

Revisiting Leakage Abuse Attacks

Laura Blackstone
Brown University
laura_blackstone@alumni.brown.edu

Seny Kamara
Brown University
seny@brown.com

Tarik Moataz*
Aroki Systems
tarik@aroki.com

Abstract—Encrypted search algorithms (ESA) are cryptographic algorithms that support search over encrypted data. ESAs can be designed with various primitives including searchable/structured symmetric encryption (SSE/STE) and oblivious RAM (ORAM). Leakage abuse attacks attempt to recover client queries *using knowledge of the client's data*. An important parameter for any leakage-abuse attack is its *known-data rate*; that is, the fraction of client data that must be known to the adversary.

In this work, we revisit leakage abuse attacks in several ways. We first highlight some practical limitations and assumptions underlying the well-known IKK (Islam et al. *NDSS '12*) and Count (Cash et al., *CCS '15*) attacks. We then design four new leakage-abuse attacks that rely on much weaker assumptions. Three of these attacks are *volumetric* in the sense that they only exploit leakage related to document sizes. In particular, this means that they work not only on SSE/STE-based ESAs but also against ORAM-based solutions. We also introduce two volumetric *injection* attacks which use adversarial file additions to recover queries even from ORAM-based solutions. As far as we know, these are the first attacks of their kind.

We evaluated all our attacks empirically and considered many experimental settings including different data collections, query selectivities, known-data rates, query space size and composition. From our experiments, we observed that the only setting that resulted in reasonable recovery rates under practical assumptions was the case of high-selectivity queries with a leakage profile that includes the response identity pattern (i.e., the identifiers of the matching documents) and the volume pattern (i.e., the size of the matching documents). All other attack scenarios either failed or relied on unrealistic assumptions (e.g., very high known-data rates). For this specific setting, we propose several suggestions and countermeasures including the use of schemes like PBS (Kamara et al, *CRYPTO '18*), VLH/AVLH (Kamara and Moataz, *Eurocrypt '19*), or the use of padding techniques like the ones recently proposed by Bost and Fouque (Bost and Fouque, *IACR ePrint 2017/1060*).

I. INTRODUCTION

The area of encrypted search is concerned with the design and analysis of cryptographic techniques to search over encrypted data. There are many ways to design encrypted search algorithms (ESA) including using fully-homomorphic encryption (FHE) [28], oblivious RAM (ORAM) [31], functional

encryption [9], property-preserving encryption [2], [7], [8] and searchable/structured symmetric encryption (SSE/STE) [69], [19], [18]. All these approaches achieve different tradeoffs between leakage, expressiveness and efficiency.

At a high level, a static ESA is composed of two algorithms: a setup algorithm and a search algorithm. Setup encrypts a data collection in such a way that it can later be queried using the search algorithm. The most basic form of search is single-keyword search which, given a keyword w , returns the documents in the collection that contain w . If the solution is dynamic, there is an additional update algorithm to modify the encrypted data collection. In this work, we will focus on SSE/STE- and ORAM-based ESAs, which we sometimes refer to as *structured ESAs* and *oblivious ESAs*, respectively.

Leakage. While it is possible to search on encrypted data with essentially no leakage with time and communication that is linear in the size of the document collection, this is however not a viable approach for datasets of practical interest. Because of this, all known ORAM-, STE- and PPE-based ESAs leak some information. This leakage comes in two forms: setup leakage, which is revealed at setup time by the encrypted dataset, and query leakage, which is revealed at query time by the encrypted dataset and the query operation. To better understand the real-world impact of this leakage, an important research direction in encrypted search has been to design leakage attacks. This line of work was initiated by Islam, Kuzu and Kantarcioglu in [38] in the context of SSE and was expanded to PPE by Naveed, Kamara and Wright [56] and to ORAM by Kellaris, Kollios, Nissim and O'Neill [46]. Since then, several works have further explored leakage attacks including [13], [75], [48], [33] in the SSE setting and [13], [22], [34] in the PPE setting.

Leakage attacks. Leakage attacks come in different forms depending on the leakage profiles they exploit, the adversarial models in which they work, the information they recover, the auxiliary information they need and the assumptions they rely on. We can categorize attacks along the following dimensions:

- **adversarial model:** *snapshot* attacks only require access to the encrypted document collection; *persistent* attacks require access to the encrypted data collection and to the transcripts of the query operations.
- **target:** *data-recovery* attacks recover information about the data collection whereas *query-recovery* attacks recover data about queries;
- **auxiliary data:** *sampled-data* attacks require a sample from a distribution that is close (e.g., in statistical distance) to the distribution of the data collection; *known-query* attacks

* : Work done in part while at Brown University.

require knowledge of a subset of the queries; *known-data* attacks require knowledge of a subset of the data collection. An important parameter for known-data attacks is the *known-data rate* which is defined as the fraction δ of documents in the client’s data collection that are known to the adversary.

- passive or active: passive attacks do not require the adversary to choose any part of the client’s data; active, or *chosen-data* attacks, require the adversary to choose some of the data in the data collection.

In this work, we focus on the persistent model since it has recently been shown that solutions with little to no leakage can be achieved in the snapshot setting using both PPE [50] and STE [3] (this doesn’t take into account possible systems-level pitfalls as pointed out in [32]). We recall that sampled-data attacks are commonly referred to as *inference* attacks, that known-data attacks are commonly referred to as *leakage abuse* attacks and that chosen-data attacks are commonly referred to as *injection* attacks.

Leakage profiles. Of course, an important characteristic of any leakage attack is what kind of leakage it exploits. For example, the IKK attack exploits the *co-occurrence pattern* which reveals, for every pair of queries, the number of times they appear in the same document. The Count attack rely on the co-occurrence pattern and the *response length pattern*; the latter of which reveals, for every query, the number of documents that contain it. The injection attacks of [75] exploit the *response identity* (also known as the *access pattern*) which reveals, for every query, the identifiers of the documents that contain it. Note that the response identity reveals the response length and the co-occurrence pattern so the IKK and Count attack can apply to any construction that leaks the response identity. An important class of leakage patterns for our purposes will be what we call *volumetric* patterns. By this we mean any leakage pattern that reveals the size of documents. Here, we will focus specifically on the *volume pattern* which reveals, for each query, the volumes of the documents that contain it; and the *total volume pattern* which reveals, for each query, the sum of the volumes of the documents that contain it. All the patterns discussed above are common in the literature. In particular, they are part of the leakage profile of all standard response-revealing multi-map encryption schemes [19], [18], [42], [41], [17], [10], which are the encrypted structures that underlie most single-keyword searchable symmetric encryption schemes. Some of these leakage patterns (e.g., such as the volume pattern and the total volume pattern) are also part of the leakage profile of ORAM-based ESAs.

Known-data attacks against structured ESAs. Known-data attacks were introduced by Cash, Grubbs, Perry and Ristenpart in [13]. In that work, they described several attacks against both SSE/STE- and PPE-based ESAs. They also introduced injection (or *chosen-data*) attacks against PPE-based ESAs. Injection attacks were later demonstrated against structured ESAs by Zhang, Katz and Papamanthou [75].

The IKK attack was first described in [38] as an *inference* attack that exploits the co-occurrence pattern. [38] reported high recovery rates but the experiments conducted had several methodological flaws. The most salient ones were that: (1) they were run on a small query space (of size 2500 out of a total of 77000 after stemming and removing stop words); and (2)

the training and test data collections were not independent. Motivated by this, Cash et al. re-evaluated the IKK attack with independent testing and training data and found that IKK could not recover any queries. IKK was then re-evaluated as a known-data attack and it was found that it could achieve reasonable recovery rates if it was given 95% or more of the client’s data. Effectively, [13] showed that IKK failed as an inference attack but worked as a known-data attack when $\delta \geq .95$. [13] then introduced a new attack called the Count attack. This attack relies on co-occurrence and response length leakage and was shown to perform better than the IKK attack. Recently, the ePrint version of [16] was updated to include a new attack that performs better than the one originally published in [13]. Throughout this work, we will refer to the first Count attack as Count v.1 and to the new attack as Count v.2.

