiple children and nodes are sorted on each level. Each node has an ordered list of *labels* of size L and an unbounded unsorted set of encrypted data called *buffer*. Parameter L controls the list size, the leaf’s buffer size, and the size of client’s working set. The insertion procedure simply adds an encrypted piece of data to the root’s buffer.

The query procedure is more complex. To answer a query, the server interacts with the client to split the tree according to the query endpoints. On a high level, for each endpoint the buffers are cleared (content pushed down to leaves), and nodes in the paths are split. After that, answering a query means sending all ciphertexts in all buffers between the two endpoint leaves.

The authors provide cost analysis of their construction. Insertions are always cheap — one round is needed to send an element to the server. Search operations are expected to require $\mathcal{O}(\log_L n)$ rounds. It must be noted that the first queries will require many more rounds, since large buffers must be sorted.

Security. This construction satisfies the security definition of frequency-hiding partial order preserving (FH-POP) protocol (introduced in the paper [55]). According to Theorem 3 from [55], after n insertions and m query operations with local storage of size L , where $mL \in o(n)$, the POPE scheme is a frequency-hiding partial-order-preserving with $\Omega\left(\frac{n^2}{mL \log_L n} - n\right)$ incomparable pairs of elements. Simply put, the construction leaks pairwise order of a *bounded* number of elements. Aside from this, the construction provably hides the frequency (i.e. equality) of the elements.

4.3.1 Analysis and implementation challenges

Protocol	I/O requests (result included)		Communication per operation (result excluded)			
			Volume (messages)		Size (bytes)	
	Construction	Query	Construction	Query	Construction	Query
B+ tree w. ORE	3	44	2	2	177	342
Kerschbaum [40]	494	7	40	86	671	1453
POPE [55] cold	1	2175	2	497722	32	9056644
POPE [55] warm		300		914		43331
Logarithmic-BRC [21]	—	40	1	2	—	391
ORAM	31	185	143	490	18254	62662

Table 3: Simulation result for protocols’ performance values

In our analysis we count each request-response communication as a round. This is different from [55] where they use *streaming* a number of elements as a single round. The rationale for our approach is that if we allow persistent channels additionally to messages, then any protocol can open a channel for each operation. Thus, we do not allow channels for all protocols in our analysis.

Also, as noted by authors, if $L = n^\epsilon$ for $0 < \epsilon < 1$, then the amortized costs become $\mathcal{O}(1)$. While this is true, in our analysis the choice of L depends on the storage volume block size for I/O optimizations, instead of the client’s volatile storage capacity. Thus, the costs remain logarithmic.

Insertion bandwidth is constant and small — one ciphertext is transferred. Search bandwidth depends heavily on the current state of the tree. When the tree is completely unsorted (the first query), all elements of the tree will be transferred to split the large root, then possibly internal node will have to be split requiring sending of $\frac{N}{L}$ elements, and so on, thus $\mathcal{O}(N + r)$. When the tree is completely sorted (after a large number of uniform queries), the bandwidth will be similar to that of a standard B+ tree — $\mathcal{O}(L \log_L N + r)$. The average case is hard to compute; however, authors prove an upper bound on bandwidth after n insertions and m queries — $\mathcal{O}(mL \log_L n + n \log Lm + n \log L(\ln n))$.

POPE tree is not optimized for I/O the way B-tree is. Insertion requires a single I/O request for the block where the server appends the element. Search complexity is more complex to analyze as is bandwidth complexity. In the worst-case (first query), all blocks need to be accessed $\mathcal{O}(\frac{N}{B} + \frac{r}{B})$. In the best-case all nodes occupy exactly one block and I/O complexity is the same as with B+ tree $\mathcal{O}(\log_L \frac{N}{B} + \frac{r}{B})$. The average case is in between and matters get worse as the node is not guaranteed to occupy a single block due to arbitrary sized buffers.

Client’s persistent storage is negligibly small — it stores the encryption key. Volatile storage is bounded by L .

For practitioners we present a number of things to consider. Buffer within one node is unsorted, so in the worst-case, L -sized chunks remain unordered. Due to this property, the query result may contain up to $2(L - 1)$ extra entries, which the client will have to discard from the response.

The first query after a large number of insertions will result in client sorting the whole N elements, and thus, POPE has different performance for cold and warm start. Also, even to navigate an already structured tree, the server has to send to the client the whole L elements and ask where to go on all levels.

Furthermore, [55] does not stress the fact that after alternating insertions and queries, it may happen that some intermediate buffers are not empty, thus returning buffers between endpoints must include intermediate buffers as well. The consequence is that the whole subtree is traversed between paths to endpoints, unlike the B+ tree case when only leaves are involved.

Finally, POPE tree is not optimized for I/O operations. Even if L is chosen so that the node fits in the block, only leaves and only after some number of searches will optimally fit in blocks. Arbitrary sized buffers of intermediate nodes and the lack of underflow requirement do not allow for I/O optimization.

4.4 Logarithmic-BRC using SSE

Demertzis et al. [21] introduced a novel protocol called “Logarithmic-BRC” whose I/O complexity depends only on the result size, regardless of the database size. The core primitive for their construction a Symmetric Searchable Encryption (SSE) scheme. An SSE scheme is a server-client protocol in which the server stores a specially encrypted keywords-to-documents map, and a client can query documents with keywords while the server learns neither keywords nor the documents (although there is access pattern leakage). Note that the map stores short document identifiers instead of actual documents, and we will use the term “documents” to mean “document identifiers” in this section.

The construction treats record values as “documents” and index ranges as “keywords” so that records can be retrieved by the ranges that include them. Specifically, a client builds a virtual binary tree over the domain of indices and assigns each record a set of keywords, which is the path from that record to the root. This way, the root keyword is associated with all documents and the leaf keyword is associated with only one record.