Discussion and overview of our contributions. The IKK and Count attacks have received a lot of attention and are commonly used to draw conclusions about various ESAs. As an example, they are often cited as a reason to prefer oblivious ESAs over structured ESAs [71], [72], [74], [52], [20], [64], [27], [25], [29], [65], [58], [53], [51], [67], [66], [73], [26], [5], [37]. The results in this work underscore that the study of and, especially, the *interpretation* of known-data attacks is more nuanced.

To address this, we present in Section III several new known-data attacks that do not have these limitations and that achieve higher recovery rates for much lower known-data rates. What is perhaps surprising about our attacks is that they work not only against structured ESAs but also against *oblivious* ESAs which contradicts the conventional wisdom that ORAM-based search is resistant to leakage-abuse attacks.

Another contribution of our work is that we demonstrate, for the first time, that injection attacks apply not only to PPE-based ESAs [13] and structured ESAs [75] but also to oblivious ESAs. In particular, we describe in Section IV two volumetric injection attacks, which contradicts the conventional wisdom that ORAM-based search is resistant to injection attacks.

We evaluated our attacks empirically and provide an overview of our results in Section V. For a complete empirical analysis we refer the reader to the full version of this work. Specifically, we evaluate the attacks in a host of different settings including different keyword selectivities, query space compositions, query space sizes and datasets. This extensive evaluation shows that the success rates of known-data attacks is very sensitive to various parameters that were not considered in previous work. This highlights the importance of common but often implicit assumptions made in the leakage attack literature.

Finally, in Section VII we propose several countermeasures to mitigate all known-data attacks, including our own.

A. Theory vs. Practice of Known-Data Attacks

Here, we revisit the state-of-the art in known-data attacks highlighting some of the practical limitations and assumptions of the currently-known attacks.

Reliance on co-occurrence. Recall that all IKK, Count v.1 and Count v.2 attacks all rely on the co-occurrence pattern. This particular leakage pattern, however, can be hidden using standard SSE/STE techniques. In fact, we describe a construction in Section II and, in more details, in Appendix A, OPQ, that does not reveal the co-occurrence pattern. As far as we know, this construction has not appeared in previous work and may be of independent interest.¹

High known-data rates. The experimental results in [13] (cf. Figure 6) and [16] (cf. Figure 6) show that the IKK, Count v.1 and Count v.2 attacks achieve non-trivial recovery rates only with very high known-data rates: IKK needs to know 70% of the data to recover 5% of the queries; Count v.1 needs to know 80% of the data to recover 40% of the queries (note that knowing even up to 75% of the data results in a query recovery rate of 0%); and Count v.2 needs to know 75% of the data to recover 40% of the queries. Given how high these known-data rates are, it is not clear whether these attacks should be considered practical. An instructive question to ask here is how exactly an adversary could, in practice, get up to, say 75%, of a client’s data? Recall that in the encrypted search setting the client encrypts its data and *outsources* it to an untrusted server. In particular, this means the client deletes the data from its system after setup which leaves the only copy on the server and in encrypted form. In such a setting, there are a few scenarios in which an adversary could recover 75% of the client’s data. One is that the client, for some reason, chooses to encrypt public data which the adversary later recovers. A second is that the client decides to release a large percentage of its data (after downloading and decrypting it from the server). A third is that the client queries its data and over time caches enough of the results to amount to a large percentage of the data. At this stage, a data breach occurs on the client and the cached data is revealed to the adversary. The first two scenarios are relatively contrived and have more to do with a misuse of the primitive: in such settings one should use private information retrieval. The third scenario is perhaps less contrived but caching 75% of one’s data locally seems to defeat the purpose of outsourced storage. Indeed, if a client is willing to store 75% of its data locally then they might as well do search locally on the 75% and use encrypted search only for the remaining 25%.

High- vs. low-selectivity keywords. An important consideration when evaluating a query-recovery attack is how exactly client queries are chosen. Most leakage attack papers make implicit assumptions about this but the query distribution has a large impact on accuracy rates. For example, the experiments in [13] assume the client queries high-selectivity keywords, where selectivity refers to the number of documents matching a particular keyword. It is not clear, however, if this is realistic. In fact, we ran the IKK and Count attacks on low-selectivity keywords over the Enron dataset and neither attack worked; even when the adversary had a complete knowledge of the client’s data. More precisely, the IKK and Count attacks had recovery rate 0 even when $\delta = 1$ for keyword selectivities lower than 20.

Known queries. The IKK and Count v.1 attacks are described in [13] as known-data attacks that do not require

knowledge of any client queries. While it is true that these attacks can achieve high recovery rates without knowledge of client queries, this only holds if the adversary has complete knowledge of the data; that is, if $\delta = 1$. If the adversary has less than full knowledge (i.e., $\delta < 1$) and no knowledge of any queries, then Count v.1 does not work.²

Theoretical vs. practical attacks. In cryptanalysis it is common to distinguish between theoretical attacks and practical attacks.³ The former are attacks that work in strong adversarial models; often relying on assumptions about the adversary’s capabilities which rarely occur in practice. Examples include related-key attacks where the adversary is allowed to make chosen-plaintext and chosen-ciphertext queries under related keys (e.g., through modification of keys). Another is the known-key model where the adversary is assumed to know the key and its goal is to distinguish ciphertexts from random. It is our belief that the IKK and Count attacks are mostly of theoretical interest since they rely on strong assumptions like known-queries and high known-data rates and have only been shown to work on high-selectivity keywords.

Should we discount theoretical attacks? Though these attacks are of theoretical interest it does not mean we should dismiss them. Even theoretical cryptanalytic results have something to teach us about the security of our constructions. For example, even attacks that require high known-data rates can be interesting if they exploit a leakage profile that, up to this point, had not been successfully cryptanalyzed. Especially since one can assume that attacks always improve. Such results serve as a warning and motivation to design schemes with better leakage profiles. It is important, however, to be clear and explicit about the limitations and implications of cryptanalytic results.

B. Our Attacks

Motivated by the discussion above, we revisit known-data attacks (i.e., leakage abuse attacks) against SSE. We introduce four new attacks and two new injection attacks in Sections III and IV. We also perform a thorough evaluation which we report on in Section V and provide in detail in the full version of this work. Our attacks achieve high recovery rates with low known-data rates and do not rely on known queries. We summarize the characteristics of our attacks in Table I. Most surprisingly, all but one of our attacks are volumetric and, therefore, apply not only to SSE/STE-based solutions but also to ORAM-based constructions. As far as we know, these are the first known-data and chosen-data attacks against ORAM. We now summarize our attacks and their performance:

- Volume analysis (VolAn): a known-data attack that exploits the *total volume pattern*. It has high recovery rates when $\delta \geq .8$ and the client queries keywords with high-selectivity (i.e., 10-13) or pseudo-low-selectivity (i.e., 1-2).

²This may seem to contradict the results presented in [13] but the experimental evaluation of Count v.1 presented in Figure 6 of [13] was incomplete. For $\delta = 1$ (i.e., complete knowledge) the attack does not need knowledge of queries but for $\delta < 1$ it does. We confirmed this with the authors who updated their manuscript with the new Count v.2 attack which does not require knowledge of queries.

³We highly recommend the paper of Aumasson [6] for an insightful discussion of these issues.

¹Oblivious RAM can also be used to design an ESA that hides the co-occurrence pattern. We describe one such design, which we call FLL, in Section II and in Appendix A.

Attack	Type			Leakage pattern	Known queries	δ for high selectivity	δ for p-low selectivity	δ for low selectivity
	Sampled	Known	Injection					
IKK [38]	\times	\checkmark	\times	co	\checkmark	$\geq 95\%$	\odot	\odot
Count [13]	\times	\checkmark	\times	co, rlen	\checkmark	$\geq 80\%$	\odot	\odot
Zhang et al. [75]	\times	\times	\checkmark	rid	\times	—	—	—
Subgraph ^{ID}	\times	\checkmark	\times	rid	\times	$\geq 5\%$	$\geq 50\%$	$\geq 60\%$
Subgraph ^{VL}	\times	\checkmark	\times	vol	\times	$\geq 5\%$	$\geq 50\%$	\odot
VolAn	\times	\checkmark	\times	tvol	\times	$\geq 85\%$	$\geq 85\%$	\odot
SelVolAn	\times	\checkmark	\times	tvoll, rlen	\times	$\geq 80\%$	$\geq 85\%$	\odot
Decoding & Binary	\times	\times	\checkmark	tvoll	\times	—	—	—

Table I: Comparison of existing leakage abuse (known-data) and injection (chosen-data) attacks. The last three columns give the known-data rate δ needed for a recovery rate of at least 20% against low-, pseudo-low- and high-selectivity keywords, respectively. \odot means the experiment was not conducted in previous work. \circ means that even a known-data rate of $\delta = 1$ does not achieve at least 20% recovery rate. All experiments were based on the Enron dataset [63] with 150 queries and a keyword space of 500 keywords.