Upon query, a client computes a cover — a set of nodes whose sub-trees cover the requested range. A client sends these keywords to the SSE server, which returns encrypted documents — result values. Of the several covering techniques suggested in the protocol [21] we have chosen the Best Range Cover (BRC), because it results in fewest nodes and does not return false-positives. Kiayias et al. [41] have proven that the worst-case number of nodes for domain of size N is $\mathcal{O}(\log N)$ and presented an efficient BRC algorithm.

Security. In a snapshot setting, this construction’s security is that of SSE. We have used [15] and [14] SSE schemes and

their leakage in a snapshot setting is the database size and at most some initialization parameters. Security of modern SSE is high enough to call them *fully hiding* in our setting. We remind that we consider only a snapshot model. Additional access pattern leakage comes up during queries; exact implications of this leakage remain an open research problem.

4.4.1 Analysis and implementation challenges

Communication involves a client sending at worst $\log_2 N$ keywords and server responding with the exact result.

For each keyword in the query set, server will query SSE, which will return r documents. Therefore, server’s I/O complexity is that of SSE.

Authors have used [15] SSE scheme in their implementation, but we have found it unacceptably slow in terms of I/O. Instead we have implemented an improved scheme [14], which directly addresses I/O optimization. We have implemented both schemes and have run the protocol with both of them and can confirm the two orders of magnitude improvement in I/O.

Both SSE schemes’ I/O complexity is linear with the result size r . Cash et al. [14] scheme makes at most one I/O per result document in the worst-case and suggests extensions to significantly improve I/O complexity. We have implemented their **pack** extension, which packs documents in blocks to fit the I/O pages.

Logarithmic-BRC is perfectly scalable as its performance does not depend on total data size and only degrades with the result size. Storage overhead, however, is significant. Each record is associated with the whole path in the binary tree — $\log_2 N$ nodes (keywords). The storage complexity is therefore $\mathcal{O}(N \log N)$, and the overhead is then a factor of $\log N$.

Updates, while addressed in the original protocol, are impractical. Authors suggest using bulk-loading for updates, maintaining merge trees and requiring the client to do a merge once in a while. I/O complexity of such approach is unclear. In our implementation we perform the construction stage only in batch mode. We also emphasize that the update routine was not implemented for evaluation in the original paper.

4.5 The two extremes

To put the aforementioned protocols in a context we introduce the baselines — an efficient and insecure construction we will refer to as *no encryption* and maximal security protocol we refer to as *ORAM*.

4.5.1 No encryption

This protocol is a regular B+ tree [4] without any ORE in it. It is the construction one can expect to see in almost any general-purpose database.

In terms of security it provides no guarantees — all data is in the clear. In terms of efficiency it is optimal. B+ tree data structure is optimal in I/O usage, indices inside nodes are smallest possible (integers) and there is no overhead in comparing elements inside the nodes as opposed to working with ORE ciphers. We present this protocol as a baseline solution in terms of efficiency over security.

4.5.2 ORAM

Oblivious RAM (ORAM) is a construction that additionally to semantic security of a snapshot setting (see Section 2) provably hides the access pattern — a sequence of reads and writes to particular memory locations. With ORAM an adversary would not be able to recognize a series of accesses to the same location and will not differentiate reads versus writes. ORAM was introduced by Goldreich and Ostrovsky [26] who also proved its lower bound (strengthened in [44]) — logarithmic overhead per request. A number of efficient ORAM constructions were designed (see [18] for a good survey) and we use the state-of-the-art construction, PathORAM [58].

A generic ORAM server responds to read and write requests for a particular address. In our baseline protocol we store B+ tree nodes in ORAM. A client works with the tree as it normally would except each time it needs to access a node, it communicates with ORAM.

In terms of security this protocol fully hides in the snapshot model and provably hides access pattern. We note that one can improve security even further by adding noise to the result obscuring communication volume. It is possible to use differential privacy to provably hide the volume, but it is outside of scope of this work. We also note that a practitioner can use a similar protocol with ORAM replaced with a trivial data store and have the tree nodes encrypted. It would be fully hiding in a snapshot setting, but we prefer the baseline that covers more than only the snapshot model.

In terms of performance this construction incurs some noticeable overhead. Regardless of specific ORAM being used, each access incurs at least logarithmic overhead according to lower bounds [26]. Combined with logarithmic complexity of the B+ tree itself, the complexity, both I/O and communication, is $\mathcal{O}(\log^2 N)$. Particular values depend on the specific construction. We found that PathORAM has good I/O performance, because its internal tree structure translates into good cache affinity.

We present this protocol as a baseline solution in terms of security over efficiency. We have not implemented standalone PathORAM, but rather a simulator which correctly reports I/O, communication and primitive usage. Surprisingly, we found that ORAM protocol’s overhead, although higher than in ORE-based protocols, is in-line with the most secure protocols in our benchmark.

5. EVALUATION

All experiments were conducted on a single machine. We use macOS 10.13.6 with 8-Core 3.2GHz Intel Xeon W processor, 32 GB DDR4 ECC main memory and 1 TB SSD disk. The main code is written in C# and runs on .NET Core 2.1.3.

Interactive website

Additionally to making our source code, compiled binaries and Docker images available, we want to let researchers interactively run small-sized simulations. We host a website [8] where one can select a protocol (including baselines, CLOZ and both SSE schemes), cache size and policy and I/O page parameter; supply one’s own data and query sets, and run the simulations. Simulations are run one at a time and usually complete within seconds. The user is then able to view the result — tables, plots, values and raw JSON, which we used to build plots for this paper. Input size on the website is limited for practical purposes and users are

encouraged to run arbitrary-size simulations using our binaries or Docker images.