- Selective volume analysis (SelVolAn): an extension of volume analysis that relies on the *total volume* and *response length* patterns. This attack has slightly higher recovery rates under the same conditions as volume analysis.
- Subgraph attacks: a framework to design known-data attacks against *atomic* leakage patterns, i.e., leakage pattern that reveals information about each matching document. We give two concrete instantiations of our framework. The first is Subgraph^{ID} which exploits the *response identity* and the second is Subgraph^{VL} which exploits the *volume pattern*. Both attacks achieve high recovery rates with very low known-data rates (i.e., with δ as low as .05) when the client queries high-selectivity keywords. The recovery rate drops significantly and reaches 0% when the client queries keywords with low selectivity (i.e., 1-2) and pseudo-low selectivity.
- the Decoding attack (Decoding): an injection attack that exploits the *total volume pattern*. This attack always recovers its target query if the adversary can inject between 4 to 16KBytes depending on the query’s selectivity.
- the Binary Search attack (Binary): an injection attack that also exploits the *total volume pattern*. The attack requires logarithmic number of (adaptive) injections. The attack recovers its target query if the adversary can inject around 8KBytes.

Remark. As described above, our evaluation shows that even our new attacks can fail to recover queries in certain settings. This finding is important as it shows that existing schemes can be good enough to use in some scenarios. In Section VI, we provide a set of takeaways summarizing our findings and conclusions.

C. Countermeasures

We propose several countermeasures and guidelines against both our new attacks and previously-known attacks.

As discussed earlier, our empirical evaluation found only a single practical setting where our attacks are successful: querying high-selectivity keywords using a scheme that leaks both the response identity and the volume patterns. The simplest countermeasure to this is to use a scheme that does not leak these patterns like the PBS construction of Kamara, Moataz and Ohrimenko [45]. In Section VII, we demonstrate empirically that PBS is resistant to all known-data attacks (even ours) as long as the client makes at least 4 queries. Specifically, we show that under this condition, the best possible attack has recovery rate 0.02% recovery rate.

To mitigate theoretical attacks, like the IKK or Count attack which require high known-data rates and exploit the co-occurrence pattern, we design a new scheme called OPQ that does not leak this pattern. We also point out that, in [12], Bost and Fouque introduced padding techniques that efficiently mitigate these attacks.

Finally, to protect against purely volumetric attacks one can use the recent constructions of Kamara and Moataz [44] which are volume-hiding.

D. Related Work

SSE/STE. SSE was introduced by Song, Wagner and Perrig [69]. Curtmola, Garay, Kamara and Ostrovsky formalized SSE in [19] and described the first sub-linear and optimal-time constructions. STE was introduced by Chase and Kamara in [18] as a generalization of SSE. Many works have explored various aspects of SSE including dynamism [30], [42], [41], [70], [10], [17], [11], [23], locality [17], [15], [4], [55], [21], expressiveness [18], [14], [59], [24], [54], [40], [43], multiple clients [19], [39], [60], [35], and leakage [27], [45], [44], [3].

ORAM. Goldreich and Ostrovsky introduced ORAM in [31], where they described constructions with amortized square-root and polylog overheads. Shi, Chan, Stefanov and Li introduced tree-based ORAMs which achieved worst-case polylog overhead in [68]. Since then, many works have improved ORAM along many dimensions including communication complexity, round complexity, client and server storage [47], [71], [64], [27].

The IKK attack [38]. The attack takes as input the co-occurrence pattern, a background matrix, and a set of known queries. The co-occurrence pattern is a matrix with rows and columns indexed by queries (not keywords) and where the element in the i th row and j th columns is the (normalized) number of documents that contain both the i th and j th query. The co-occurrence matrix captures the information leaked to the adversary. The background matrix is a matrix with rows and columns indexed by keywords and where the element in the i th row and j th column is the (normalized) number of documents that contain both keywords with noise added. This background matrix is meant to represent auxiliary information available to the adversary. The IKK attack solves an optimization problem to find a mapping between the two matrices which leads to a mapping between queries and keywords. The paper shows that this optimization problem is NP-complete, but can be efficiently approximated using simulated annealing.

Count v.1 [13]. The count attack takes as input the co-occurrence pattern, the response length pattern (called *counts* in [13]), and a subset of the user’s data. It starts by mapping the queries in the query sequence to their response lengths and keywords in the known dataset to their response lengths. It then maps all the queries with *unique* response lengths to the keyword with the same response length. These recovered keywords then serve as *anchors* for the second step of the attack where each remaining query is mapped to a keyword with the same response-length and the same co-occurrences with respect to the recovered keywords. The attack was shown to achieve high recovery rates when the known-data rate $\delta = 1$. When $\delta < 1$, the attack also needs knowledge of some fraction of the user’s queries.

Count v.2 [16]. Count v.2 is similar to Count v.1 except that it does not require knowledge of user queries when $\delta < 1$. The main challenge in this setting is in identifying the anchors since the response lengths computed from the known data are no longer accurate which leads to failure of the Count v.1. Count v.2, however, uses a new mechanism that takes into account the distribution of the keywords in the corpus. The resulting anchors are only accurate with a certain probability so the accuracy of the overall attack is probabilistic. The rest of the attack is similar to v.1.

Again, we note that all three attacks are subject to the limitations discussed in Section I-A.

II. PRELIMINARIES

Notation. The set of all binary strings of length n is denoted as $\{0, 1\}^n$, and the set of all finite binary strings as $\{0, 1\}^*$. $[n]$ is the set of integers $\{1, \dots, n\}$, and $2^{[n]}$ is the corresponding power set. The output x of an algorithm \mathcal{A} is denoted by $x \leftarrow \mathcal{A}$. Given a sequence \mathbf{q} of n elements, we refer to its i th element as q_i or $\mathbf{q}[i]$. If S is a set then $\#S$ refers to its cardinality. If s is a string then $|s|_2$ refers to its bit length. Throughout, k will denote the security parameter.

The word RAM. Our model of computation is the word RAM. In this model, we assume memory holds an infinite number of w -bit words and that arithmetic, logic, read and write operations can all be done in $O(1)$ time. We denote by $|x|_w$ the word-length of an item x ; that is, $|x|_w = |x|_2/w$. Here, we assume that $w = \Omega(\log k)$.

Document collections. Let \mathbb{W} be a keyword space. A document collection $\mathbf{D} = (D_1, \dots, D_n)$ over \mathbb{W} consists of n documents, each of which is a subset of \mathbb{W} . We denote by \mathbb{D} the space of all document collections over \mathbb{W} . We assume each document $D \in \mathbf{D}$ has a unique identifier that is independent from its contents. For ease of exposition, we assume these identifiers are the integers 1 through n and that they are assigned to documents uniformly at random. For ease of exposition, it will be helpful to consider the following functions. The identifier function $\text{id} : \mathbf{D} \rightarrow [n]$ that maps a document D_i to its identifier i . The function $\mathbf{D} : \mathbb{W} \rightarrow 2^{\mathbf{D}}$ that maps a keyword w to the documents that contain it. The identifiers function $\text{ids} : \mathbb{W} \rightarrow 2^{[n]}$ that maps a keyword w to the identifiers of the documents that contain it. We refer to the word-length of a document $|D|_w$ as its *volume*.

Throughout this work, we denote the adversary’s known dataset by \mathbf{D} and assume it is a subset of the client’s document collection \mathbf{D} chosen uniformly at random.