5.1 Implementation

We have implemented most of the primitives, data structures, and constructions ourselves. For some primitives and all schemes we provided the first open-sourced cross-platform C# implementation. We note that neither primitives, nor schemes are production-ready; however, we believe they can be used in research projects and prototypes. We also emphasize that the B+ tree implementation we are using, although our own with instrumentation in it, is not custom in any way, but rather standard as defined in the original paper [4] with deletion algorithm by [33].

This software project is documented and tested (over 97% coverage). All code including primitives, data structures, schemes, protocols, simulation logic, benchmarks, build scripts and tests is published on GitHub [7] under CC BY-NC 4.0 license. Additionally, we have published parts of the projects as stand-alone .NET Core (nuGet) packages, and we host a web-server where users can run simulations for small inputs (Section 5).

5.1.1 Primitives

All schemes and protocols use the same primitives most of which we implemented ourselves. All primitives rely on the default .NET Core AES implementation. .NET Core uses platform-specific implementation of AES, thus leverages AES-NI instruction. In our project all keys' sizes are 128 bits, as is AES block size.

The most used primitive is a pseudo-random generator (PRG). We implemented AES-based PRG which uses AES in CTR mode and caches unused entropy. This way we can supply seed as large as 128 bits (AES key). PRG generates entropy in 128 bits chunks, and if only a fraction of the entropy is requested, the residue is carried over to the next call. When converting entropy to integers, we use techniques to reduce bias (e.g. discarding the entropy in some cases).

The second most used primitive is a PRF. We implemented it using AES without initialization vector in ECB mode. Since we used a single block, this approach is still secure. For symmetric encryption we use AES with initialization vector in CBC mode. For hash we use default .NET Core SHA2 implementation. If keyed hash is used, we push the input through a PRF before supplying it to SHA2.

For pseudo-random permutation (PRP) we have two implementations. We implemented unbalanced Feistel networks [56] using our PRF and .NET Core `BitArray` class. We have a regular (3 round) and strong (4 rounds) version. The second implementation of PRP is for small inputs. We generate a permutation table using Knuth shuffle [42] and cache it for the next call. This implementation is more secure for small inputs and is optimized if the whole permutation is needed. Both schemes (Lewi-Wu and CLOZ) that use PRP, use the Knuth shuffle implementation because of the small inputs.

BCLO [9] relies on special primitives. We implemented LF-PRF (TapeGen in [9]) using our implementations of PRF and PRG as suggested in [9]. For implementation of hypergeometric sampler we used algorithm [34] and code from GitHub [54].

5.1.2 Schemes and protocols

Scheme	Encryption	Comparison	Size (bits)
BCLO [9]	41 HG	none	64
CLWW [19]	32 PRF	none	64
Lewi-Wu [45]	32 PRP 160 PRF 64 Hash	9 Hash	2816
CLOZ [16]	32 PRF 32 PPH 1 PRP	1046 PPH	4096
FH-OPE [38]	1 Traversal	1 Traversal	86842

Table 4: Simulation result for ORE schemes primitive usage

We implemented schemes and protocols precisely as in the original papers. When we found problems or improvements, we put those in implementation challenges notes, but did not alter the original designs in our code, unless explicitly stated. Each ORE scheme implements a C# interface; thus our own implementation of B+ tree operates on a generic ORE. For *no encryption* baseline, we have a stub implementation of the interface, which has identity functions for encryption and decryption. It is important to note that all schemes and protocols use exclusively our implementations of primitives. Thus we rule out the possible bias of one primitive implementation being faster than the other.

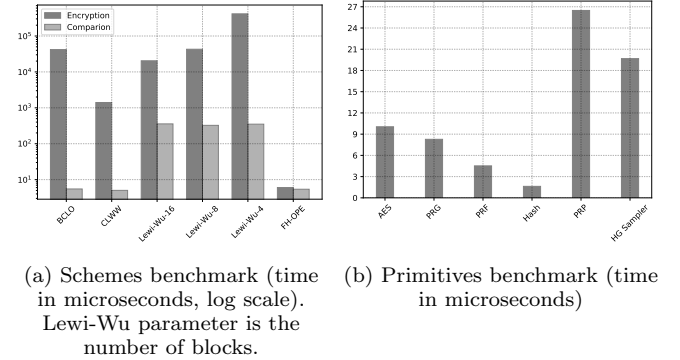


Figure 1: Schemes and primitives benchmarks

5.1.3 Simulations

We have four types of simulations.

Protocol simulation runs both protocol stages — construction and search — on supplied data for all protocols including all schemes coupled with B+ tree. In this simulation we measure the primitive usage, number of ORE scheme operations (when applies), communication volume and size, and the number of I/O requests. We intentionally do not measure elapsed time, since it would be extremely inaccurate in this setting — simulation and measurement routines take substantial fraction of time.

Scheme simulation runs all five ORE schemes and tracks only the primitive usage.

The scheme benchmark, however, is designed to track time. We use Benchmark.NET [1] to ensure that reported time is accurate. This tool handles things like cold / warm start, elevating process' priority, and performing enough

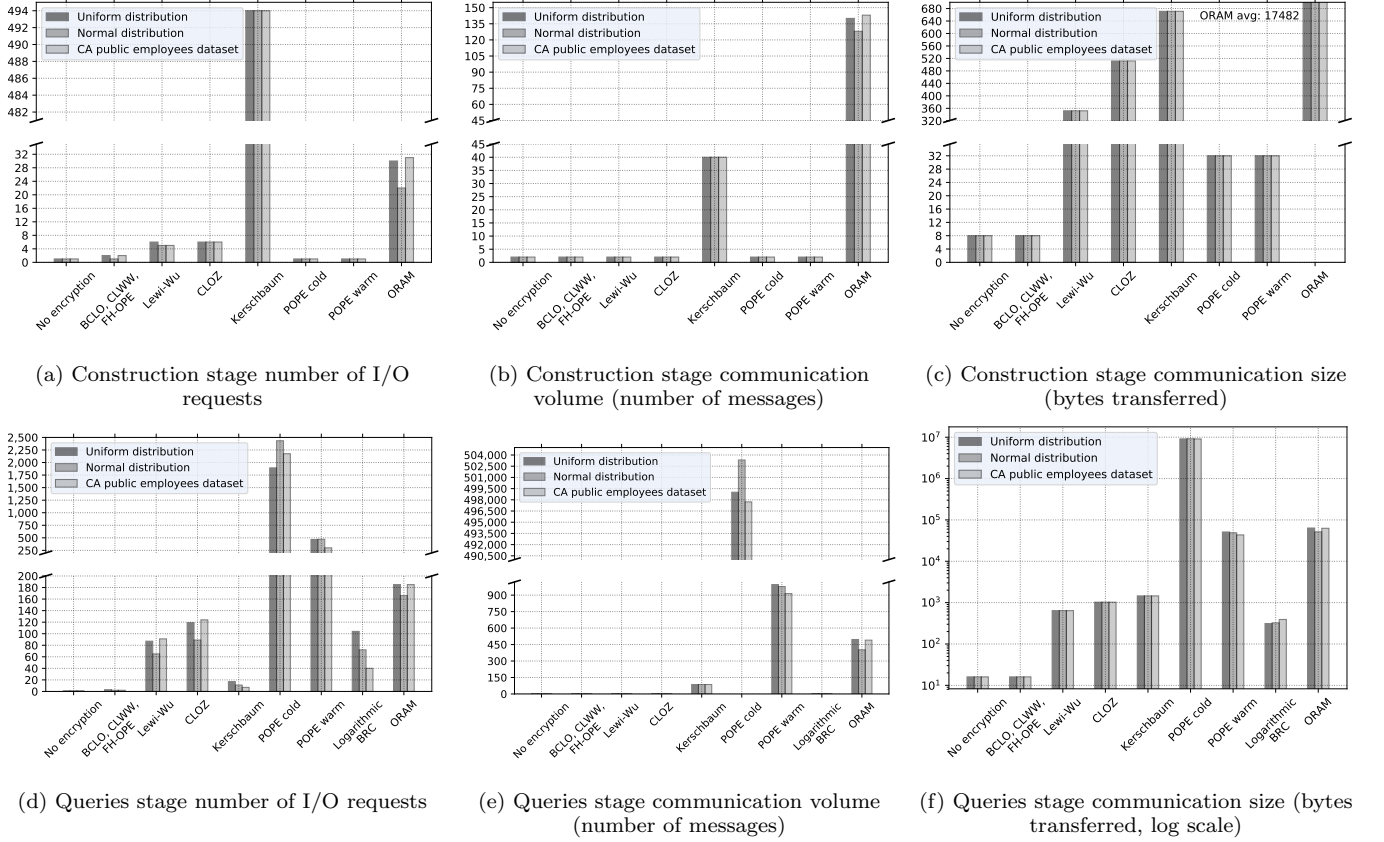


Figure 2: Performance values for different data distributions

runs to draw statistically sound conclusions. This benchmark reports elapsed time up to nanoseconds for all four schemes (excluding CLOZ) and their variants.

Finally, primitive benchmark uses the same tool, but benchmarks the primitives. We use it to compare different implementations of primitives (e.g. Feistel PRP vs pre-generated permutation) and to have basis to approximate scheme and protocol time consumption based on primitive usage.

Simulation and benchmark routines are documented, tested and published as the rest of the software.

5.2 Setup

For our simulations, we have used three datasets. Two synthetic distributions, that are uniform (range is 500) and normal (with mean 0 and standard deviation 10). The real datasets is California public employees salaries [60] (“total pay and benefits” column). Synthetic datasets and subsets of the real dataset are generated pseudo-randomly.

5.3 Results

5.3.1 Primitive usage by schemes

In Table 4 we show the simulation-derived values of each OPE and ORE scheme’s primitive usage. Each scheme is given 1000 data points of each dataset. First, scheme encrypts each data point, then decrypts each ciphertext and then performs five comparisons (all possible types) pairwise. This micro-simulation is repeated 100 times. Resulting values for primitive usage are averaged for each scheme. State


and ciphertext sizes are calculated after each operation and the values are averaged.

The FH-OPE number of traversals for comparison is one since endpoints are found once and then five comparisons are made for the same two ciphertexts. The FH-OPE state size is smaller than expected due to compactations described in [38]. Ciphertext size of CLOZ assumes that the output size of PPH hash is 128 bits. Please note that the simulated values are consistent with the theoretical calculations.


5.3.2 Benchmarks of schemes and primitives

Using the Benchmark.NET tool [1], we have accurately tracked the performance of the schemes and primitives running of different parameters (see Figure 1). On each run, ORE schemes were supplied 100 numbers and they encrypted them, compared each one with the next one (all five comparison operators) and decrypted them. Primitives were given randomly generated byte inputs and keys of different sizes (e.g. PRP of 2 to 32 bits). Benchmark.NET decides how many times to run the routine to get statistically sound results. For example, large variance results in more runs. To improve the accuracy, each run is compiled in release mode as separate project and runs in a separate process with high-priority.