Structured ESAs. A static structured encrypted search algorithm $\text{ESA} = (\text{Setup}, \text{Search})$ consists of two efficient algorithms. Setup takes as input a security parameter 1^k and a document collection $\mathbf{D} = (D_1, \dots, D_n)$ and outputs a secret key K and an encrypted data collection $(\text{EDB}, \text{ct}_1, \dots, \text{ct}_n)$. Search is a two-party protocol between a client and a server. The client inputs its secret key K and a keyword w and the server inputs an encrypted collection $(\text{EDB}, \text{ct}_1, \dots, \text{ct}_n)$. The client receives a set of encrypted documents $\{\text{ct}_i\}_{i \in \text{ids}(w)}$ and the server receives \perp . Structured ESAs are constructed using structured encryption (STE) [18] and, in particular, using a multi-map or dictionary encryption schemes. We now describe two examples BSL and OPQ which we further detail in Appendix A.

Given a document collection, BSL first generates a multi-map (also known as an inverted index) that maps each keyword to the identifiers of the documents that contain the keyword. This structure is then encrypted using a response-revealing multi-map encryption scheme [19], [18], [42], [41], [17], [10], [3], while the documents are encrypted using a symmetric-key encryption scheme. This scheme has optimal search and storage complexities. We refer the reader to [3] for the most recent construction.

The OPQ construction first generates a multi-map that maps each keyword to the *documents* that contain it—as opposed to the identifiers of the documents that contain it. The multi-map is then encrypted using a response-hiding multi-map encryption scheme. While the underlying idea is straightforward, as far as we know, this scheme is new and has never appeared in prior work. It has a better leakage profile than BSL, has the same query complexity but incurs additional storage overhead since entire documents can be replicated instead of just identifiers.

Oblivious ESAs. ESAs can also be designed using ORAM. The simplest approach is similar to the structured ESA construction described above but where EDB is replaced with an oblivious RAM ORAM that stores a search structure (e.g., a multi-map). The Search algorithm then executes the structure’s query algorithm and replaces each read operation with a call to the ORAM’s access protocol. We describe some examples of oblivious ESA constructions in Appendix A, and provide a high level overview below.

Given a document collection, SMI creates a multi-map that maps each keyword to the identifiers of the documents that contain it. This multi-map is then stored and managed in an ORAM (or stored in a custom oblivious multi-map structure [74]), while the documents are encrypted using a symmetric-key encryption scheme. This construction has a better leakage profile than BSL but is less efficient than both BSL and OPQ. In particular, the oblivious structure incurs a logarithmic multiplicative overhead in communication complexity and multiple rounds of interaction.

FLL is similar in that it stores a similar multi-map in an ORAM but, unlike SMI, it also stores blocks of the documents in a multi-map that is itself stored in an ORAM. This solution

has a better leakage profile than all the schemes above but also incurs a logarithmic multiplicative overhead and multiple rounds of communication.

Modeling leakage. Each ESA operation is associated with leakage which itself can be composed of multiple *leakage patterns*. The collection of all these leakage patterns forms the scheme’s *leakage profile*. Leakage patterns are (families of) functions over the various spaces associated with the underlying data collection. For concreteness, we recall some well-known leakage patterns:

- the *query equality pattern* is the function family $\text{qeq} = \{\text{qeq}_{k,t}\}_{k,t \in \mathbb{N}}$ with $\text{qeq}_{k,t} : \mathbb{D}_k \times \mathbb{W}_k^t \rightarrow \{0, 1\}^{t \times t}$ such that $\text{qeq}_{k,t}(\mathbf{D}, w_1, \dots, w_t) = M$, where M is a binary $t \times t$ matrix such that $M[i, j] = 1$ if $w_i = w_j$ and $M[i, j] = 0$ if $w_i \neq w_j$. The query equality pattern is referred to as the search pattern in the SSE literature;
- the *identifier pattern* is the function family $\text{rid} = \{\text{rid}_{k,t}\}_{k,t \in \mathbb{N}}$ with $\text{rid}_{k,t} : \mathbb{D}_k \times \mathbb{W}_k^t \rightarrow [2^{[n]}]^t$ such that $\text{rid}_{k,t}(\mathbf{D}, w_1, \dots, w_t) = (\text{ids}(w_1), \dots, \text{ids}(w_t))$. The identifier pattern is referred to as the access pattern in the SSE literature;
- the *response length pattern* is the function family $\text{rlen} = \{\text{rlen}_{k,t}\}_{k,t \in \mathbb{N}}$ with $\text{rlen}_{k,t} : \mathbb{D}_k \times \mathbb{W}_k^t \rightarrow \mathbb{N}$ such that $\text{rlen}_{k,t}(\mathbf{D}, w_1, \dots, w_t) = (\#\mathbf{D}(w_1), \dots, \#\mathbf{D}(w_t))$;
- the *volume pattern* is the function family $\text{vol} = \{\text{vol}_{k,t}\}_{k,t \in \mathbb{N}}$ with $\text{vol}_{k,t} : \mathbb{D}_k \times \mathbb{W}_k^t \rightarrow \mathbb{N}^t$ such that

$$\text{vol}_{k,t}(\mathbf{D}, w_1, \dots, w_t) = \left(\left(|\mathbf{D}|_w \right)_{\mathbf{D} \in \mathbf{D}(w_1)}, \dots, \left(|\mathbf{D}|_w \right)_{\mathbf{D} \in \mathbf{D}(w_t)} \right).$$

- the *total volume pattern* is the function family $\text{tvol} = \{\text{tvol}_{k,t}\}_{k,t \in \mathbb{N}}$ with $\text{tvol}_{k,t} : \mathbb{D}_k \times \mathbb{W}_k^t \rightarrow \mathbb{N}^t$ such that
- $$\text{tvol}_{k,t}(\mathbf{D}, w_1, \dots, w_t) = \left(\sum_{\mathbf{D} \in \mathbf{D}(w_1)} |\mathbf{D}|_w, \dots, \sum_{\mathbf{D} \in \mathbf{D}(w_t)} |\mathbf{D}|_w \right).$$

Each operation of an ESA (e.g., setup, query) generates leakage which is the *direct product* of one or more leakage patterns.

We say that a leakage pattern is *atomic* if it reveals information about each individual matching document. For example, rid and vol are atomic whereas tvol is not. We say that a leakage pattern is *volumetric* if it reveals size information about the matching documents. For example, vol and tvol are volumetric. Attacks that rely on volumetric leakage are particularly interesting because they apply to almost all constructions, including ORAM-based constructions.⁴

Security. In encrypted search we consider *persistent* and *snapshot* adversaries. A persistent adversary receives: (1) the encrypted data; and (2) the transcripts of the interaction between the client and the server when a query is made. A snapshot adversary, on the other hand, only receives the encrypted data after a query has been executed.

The security of ESAs can be formalized using “leakage-parameterized” definitions following [19], [18]. In this framework, a construction is proven secure with respect to a security

⁴Note that different schemes have different leakage profiles but, until recently, all known constructions leaked one of the patterns described above. The total volume pattern in particular is very difficult to suppress and is leaked by all known constructions except for the constructions recently proposed in [44].

definition that is parameterized with a specific leakage profile. Leakage-parameterized definitions for persistent adversaries were given in [19], [18] and for snapshot adversaries in [3].⁵ We recall these definitions here informally and refer the reader to [19], [18], [3] for the formal definitions.