Please note the logarithmic scale of schemes’ performances. FH-OPE is fast since it does not perform CPU-heavy operations and works in main memory. CLWW is the fastest stateless scheme. Its comparison is so simple that it works almost as fast as regular integer comparison. Lewi-Wu per-




(a) Construction stage number of I/O requests




(b) Construction stage number of messages



(a) Queries stage number of I/O requests for different query sizes



(c) Queries stage number of I/O requests



(d) Queries stage number of messages


Figure 3: Protocol scalability

formance degrades exponentially with the increase of block size mainly due to exponential number of PRF executions and the performance of PRP degrading exponentially. Note also that Lewi-Wu comparison takes noticeable time due to Hash primitive usage.

In primitives benchmark it is clear that most primitives use AES under the hood. PRG and PRF take less than AES because they do not include initialization vector generation needed for symmetric encryption. PRP is implemented as a Knuth shuffle [42] and its complexity is exponential in input bit length. Input size of 2 bits is shown on Figure 1. Both PRG and PRP make use of internal cache in our implementations. PRG does not discard the entropy generated by AES cycle, so one AES cycle can supply four 32-bit integers. PRP generates the permutation table once and does not regenerate it if the same key and number of bits are supplied. Hypergeometric sampler uses PRG internally and therefore its performance depends heavily on the PRG implementation. From this plot it is clear why a linear number of sampler usage in BCLO may be better than exponential number of PRP usage in Lewi-Wu.

5.3.3 Protocols

In this experiment we have run each protocol with each of the three datasets. Dataset sizes are 247000 (bounded by California Employees dataset size) and the number of queries is 1000. Queries are generated uniformly at random with a fixed range — 0.5% of data range. The cache size is fixed to 128 blocks, and the B+ tree branching factor as well as block sizes for other protocols are set such that the page



(b) Number of I/O requests over time (number of executed queries).

Figure 4: Number of I/O requests for different queries

size is 4 kilobytes. It effectively means that BCLO, CLWW and FH-OPE branching factors are 512, Lewi-Wu gets 11, CLOZ gets 8 and *no encryption* gets 1024. Kerschbaum with POPE get 256 elements per page, logarithmic-BRC gets 128 and ORAM gets 2 elements per page. The values we are measuring are the number of I/O operations, communication volume, and size for both construction and query stages.

See Table 3 for the snapshot for particular distribution (CA employees). Figure 2 shows all values we tracked for all protocols and distributions. Values for ORE based protocols are averaged. Being “cold” in our simulations means executing the first query and being warm means the first query has been previously executed. This difference makes sense only for POPE as its first query incurs disproportionately large overhead by design.

Note that all ORE based protocols behave the same except when ciphertext size matters. Thus, since BCLO, CLWW and FH-OPE have the same ciphertext size, they create B+ trees with the same page capacity and have the same number of I/Os for different operations. Lewi-Wu and CLOZ schemes have relatively large ciphertexts and thus induce larger traffic (see Subfigure 2c) and smaller B+ tree branching factor resulting in greater number of I/O requests (see Subfigure 2d).

Kerschbaum protocol requires high number of I/O requests during construction since it needs to insert an element into the arbitrary place in an array and rotate the data structure on a disk. It also induces large communica-

tion volume since the insertion is interactive unlike in other protocols.

POPE suffers huge penalty on the first query (see Subfigures 2d, 2e and 2f) since it reads and sends all blocks to the client for sorting. POPE performance improves as more queries are executed.

Logarithmic-BRC does not support interactive insertions and thus its construction stage is not benchmarked. Otherwise it is the most performant of all non-ORE protocols. Note, however, that its performance depends on the result size, not data size.

As expected, ORAM performs worse than the ORE-based protocols, but its performance is in-line with the non-ORE protocols. It may seem that ORAM does especially bad in construction communication (Subfigures 2e, 2f), but it is only because POPE has a shortcut in construction. This “debt” is being paid off during queries (Subfigure 2f).

Note that the values do not vary a lot among different data distributions except for I/O requests. I/O performance depends on the result size for queries, and is therefore more sensitive to data distribution.

Also note that using an ORE scheme in B+ tree does not add any substantial I/O overhead (see “No encryption”).

On Figure 4a it is clear that query performance does not depend substantially on the query size, except for Logarithmic-BRC, for which the relation is linear. ORE schemes with large ciphertext sizes are little sensitive to a query size since the number of blocks needed to answer the query increases.

Figure 3 shows Table 2 asymptotic values. The simulation was run for uniform dataset is 247000 records (hundred percent), 1000 queries, 0.5% query range and 128 blocks cache size. Kerschbaum construction I/Os and cold POPE query values grow linearly with inputs, while the other protocols grow logarithmically, square-logarithmically, or do not grow.

On Figure 4b protocols’ performance over time is shown. Evidently, POPE is the only protocol where cold vs warm makes a difference.

6. REMARKS AND CONCLUSION

Having done theoretical and practical evaluations of the protocols, we have found that primitive usage is a much better performance measure than the plain time measurements. We have also found that I/O optimization is a vital characteristic of a protocol and cannot be neglected. When it comes to practical use, the observed time of a query execution is a mix of a number of factors and I/O requests can slow the system down dramatically.

ORE-based B+ tree protocol is provably I/O optimal and can potentially be extended by using another data structure with ORE. Its security / performance trade off is tunable by choosing and parametrizing the underlying ORE scheme. Each scheme we considered has its own unique advantages and drawbacks. BCLO [9] is the least secure scheme in the benchmark, but is stateless and produces numerical ciphertexts, so it may be used in the databases without any modifications. Frequency-hiding OPE [38] also has this property, hides the frequency of the ciphertexts, but is stateful and requires uniformity of the input. Lewi-Wu [45] is easily customizable in terms of tuning performance to security ratio, and it offers the security benefits of left / right framework — particularly useful for B+ tree. CLWW [19] provides weaker

security guarantees but is the fastest scheme in the benchmark.

Kerschbaum protocol [40] offers semantically secure ciphertexts, hiding the location of the smallest and largest of them, and has a simple implementation. The protocol is well-suited for bulk insertions and scales well.

POPE [55] offers a “deferred” B+ tree implementation. By deferring the sorting of its ciphertexts, POPE remains more secure for the small number of queries. POPE has the fastest insertion routine and does not reveal the order of most of its ciphertexts. It will be more performant for the systems where there are a lot more insertions than queries. We would also recommend to “warm up” the structure to avoid a substantial delay upon the first query.

Logarithmic-BRC is a perfect choice for huge datasets where query result size is limited. It is the only protocol with substantial space overhead, but it offers scalability and perfect (in a snapshot setting) security.

ORAM has shown the most interesting result. Its performance is not only adequate, but also in-line with the other even less secure protocols. With this empirical result, we expect more interest in ORAM research, possibly discovering tighter bounds, faster constructions and efficient ways to use the schemes. On the other hand, this construction’s performance is in some sense an upper bound on performance of less secure (access pattern revealing) protocols, as practitioner will choose ORAM over both less secure and less performant solutions.

We found our framework to be a powerful tool for analyzing the protocols. We encourage protocol developers to contribute their implementations and run the corresponding simulations.

An important future work is to understand better the meaning of the different leakage profiles and their implications. Furthermore, another direction is to try to improve the performance of the most secure schemes (e.g. [16]).

7. ACKNOWLEDGMENTS

We would like to thank Adam O’Neill and George Kellaris for helpful discussions. George Kollios and Dmytro Bogatov were supported by an NSF SaTC Frontier Award CNS-1414119. Leonid Reyzin was supported in part by NSF grant 1422965.

References

- [1] .NET Foundation. *Benchmark.NET*. <https://github.com/dotnet/BenchmarkDotNet>. 2018.
- [2] R. Agrawal, J. Kiernan, R. Srikant, and Y. Xu. “Order Preserving Encryption for Numeric Data”. In: *Proceedings of the 2004 ACM SIGMOD International Conference on Management of Data*. SIGMOD ’04. ACM, 2004, pp. 563–574.
- [3] L. Arge. “The Buffer Tree: A Technique for Designing Batched External Data Structures”. In: *Algorithmica* 37.1 (Sept. 2003), pp. 1–24.
- [4] R. Bayer and E. McCreight. “Organization and Maintenance of Large Ordered Indices”. In: *Proceedings of the 1970 ACM SIGFIDET (Now SIGMOD) Workshop on Data Description, Access and Control*. SIGFIDET ’70. ACM, 1970, pp. 107–141.
- [5] V. Bindschaedler, P. Grubbs, D. Cash, T. Ristenpart, and V. Shmatikov. “The Tao of Inference in Privacy-protected Databases”. In: *Proc. VLDB Endow.* 11 (2018), pp. 1715–1728.
- [6] T. Boelter, R. Poddar, and R. A. Popa. “A Secure One-Roundtrip Index for Range Queries”. In: *IACR Cryptology ePrint Archive* (2016), p. 568.
- [7] D. Bogatov. *ORE Benchmark*. <https://github.com/dbogatov/ore-benchmark>. 2018.
- [8] D. Bogatov. *Interactive Secure Range Queries Simulations*. <https://ore.dbogatov.org/>. 2019.
- [9] A. Boldyreva, N. Chenette, Y. Lee, and A. O’Neill. “Order-Preserving Symmetric Encryption”. In: *Advances in Cryptology - EUROCRYPT 2009*. Springer Berlin Heidelberg, 2009, pp. 224–241.
- [10] A. Boldyreva, N. Chenette, and A. O’Neill. “Order-Preserving Encryption Revisited: Improved Security Analysis and Alternative Solutions”. In: *Advances in Cryptology - CRYPTO 2011*. Springer Berlin Heidelberg, 2011, pp. 578–595.
- [11] D. Boneh, K. Lewi, M. Raykova, A. Sahai, M. Zhandry, and J. Zimmerman. “Semantically Secure Order-Revealing Encryption: Multi-input Functional Encryption Without Obfuscation”. In: *Advances in Cryptology - EUROCRYPT 2015*. Springer Berlin Heidelberg, 2015, pp. 563–594.
- [12] M. Bun and M. Zhandry. “Order-Revealing Encryption and the Hardness of Private Learning”. In: *Theory of Cryptography*. Springer Berlin Heidelberg, 2016, pp. 176–206.
- [13] D. Cash, P. Grubbs, J. Perry, and T. Ristenpart. “Leakage-Abuse Attacks Against Searchable Encryption”. In: *Proceedings of the 22Nd ACM SIGSAC Conference on Computer and Communications Security*. ACM, 2015, pp. 668–679.
- [14] D. Cash, J. Jaeger, S. Jarecki, C. Jutla, H. Krawczyk, and M. Steiner. “Dynamic searchable encryption in very-large databases: Data structures and implementation”. In: *In Network and Distributed System Security Symposium (NDSS ’14)*. 2014.
- [15] D. Cash, S. Jarecki, C. Jutla, H. Krawczyk, M.-C. Rosu, and M. Steiner. “Highly-Scalable Searchable Symmetric Encryption with Support for Boolean Queries”. In: Springer Berlin Heidelberg, 2013, pp. 353–373.
- [16] D. Cash, F.-H. Liu, A. O’Neill, M. Zhandry, and C. Zhang. “Parameter-Hiding Order Revealing Encryption”. In: *Advances in Cryptology - ASIACRYPT 2018*. 2018. Forthcoming.
- [17] D. Cash, F.-H. Liu, A. O’Neill, and C. Zhang. *Reducing the Leakage in Practical Order-Revealing Encryption*. Cryptology ePrint Archive, Report 2016/661. 2016.
- [18] Z. Chang, D. Xie, and F. Li. “Oblivious RAM: A Dissection and Experimental Evaluation”. In: *Proc. VLDB Endow.* 9.12 (2016), pp. 1113–1124.
- [19] N. Chenette, K. Lewi, S. A. Weis, and D. J. Wu. “Practical Order-Revealing Encryption with Limited Leakage”. In: *Fast Software Encryption*. Springer Berlin Heidelberg, 2016, pp. 474–493.
- [20] CipherCloud. <https://www.ciphercloud.com/>.
- [21] I. Demertzis, S. Papadopoulos, O. Papapetrou, A. Deligiannakis, and M. Garofalakis. “Practical Private Range Search Revisited”. In: ACM, 2016, pp. 185–198.
- [22] I. Demertzis, S. Papadopoulos, O. Papapetrou, A. Deligiannakis, and M. N. Garofalakis. “Practical Private Range Search Revisited”. In: *Proceedings of the 2016 International Conference on Management of Data*. 2016, pp. 185–198.
- [23] F. B. Durak, T. M. DuBuisson, and D. Cash. “What Else is Revealed by Order-Revealing Encryption?” In: *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*. ACM, 2016, pp. 1155–1166.
- [24] Y. Elovici, R. Waisenberg, E. Shmueli, and E. Gudes. “A Structure Preserving Database Encryption Scheme”. In: *Secure Data Management*. Springer Berlin Heidelberg, 2004, pp. 28–40.
- [25] J. Eom, D. H. Lee, and K. Lee. “Multi-Client Order-Revealing Encryption”. In: *IEEE Access* (2018), pp. 45458–45472.
- [26] O. Goldreich and R. Ostrovsky. “Software Protection and Simulation on Oblivious RAMs”. In: *J. ACM* 43.3 (May 1996), pp. 431–473.
- [27] P. Grubbs, T. Ristenpart, and V. Shmatikov. “Why Your Encrypted Database Is Not Secure”. In: *Proceedings of the 16th Workshop on Hot Topics in Operating Systems*. ACM, 2017, pp. 162–168.
- [28] P. Grubbs, K. Sekniqi, V. Bindschaedler, M. Naveed, and T. Ristenpart. “Leakage-Abuse Attacks against Order-Revealing Encryption”. In: *2017 IEEE Symposium on Security and Privacy (SP)* (2016), pp. 655–672.
- [29] H. Haagh, Y. Ji, C. Li, C. Orlandi, and Y. Song. “Revealing Encryption for Partial Ordering”. In: *Cryptography and Coding*. Springer International Publishing, 2017, pp. 3–22.
- [30] V. T. Hoang and P. Rogaway. “On Generalized Feistel Networks”. In: *Proceedings of the 30th Annual Conference on Advances in Cryptology*. Springer-Verlag, 2010, pp. 613–630.