Definition II.1 (Security vs. persistent adversary (Informal)). *Let $\Lambda = (\mathcal{L}_S, \mathcal{L}_Q) = (\text{patt}_1, \text{patt}_2)$ be a leakage profile. An encrypted search algorithm ESA is Λ -secure if there exists a PPT simulator that, given $\text{patt}_1(\mathbf{D})$ for an adversarially-chosen document collection \mathbf{D} and $\text{patt}_2(\mathbf{D}, q_1, \dots, q_t)$ for adaptively-chosen queries q_i , can simulate the view of any PPT adversary. Here, the view includes the encrypted data collection and the transcript of the queries.*

Known leakage profiles. There are many ways to design ESAs and each one provides a tradeoff between leakage and efficiency. Here, we summarize some of the most common leakage profiles and refer the reader to Appendix A for an overview of how these profiles can be achieved (and their cost). For ease of exposition, we will ignore setup leakage in this work and just denote it by \star . This is justified since none of our attacks rely on it and, moreover, there are no known attacks that leverage it. As we show in Appendix A, there are structured ESAs with leakage profiles,

$$\Lambda_{\text{BSL}} = (\mathcal{L}_S, \mathcal{L}_Q) = (\star, (\text{qeq}, \text{rid}, \text{vol}))$$

and

$$\Lambda_{\text{OPQ}} = (\mathcal{L}_S, \mathcal{L}_Q) = (\star, (\text{qeq}, \text{tvol})),$$

and oblivious ESAs with leakage profiles,

$$\Lambda_{\text{SMI}} = (\mathcal{L}_S, \mathcal{L}_Q) = (\star, (\text{rlen}, \text{rid}, \text{vol}))$$

and

$$\Lambda_{\text{FLL}} = (\mathcal{L}_S, \mathcal{L}_Q) = (\star, (\text{rlen}, \text{tvol})).$$

Adversarial model. As discussed in Section I, we consider two kinds of attacks each of which is carried out by different adversaries. Known-data attacks (i.e., leakage-abuse attacks) are carried out by a passive adversary that: (1) observes all query operations and therefore sees the query leakage; (2) knows a fraction of the client’s data; and (3) knows the universe of keywords from which the queries are drawn. Chosen-data attacks (i.e., injection attacks) are carried out by an active adversary that can add arbitrary documents either adaptively (i.e., as a function of previously-observed search results and/or leakage) or non-adaptively.⁶

III. VOLUMETRIC KNOWN-DATA ATTACKS

In this section we present four new known-data attacks, three of which are volumetric. The first is *volume analysis* which exploits the total volume pattern tvol . The second attack is *selective volume analysis* which exploits the total volume and response length patterns (tvol, rlen). The last two are concrete instantiations of an attack framework we refer to as *subgraph* attacks. The first instantiation, the *volumetric*

⁵Even though parameterized definitions were introduced in the context of SSE and STE, they can be (and have been) applied to other primitives, including to FHE-, PPE-, ORAM- and FE-based solutions.

⁶Note that for chosen-data attacks, one assumes a trivial client that accepts all injected documents without any filtering.

subgraph attack, exploits the volume pattern vol . The second instantiation, the *identifier subgraph attack* exploits the identifier pattern rid .

Remark on queries vs. keywords. For ease of exposition, we will use the term *query* when referring to a keyword that is unknown to the attacker. In particular, given leakage $\text{patt}(\mathbf{D}, q_1, \dots, q_t)$ on a sequence of t queries, the adversary’s goal will be to match each query q_i to a keyword $w \in \mathbb{W}$.

A. Volume Analysis

Volume analysis exploits the total volume pattern. It can be viewed as a volume-based analogue of frequency analysis. Given the total volume pattern, it maps every unknown query to the keyword in the known dataset that has the *closest* total volume. The intuition here is that keywords usually belong to documents with different volumes so the sum of these volumes can be used as a unique signature for a keyword.

Volume analysis takes as auxiliary input a known dataset $\tilde{\mathbf{D}}$ and as leakage

$$\text{tvol}(\mathbf{D}, q_1, \dots, q_t) = (v_1, \dots, v_t)$$

where $v_i = \sum_{\mathbf{D} \in \mathbf{D}(q_i)} |\mathbf{D}|_w$. The attack then maps the i th query q_i to the keyword $w \in \mathbb{W}$ that has the closest *known* volume to v_i . More precisely, the attack maps q_i to

$$\text{argmax}_{w \in \mathbb{W}} \left\{ f(w) : f(w) \leq v_i \right\},$$

where f is the function

$$f(w) = \sum_{\tilde{\mathbf{D}} \in \tilde{\mathbf{D}}(w)} |\tilde{\mathbf{D}}|_w,$$

that maps each keyword to its known volume. The pseudo-code of the attack appears in the full version of this work.

Efficiency. The attack runs in $O(t \cdot \#\mathbb{W})$ time.

In the full version of this work, we provide an extension of the *volume analysis* attack we call *selective volume analysis*. It exploits the response length pattern in addition to the total volume pattern.

B. Subgraph Attacks

In this section, we present our subgraph attack framework.

Overview. Contrary to previous attacks, our subgraph attacks exploit leakage patterns that reveal information on each matching document. This includes the volume pattern and the response identity pattern and we refer to the information revealed about a document as its handle. At a high level, the attack models the leakage pattern and the known dataset as bi-partite graphs. For the leakage graph, the top vertices are queries and the bottom vertices are handles, e.g., the document identifiers or the volumes. An edge is added between a query and a handle if and only if the handle is part of that query’s observed leakage. For the known-data graph, the top vertices are keywords and the bottom vertices are handles again, e.g., the document identifiers or volumes. Similarly, an edge exists between a keyword and a handle if and only if the handle corresponds to the keyword, e.g., if the document contains the keyword or if the keyword has that volume. Note that

the bottom vertices are the same in both graphs. The attack can be thought of as a subgraph mapping problem in which the adversary’s goal is to map the known-data graph into the leakage graph by leveraging the edge distribution of the graphs. As an example, if in each graph there is only one top vertex that is adjacent to a given set of bottom vertices, then it is very likely that the vertices represent the same keyword. The attack builds on this idea and a series of filtering steps. Due to space constraints, the pseudo-code appears in the full version of this work.

Details. Subgraph attacks are query-recovery attacks with known-data that can exploit any *atomic* leakage pattern; that is, any pattern that reveals a function of each matching document. Formally, a pattern patt is atomic if there exists a function $h : \mathbf{D} \rightarrow \mathbb{Y}$ such that

$$\text{patt}(\mathbf{D}, q_1, \dots, q_t) = (L_1, \dots, L_t)$$

where, for all $i \in [t]$, L_i is a tuple $(h(\mathbf{D}))_{\mathbf{D} \in \mathbf{D}(q_i)}$. Atomic patterns are relatively common and include the volume pattern vol and the identifier pattern rid . In the case of the volume pattern, h is the function $|\cdot|_w$. In the case of the identifier pattern, h is the function ids . In the following, we refer to the value $h(\mathbf{D})$ as \mathbf{D} ’s *handle*. We stress that just because a subgraph attack can be defined with respect to any atomic leakage pattern, it does not necessarily mean that it will be successful against that pattern. Its accuracy has to be verified experimentally, as we do in Section V. The attack works as follows.

Bipartite graphs. The attack takes as input an auxiliary data set $\tilde{\mathbf{D}} \subseteq \mathbf{D}$ and query leakage (L_1, \dots, L_t) defined as above. It starts by creating two bipartite graphs $\tilde{\mathbf{G}} = ((\tilde{\mathbf{L}}, \mathbb{W}), \tilde{\mathbf{E}})$ and $\mathbf{G} = ((\mathbf{L}, \mathbf{Q}), \mathbf{E})$ from the auxiliary data and the leakage, respectively. The vertex set $\tilde{\mathbf{L}}$ is composed of the handles of the known documents; that is,

$$\tilde{\mathbf{L}} = \left\{ h(\tilde{\mathbf{D}}) \right\}_{\tilde{\mathbf{D}} \in \tilde{\mathbf{D}}}.$$

For all keywords $w \in \mathbb{W}$ and documents $\tilde{\mathbf{D}} \in \tilde{\mathbf{D}}$, $\tilde{\mathbf{E}}$ includes an edge $(w, h(\tilde{\mathbf{D}}))$ if $w \in \tilde{\mathbf{D}}$. The second bipartite graph $\mathbf{G} = ((\mathbf{L}, \mathbf{Q}), \mathbf{E})$ is constructed as follows. The vertex set \mathbf{L} is composed of the observed document handles; that is,

$$\mathbf{L} = \left\{ h(\mathbf{D}) \right\}_{\mathbf{D} \in \bigcup_{i=1}^t \mathbf{D}(q_i)}.$$

The vertex set $\mathbf{Q} = (q_1, \dots, q_t)$ is composed of the queries q_1 through q_t . The edges \mathbf{E} are created using the observed leakage by adding an edge $(q_j, h(\mathbf{D}))$ if $h(\mathbf{D}) \in L_j$, for all $j \in [t]$ and $\mathbf{D} \in \bigcup_{i=1}^t \mathbf{D}(q_i)$.

Refinement. Given these two graphs, the algorithm’s goal is to match each q_i in \mathbf{Q} to some w in \mathbb{W} . For each q_i , the attack will build a set of *potential* keyword matches $S_i \subseteq \mathbb{W}$ and keep refining it using several filtering steps. We will denote by $S_i^{(j)}$ the set of q_i ’s potential keyword matches after the j th refinement step. Let $S_i^{(0)} = \mathbb{W}$ and let w_i^* be q_i ’s correct match. The first filtering step is based on the observation that w_i^* ’s matching documents in $\tilde{\mathbf{D}}$ have to be a subset of q_i ’s

matching documents in \mathbf{D} . More formally, we have

$$S_i^{(1)} = \left\{ w \in \mathbb{W} : N_{\tilde{G}}(w) \subseteq N_G(q_i) \right\},$$

where $N_{\tilde{G}}(w)$ and $N_G(w)$ are the neighbors of w and q_i in \tilde{G} and G , respectively. The second filtering step is based on the observation that the selectivity of w^* in $\tilde{\mathbf{D}}$ should be a δ fraction of its selectivity in \mathbf{D} , where δ is the known-data rate. Based on this we have

$$S_i^{(2)} = \left\{ w \in S_i^{(1)} : \#N_{\tilde{G}}(w) \geq \delta \cdot \#N_G(q_i) - \varepsilon \right\},$$

where ε is an error parameter we set experimentally.

Cross filtering. The next filtering step is optional and can be used only if the function $h : \mathbf{D} \rightarrow Y$ is a bijection and if the adversary knows a large enough fraction of \mathbf{D} . The observation we rely on is that the correct keyword in a potential set $S_i^{(2)}$ must be contained in all the documents in the set $h^{-1}(L_i)$. As a concrete example, consider the case where h is the document ID function id . In this case, the observation above translates to the fact that the correct keyword w_i^* in $S_i^{(2)}$ must be contained in all the documents $(D_\alpha)_{\alpha \in L_i}$. This follows from the correctness of the ESA scheme. More formally, we have that

$$w_i^* \in S_i^{(2)} \cap \left(\bigcap_{D \in h^{-1}(L_i)} D \right).$$

Notice, however, that we may not be able to compute the above subset because it requires us to invert h and recover all the documents that matched the query. In particular, we can only invert h if our auxiliary dataset is complete, i.e., $\delta = 1$. Nevertheless, we can approximate the set h^{-1} even when $\delta < 1$ as follows. We use the set $\tilde{L} \cap L_i$ which is the subset of observed handles of documents that we know. We then compute

$$S_i^{(3)} = S_i^{(2)} \cap \left(\bigcap_{\alpha \in \tilde{L} \cap L_i} \tilde{D}_\alpha \right).$$

At the end of this step, if the size of the set $S_i^{(3)}$ is equal to 1, then this means that we found a match for the i th query q_i , and therefore update the map α accordingly.

Iterative elimination. The final step of the attack relies on the observation that if some w is the correct match for a query q_i then it cannot be the correct match for another query q_j , where $i \neq j$. In other words, if w is the unique element of some potential set $S_i^{(\ell)}$, for $\ell \in [4, t+3]$, then w cannot be the matching keyword in some other potential set $S_j^{(\ell)}$ and, therefore, we can remove it from $S_j^{(\ell)}$. If this removal leads to $S_j^{(\ell)}$ having a single element, then we can in turn remove that element from other potential sets. This process keeps going until the potential sets stabilize. Note that while the algorithm will terminate it may not find matches for all q_i .

Efficiency. The first filtering step is $O(t \cdot \#\mathbb{W})$, the cross filtering step is

$$O\left(t + \sum_{i=1}^t \sum_{D \in \mathbf{D}(q_i)} \#\mathbf{D}\right)$$

and the iterative elimination step is $O(t^2)$. In total, the algorithm runs in time

$$O\left(t \cdot \#\mathbb{W} + \sum_{i=1}^t \sum_{D \in \mathbf{D}(q_i)} \#\mathbf{D}\right).$$

The volumetric subgraph attack. As discussed above, subgraph attacks can exploit any atomic leakage pattern by properly instantiating the handle function h . The volumetric subgraph attack results from instantiating h with the function $|\cdot|_w$ which maps each document to its volume. Note that $|\cdot|_w$ is not bijective so this instantiation cannot use the cross filtering step.

The identifier subgraph attack. Our subgraph framework can also be used to exploit the identifier pattern by instantiating h with the function ids that maps keywords to the identifiers of the documents that contain it. Note that because ids is bijective, we can use the cross filtering step.

IV. VOLUMETRIC INJECTION ATTACKS

Injection attacks were first proposed by Cash et al. [13] in the context of the PPE-based ShadowCrypt and Mimesis [36], [49] systems. The first injection attacks on structured ESAs were described by Zhang, Katz and Papamanthou in [75]. In that work, two attacks are described, each of which exploits the identifier pattern. In this section, we describe new volumetric injection attacks. In particular, our attacks exploit the *total volume pattern* and, therefore, can even be used against oblivious ESAs.

A. The Decoding Attack

We now describe our decoding attack. At a high-level, the attack first observes user queries and their associated volume. It then uses this information to create documents with carefully-chosen sizes for the purpose of injection. More precisely, the sizes of an injected documents are chosen so that the volume of each keyword becomes unique. This unique volume can then be used as a signature when the keyword is queried again. Due to space constraints, the pseudo-code of this attack appears in the full version of this work.

The decoding attack works in two phases: baseline and recovery. In the *baseline* phase, the adversary waits until it has observed the total volumes for all keywords in \mathbb{W} .⁷ During the recovery phase, the adversary observes an additional sequence of $t \geq 1$ client queries $\mathbf{q} = (q_1, \dots, q_t)$ with total volumes $\mathbf{v} = (v_1, \dots, v_t)$. The attack will recover all queries in \mathbf{q} . We now describe each phase in more detail.

Baseline. During the baseline phase, the adversary observes queries until it holds the volumes $\mathbf{b} = (b_1, \dots, b_m)$ of all the keywords $w_1, \dots, w_\ell \in \mathbb{W}$. It then creates a set of documents to inject as follows. It first computes an *offset* γ defined as

$$\gamma = \min \left\{ \gamma \in \mathbb{N} : \forall i, j \in [m], \gamma \nmid b_i - b_j \right\}.$$

For all keywords $w_i \in \mathbb{W}$, the adversary injects a document with volume $i \cdot \gamma$ filled with w_i . Intuitively, this step increases w_i 's volume by $i \cdot \gamma$.

⁷Note that the attack works similarly if the user only queries a strict subset of \mathbb{W} . This will lead to fewer injections.

Recovery. During the recovery phase, the adversary will observe volumes $\mathbf{v} = (v_1, \dots, v_t)$ on queries $\mathbf{q} = (q_1, \dots, q_t)$. Note that for all $i \in [t]$, the volume of q_i can be written as

$$v_i = b_j + u \cdot \gamma$$

for some $j \in [m]$ and $u \in [\ell]$. The adversary’s goal now is to map q_i to some keyword $w \in \mathbb{W} = (w_1, \dots, w_\ell)$. It does this as follows. It checks if there exists a $u \in [\ell]$ such that $v_i - u \cdot \gamma$ is equal to one of the baseline volumes (b_1, \dots, b_m) . If this is the case, then the adversary maps q_i to keyword w_u . Note that there can be at most a single baseline volume that satisfies this condition. To see why, suppose there exists two baseline volumes $\beta_1 \neq \beta_2$ and two values $z \neq z'$ in $[\ell]$ such that

$$v_i - z \cdot \gamma = \beta_1 \quad \text{and} \quad v_i - z' \cdot \gamma = \beta_2.$$

But this implies that $\gamma \cdot (z' - z) = \beta_1 - \beta_2$ which is a contradiction since $\gamma \nmid b_i - b_j$, for all $i, j \in [m]$.

Correctness & efficiency. The decoding attack recovers all queries in \mathbf{q} as long as the user did not add any documents of its own. As described, the attack recovers all queries in \mathbf{q} by injecting $\#\mathbb{W}$ documents with a total volume of $O(\gamma \cdot \#\mathbb{W}^2)$. Note, however, that the attack can be made a lot more efficient if the adversary is only interested in recovering queries within some target set $T \subset \mathbb{W}$. In this case, in the baseline phase the adversary still needs to gather the baseline volumes for queries in \mathbb{W} , but only needs to inject documents for keywords in T . With this modification, the attack recovers queries in $\mathbf{q} \cap T$ by injecting T documents with a total volume of $O(\gamma \cdot \#T^2)$. Our evaluation demonstrates that the offset γ for different subsets of the Enron dataset is between 4 and 16 KBytes depending on the query selectivity (see Appendix D).