- [31] M. S. Islam, M. Kuzu, and M. Kantarcioglu. "Access Pattern disclosure on Searchable Encryption: Ramification, Attack and Mitigation". In: *19th Annual Network and Distributed System Security Symposium, NDSS 2012, San Diego, California, USA, February 5-8, 2012*. 2012.
- [32] M. S. Islam, M. Kuzu, and M. Kantarcioglu. "Inference attack against encrypted range queries on outsourced databases". In: *Fourth ACM Conference on Data and Application Security and Privacy, CODASPY'14, San Antonio, TX, USA - March 03 - 05, 2014*. 2014, pp. 235–246.
- [33] J. Jannink. "Implementing Deletion in B+-trees". In: *SIGMOD Rec.* 24.1 (Mar. 1995), pp. 33–38.
- [34] V. Kachitvichyanukul and B. Schmeiser. "ALGORITHM 668: H2PEC: sampling from the hypergeometric distribution". In: 14 (Dec. 1988), pp. 397–398.
- [35] H. Kadhemi, T. Amagasa, and H. Kitagawa. "MV-OPES: Multivalued-Order Preserving Encryption Scheme: A Novel Scheme for Encrypting Integer Value to Many Different Values". In: (2010), pp. 2520–2533.
- [36] H. Kadhemi, T. Amagasa, and H. Kitagawa. "Optimization Techniques for Range Queries in the Multivalued-partial Order Preserving Encryption Scheme". In: *Knowledge Discovery, Knowledge Engineering and Knowledge Management*. Springer Berlin Heidelberg, 2013, pp. 338–353.
- [37] G. Kellaris, G. Kollios, K. Nissim, and A. O'Neill. "Generic Attacks on Secure Outsourced Databases". In: *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*. ACM, 2016, pp. 1329–1340.
- [38] F. Kerschbaum. "Frequency-Hiding Order-Preserving Encryption". In: *Proceedings of the 22Nd ACM SIGSAC Conference on Computer and Communications Security*. ACM, 2015, pp. 656–667.
- [39] F. Kerschbaum and A. Schroepfer. "Optimal Average-Complexity Ideal-Security Order-Preserving Encryption". In: *Proceedings of the 2014 ACM SIGSAC Conference on Computer and Communications Security*. ACM, 2014, pp. 275–286.
- [40] F. Kerschbaum and A. Tueno. "An Efficiently Searchable Encrypted Data Structure for Range Queries". In: *arXiv preprint arXiv:1709.09314* (2017).
- [41] A. Kiayias, S. Papadopoulos, N. Triandopoulos, and T. Zacharias. "Delegatable Pseudorandom Functions and Applications". In: *Proceedings of the 2013 ACM SIGSAC Conference on Computer & Communications Security*. ACM, 2013, pp. 669–684.
- [42] D. E. Knuth. *Seminumerical algorithms*. 3rd ed. Vol. 2. Addison-Wesley, 2016, pp. 145–146.
- [43] M. Lacharite, B. Minaud, and K. G. Paterson. "Improved Reconstruction Attacks on Encrypted Data Using Range Query Leakage". In: *2018 IEEE Symposium on Security and Privacy (SP)*. 2018, pp. 297–314.
- [44] K. G. Larsen and J. B. Nielsen. "Yes, There is an Oblivious RAM Lower Bound!" In: *Advances in Cryptology - CRYPTO 2018*. 2018, pp. 523–542.
- [45] K. Lewi and D. J. Wu. "Order-Revealing Encryption: New Constructions, Applications, and Lower Bounds". In: ACM, 2016, pp. 1167–1178.
- [46] D. Liu and S. Wang. "Nonlinear order preserving index for encrypted database query in service cloud environments". In: *Concurrency and Computation: Practice and Experience* (), pp. 1967–1984.
- [47] D. Liu and S. Wang. "Programmable Order-Preserving Secure Index for Encrypted Database Query". In: *Proceedings - 2012 IEEE 5th International Conference on Cloud Computing, CLOUD 2012*. 2012, pp. 502–509.
- [48] Z. Liu, K.-K. R. Choo, and M. Zhao. "Practical-oriented protocols for privacy-preserving outsourced big data analysis: Challenges and future research directions". In: *Computers & Security* 69 (2017), pp. 97–113.
- [49] B. Lynn. *Pairings-based Crypto (PBC)*. 2018. URL: <https://crypto.stanford.edu/pbc/> (visited on 08/15/2018).
- [50] M. Maffei, M. Reinert, and D. Schröder. "On the Security of Frequency-Hiding Order-Preserving Encryption". In: *Proceedings of the International Conference on Cryptology and Network Security*. Springer, 2017.
- [51] M. Naveed, S. Kamara, and C. V. Wright. "Inference Attacks on Property-Preserving Encrypted Databases". In: *Proceedings of the 22Nd ACM SIGSAC Conference on Computer and Communications Security*. ACM, 2015, pp. 644–655.
- [52] G. Özsoyoglu, D. A. Singer, and S. S. Chung. "Anti-Tamper Databases: Querying Encrypted Databases". In: *Data and Applications Security XVII: Status and Prospects, IFIP TC-11 WG 11.3 Seventeenth Annual Working Conference on Data and Application Security, August 4-6, 2003, Estes Park, Colorado, USA*. 2003, pp. 133–146.
- [53] R. Popa, F. Li, and N. Zeldovich. "An Ideal-Security Protocol for Order-Preserving Encoding". In: *IEEE Symposium on Security and Privacy*. 2013, pp. 463–477.
- [54] R. Popa, C. M. S. Redfield, N. Zeldovich, and H. Balakrishnan. "CryptDB: Protecting Confidentiality with Encrypted Query Processing". In: *Proceedings of the Twenty-Third ACM Symposium on Operating Systems Principles*. SOSP '11. ACM, 2011, pp. 85–100.
- [55] D. S. Roche, D. Apon, S. G. Choi, and A. Yerukhimovich. "POPE: Partial Order Preserving Encoding". In: *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*. ACM, 2016, pp. 1131–1142.
- [56] B. Schneier and J. Kelsey. "Unbalanced Feistel networks and block cipher design". In: *Fast Software Encryption*. Springer Berlin Heidelberg, 1996, pp. 121–144.
- [57] *Skyhigh Networks*. <https://www.skyhighnetworks.com/>.
- [58] E. Stefanov, M. van Dijk, E. Shi, C. Fletcher, L. Ren, X. Yu, and S. Devadas. "Path ORAM: An Extremely Simple Oblivious RAM Protocol". In: *Proceedings of the 2013 ACM SIGSAC Conference on Computer Communications Security*. ACM, 2013, pp. 299–310.