B. The Binary Search Attack

Overview. Contrary to the decoding attack where the adversary recovers all queries, the *binary search attack* is a targeted attack in which the adversary’s goal is to recover one specific query. At a high-level, the attack first observes user queries and their associated volume. It then uses this information to create a document that contains half the keyword space and that has a carefully-chosen size. When this document is injected, it modifies the volume of half of the keywords (the ones contained in the injected document) in a unique manner that is detectable. The adversary then observes more queries and uses the presence or absence of the unique volume to infer which half of the keyword space the target query is in. The attack then recurs on that half. The base case is a single-element set which it outputs as the keyword. Due to space constraints, the pseudo-code of this attack appears in the full version of this work.

Details. The binary search attack works in three phases: baseline, targeting and recovery. In the *baseline* phase, the adversary waits until it has observed the total volumes for all keywords in \mathbb{W} . During the *targeting* phase the adversary observes more client queries until it decides on a query q_0 , with total volume v_0 , that it wishes to target. In the *recovery* phase, the adversary observes an additional sequence of $t > \log \#\mathbb{W}$ client queries $\mathbf{q} = (q_1, \dots, q_t)$ with volumes $\mathbf{v} = (v_1, \dots, v_t)$ that it will use to recover its target q_0 . We now describe each phase in more detail.

Baseline. During the baseline phase, the adversary observes queries until it holds the volumes $\mathbf{b}^{(0)} = (b_1, \dots, b_m)$ for each keyword $w_1, \dots, w_\ell \in \mathbb{W}^{(0)}$, where $\mathbb{W}^{(0)} = \mathbb{W}$.

Targeting. During the targeting phase, the adversary observes queries until it decides on a target query q_0 with total volume v_0 . It then partitions $\mathbb{W}^{(0)}$ into two equal-sized sets, $\mathbb{W}_0^{(0)}$ and $\mathbb{W}_1^{(0)}$ and computes an offset

$$\gamma = \min \left\{ \gamma \in \mathbb{N} : \forall j \in [m], \gamma \neq |v_0 - b_j^{(0)}| \wedge \gamma \geq |\mathbb{W}_1^{(0)}|_w \right\}.$$

The adversary then injects a document with volume γ that contains all the keywords in $\mathbb{W}_1^{(0)}$.

Recovery. In the first round of the recovery phase, the adversary observes the total volume v_1 for query q_1 . It then uses γ to decide on one of three cases:

- if $v_1 = v_0 + \gamma$ then the adversary concludes that $q_1 = q_0$ and that $q_0 \in \mathbb{W}_1^{(0)}$;
- if $v_1 = v_0$ then the adversary concludes that $q_1 = q_0$ and that $q_0 \in \mathbb{W}_0^{(0)}$;
- if $v_1 \neq v_0$ and $v_1 \neq v_0 + \gamma$ then the adversary concludes that $q_1 \neq q_0$.

If $q_0 \in \mathbb{W}_1^{(0)}$ the adversary sets $\mathbb{W}^{(1)} = \mathbb{W}_1^{(0)}$. If $q_0 \in \mathbb{W}_0^{(0)}$ it sets $\mathbb{W}^{(1)} = \mathbb{W}_0^{(0)}$. For both these cases, before moving to the next round, the adversary re-injects a document with volume γ such that

$$\gamma = \min \left\{ \gamma \in \mathbb{N} : \forall j \in [m], \gamma \neq |v_1 - b_j^{(1)}| \wedge \gamma \geq |\mathbb{W}_1^{(1)}|_w \right\},$$

where $\mathbf{b}^{(1)} = (b_1^{(0)}, b_1^{(0)} + \gamma, \dots, b_m^{(0)} + \gamma)$, $\mathbb{W}_0^{(1)}$ and $\mathbb{W}_1^{(1)}$ are the two partitions of $\mathbb{W}^{(1)}$. If otherwise $q_1 \neq q_0$, the adversary moves to the next round without changing $\mathbb{W}^{(0)}$ nor injecting a new file.

Correctness & efficiency. The binary search attack will recover q_0 as long as: (1) it appears $\log \#\mathbb{W}$ times in $\mathbf{q} = (q_1, \dots, q_t)$; (2) it has a unique volume in the baseline volumes; and (3) the user did not add any document of its own. The attack needs to inject $\log \#\mathbb{W}$ files with a total volume of $\Omega(\#\mathbb{W})$. Our evaluation demonstrates that the size of the query space was the main factor that determines the total injected volume (i.e., $\gamma \simeq \#\mathbb{W}/2^u$ for all $u \leq \log \#\mathbb{W}$ for different subsets of the Enron dataset). The total injected volume was around 8KBytes (see Appendix D).

V. EMPIRICAL EVALUATION

To evaluate the effectiveness of our attacks, we implemented and evaluated them under different conditions. The results are perhaps surprising and provide a new and more nuanced perspective on the potential impact and limitations of leakage abuse attacks.

Document collections. We build our document collections using the Enron Email dataset [63]. This dataset is composed of 150 folders with a total of 520,901 files. Each folder corresponds to the email account of a single individual and is itself composed of several folders including, for example, inbox, sent, contacts, discussion threads, etc.

Starting from the entire Enron dataset, we generated different subsets that capture different settings of interest:

- *single user* (SU): is a document collection composed of one individual’s email account. For this dataset, we picked the `arnold-j` folder which is 11.6MBytes and is composed of 4,944 files. The total number of keywords is 40,363. This collection models the traditional single user ESA setting where a single client uses an ESA to encrypt and privately access its own dataset.
- *small multiple users* (S-MU): is a document collection composed of multiple email accounts. For this dataset, we picked 5 folders with a total size equal to 26MBytes. This document collection is composed of the following email accounts: `baughman-d`, `gay-r`, `heard-m`, `hendrickson-s` and `linder-e`. The dataset is composed of 9,416 files in total. The total number of keywords is 77,762. This collection models the multi-user ESA setting in which there is one party (e.g., a company) that uses an ESA to encrypt and store multiple users’ documents. Here, queries can be performed by a subset of authorized users. Note that the users were arbitrarily picked.
- *medium multiple user* (M-MU): is the same as above except that we increase the number of folders to 10 with a total size of 49MBytes. The data collection is composed of the same email accounts as above plus: `allen-p`, `buy-r`, `forney-j`, `hyvl-d` and `keiser-k`. The total number of keywords is 115,679. Similarly, the additional folders were picked arbitrarily.

The purpose of the first two collections is to see whether the effectiveness of our attacks will vary as a function of different data distributions. The third collection is used to understand if increasing the size of the dataset impacts the effectiveness of the attacks. We note that we performed evaluations which we do not report here in order to avoid redundancy. In particular, we evaluated the attacks on a larger document collection from Enron of size 106MBytes and the results were similar to the ones on M-UM. We also evaluated our attacks on a subset of the TREC 2007 Public Corpus dataset [57], an email dataset composed of 75,419 emails (including spam), and the results were similar.

Data indexing. We indexed each data collection using Apache Lucene [1]. We removed 224 stop words listed in the SnowBall list [62]. Furthermore, we used the Porter Stemming implementation of Lucene so all words that share the same stem are mapped to the same root.

Query frequency. We observed that previous works on leakage abuse attacks [38], [13] evaluated the effectiveness of their attacks on the most frequent keywords in the dataset. While this assumption might hold in some settings, it is far from clear if this a reasonable assumption in practice. In fact, it might seem that users would more often search for keywords that occur less frequently in their dataset rather than for keywords that occur frequently. With this in mind, we evaluated all of our attacks in three different settings:

- *high selectivity*: the queries are sampled from the set of keywords with the highest selectivities (i.e., that appear in the largest number of documents). We noticed that keyword selectivities in Enron are power-law distributed

(see Appendix C) which implies that high-selectivity keywords tend to have *unique* selectivities while low-selectivity keywords tend to have less unique selectivities (in our datasets it was none).