- [59] I. Teranishi, M. Yung, and T. Malkin. “Order-Preserving Encryption Secure Beyond One-Wayness”. In: *Advances in Cryptology – ASIACRYPT 2014*. Springer Berlin Heidelberg, 2014, pp. 42–61.
- [60] Transparent California. *2017 salaries for State of California*. <https://transparentcalifornia.com/salaries/2017/state-of-california/>. 2017.
- [61] A. J. Walker. “An Efficient Method for Generating Discrete Random Variables with General Distributions”. In: *ACM Trans. Math. Softw.* 3.3 (Sept. 1977), pp. 253–256.
- [62] X. Wang and Y. Zhao. “Order-Revealing Encryption: File-Injection Attack and Forward Security”. In: *Computer Security*. Springer International Publishing, 2018, pp. 101–121.
- [63] S. Wozniak, M. Rossberg, S. Grau, A. Alshawish, and G. Schaefer. “Beyond the Ideal Object: Towards Disclosure-resilient Order-preserving Encryption Schemes”. In: *Proceedings of the 2013 ACM Workshop on Cloud Computing Security Workshop*. ACM, 2013, pp. 89–100.
- [64] L. Xiao and I.-l. Yen. *A Note for the Ideal Order-Preserving Encryption Object and Generalized Order-Preserving Encryption*.
- [65] L. Xiao, I.-L. Yen, and D. T. Huynh. “Extending Order Preserving Encryption for Multi-User Systems”. In: *IACR Cryptology ePrint Archive* (2012), p. 192.