- *low selectivity*: the queries are sampled from the set of keywords with the lowest selectivities. In our datasets, all the low-selectivity keywords had selectivity 1.
- *pseudo-low selectivity*: in this case we consider low-selectivity keywords with a slightly higher selectivity than above. In particular, we consider the case where the selectivity ranges from 10 to 13. Note that these values are only examples and that pseudo-low selectivity can be defined using other values.

Size and composition of the query space. We first fix the size of the query space, \mathbb{Q} , to be 500. We then study the impact of increasing $\#\mathbb{Q}$. In particular, we increased $\#\mathbb{Q}$ up to 5000 keywords.⁸ We also consider two ways of instantiating the query space. The first consists of populating \mathbb{Q} with keywords that only exist in the known-data collection \mathbb{D} . This guarantees that all $q \in \mathbb{Q}$ must exist in \mathbb{D} . As a consequence, the attacker would know that any client query will match at least one document in \mathbb{D} . The second approach consists of populating \mathbb{Q} from keywords that exist in the client’s collection \mathbb{D} .

Experimental setting. As described above, there are several variables that can impact the effectiveness of our attacks. In our evaluation, we considered many different combinations of these variables and organized them in three main categories:

- (C1) *single keyword queries*: we consider the SU dataset and fix the size of the query space to be 500. We then vary the query selectivity and query space composition; refer to Figures 1a and 1b.
- (C2) *size of the query space*: this second category is the same as the first except that we increase the size of the query space to be 5000; refer to Figures 1c and 1d.
- (C3) *varying the datasets*: the third category is similar to the first category except that we replace the SU dataset with the S-MU and M-MU datasets; refer to Figures 1e, 1f, and Figures 1f.

All our attacks are evaluated against a query sequence \mathbf{q} of size $t = 150$. The queries are sampled uniformly at random from the corresponding keyword space \mathbb{Q} whose composition varies depending on the chosen approach (see above). In all our experiments, we start with the adversary knowing the entire client collection and then gradually decrease this knowledge until it knows only 5% of the documents (chosen uniformly at random). For each attack, we report the *recovery rate*, i.e., the number of queries recovered correctly over the total number of queries. We run all experiments 5 times and report the minimum, median and maximum of the recovery rate. All attacks are implemented in Java and the experiments were run on a MacBook 3.1 GHz Intel Core i7 with 16GBytes of RAM.

Baseline. In all our experiments we have included the *count-only* attack which is simply the Count attack [13] without the co-occurrence matrix. This depicts constructions

⁸The values we picked are similar to the ones used in existing works [38], [13].

for which we have suppressed the co-occurrence leakage, refer to Appendix A.

Overview of results. We noticed that the recovery rate of our attacks is impacted the most by the *selectivity* of the queries. In fact, it seems that evaluating known-data attacks on high- vs. low-selectivity queries leads to completely opposite conclusions.⁹ Moreover, we noticed that changing the datasets or increasing the size of the query space led to some fluctuations but the overall trends remained the same. This shows that our attacks work across different settings and different dataset compositions. We found, however, that if client queries are not in the adversary’s known dataset then the recovery rate is always low. This simply follows from the fact that there are several queries for which the adversary does not hold any part of the response. Below, we provide more detailed comments on the results of our evaluations:

- When the query space is composed of high-selectivity keywords, both $\text{Subgraph}^{\text{ID}}$ and $\text{Subgraph}^{\text{VL}}$ have a recovery rate of about 70% even with a known-data rate of 5% (see Figure 1e). We believe that the low known-data rate makes these attacks practical for high-selectivity keywords. However, if the query space is composed of low-selectivity keywords, then the recovery rate drops significantly. In fact, we found that $\text{Subgraph}^{\text{VL}}$ does not work at all while $\text{Subgraph}^{\text{ID}}$ has a recovery rate of about 20% even when $\delta = 1$; that is, with full knowledge of the client’s data. With a known-data rate of $\delta = 1/2$, $\text{Subgraph}^{\text{ID}}$ only has 10% recovery rate. It tends to 0 for known-data rates smaller than 10%.
- For VolAn and SelVolAn, our evaluation shows that both attacks work only when: (1) the query space consists of high-selectivity keywords; and (2) the adversary has a high known-data rate, often at least .85. We refer the reader to Figure 1a for an example. In the case where the query space consists of low-selectivity keywords, the recovery rate is very low: around 10% even when $\delta = 1/2$. It drops significantly when δ gets smaller.
- When the query space consists of pseudo-low-selectivity keywords (i.e., with selectivity between 10 and 13), both the VolAn and SelVolAn attacks have high recovery rate only when the known-data rate $\delta \geq .8$ (see Figure 2). As δ decreases, the recovery rate stabilizes at around 18% and starts to decrease again to 0 when $\delta \leq 0.15$. Note that this recovery rate is the highest among the three selectivity classes for these two attacks. The recovery rates of both $\text{Subgraph}^{\text{ID}}$ and $\text{Subgraph}^{\text{VL}}$ against pseudo-low-selectivity queries are slightly better than the recovery rates of SelVolAn and VolAn. $\text{Subgraph}^{\text{ID}}$ and $\text{Subgraph}^{\text{VL}}$ also do better on pseudo-low-selectivity keywords than they do on low-selectivity keywords. However, they do a lot worse than they do against high-selectivity keywords.¹⁰

Impact of known-data rates. Our evaluation demonstrates that the known-data rate has a big impact on the

⁹We recall that the experiments in [13] were done exclusively on high-selectivity keywords.

¹⁰We also conducted experiments in which the queries were sampled uniformly at random from the entire keyword space. The results were similar to the low-selectivity case (the $\text{Subgraph}^{\text{ID}}$ recovery rate was slightly higher though).

recovery rate. For all attacks, the larger δ is the higher the recovery rate. This is natural since all the attacks exploit some correlation between documents and keywords. As discussed in Section I, choosing what constitutes a “safe” known-data rate requires more cryptanalysis but what we can say is that a known-data rate of $\delta < 0.05$ could be considered safe against the attacks in this work. If a dataset is composed of one million documents, then even when $\delta = 0.05$ the adversary would still need to know 50,000 documents.

Impact of query selectivity. Our evaluation also shows that query selectivity is an important factor in the query recovery rate. In particular, higher query selectivity implies that the keyword is present in a higher number of documents which increases the accuracy of the attacks. For example, both volume analysis and selective volume analysis can build stronger signatures on high-selectivity queries and therefore obtain better recovery rates when $\delta \approx 1$. Similarly, on high-selectivity queries, the subgraph attacks have richer neighbor sets (i.e., more unique distribution of edges). Of course, whether a user queries high- or low-selectivity keywords is application specific but we believe that varying this parameter is important when evaluating the recovery rate of an attack.

Impact of the composition of the query space. We considered two cases: *entire* or *partial* query space. For the former, the query space consists of the keywords in the entire dataset, while for the latter the query space is composed of the keywords in the known dataset. We believe that the partial case is more realistic for unstructured data but we included both. For structured data (e.g., a medical database) the values of all attributes are often public (e.g., the possible ages, illness codes etc.) so entire or partial knowledge of the data does not matter in this case.

Impact of indexing. Indexing is a parameter that we, unfortunately, did not investigate much in this work. We performed our evaluations using an “aggressive” indexing strategy where we used porter stemming before indexing *all* the keywords in the dataset. If the adversary is not aware of the stemming algorithm, it could end up with a completely different query space than the user which would impact the recovery rate.

Impact of the composition of the dataset. Among all the parameters, we believe dataset composition is the one that has the least impact on the recovery rate. Our results show that any reasonable dataset composition of the Enron datasets or the TREC 2007 Public Corpus dataset leads to similar results.

Results on chosen-data attacks. Since our injection attacks always succeed we do not evaluate their success rate empirically. We report, however, that in order to succeed we set the size of the keyword universe \mathbb{W} to 500. Also, for the Decoding attack, one has to inject between 4 and 16 KBytes to recover one keyword depending on the type of the document collection and the keyword selectivity. For the Binary attack, the adversary has to inject around 8 KBytes for all document collections and this holds independently of the selectivity of the keyword. We provide more details about our evaluation in Appendix D.

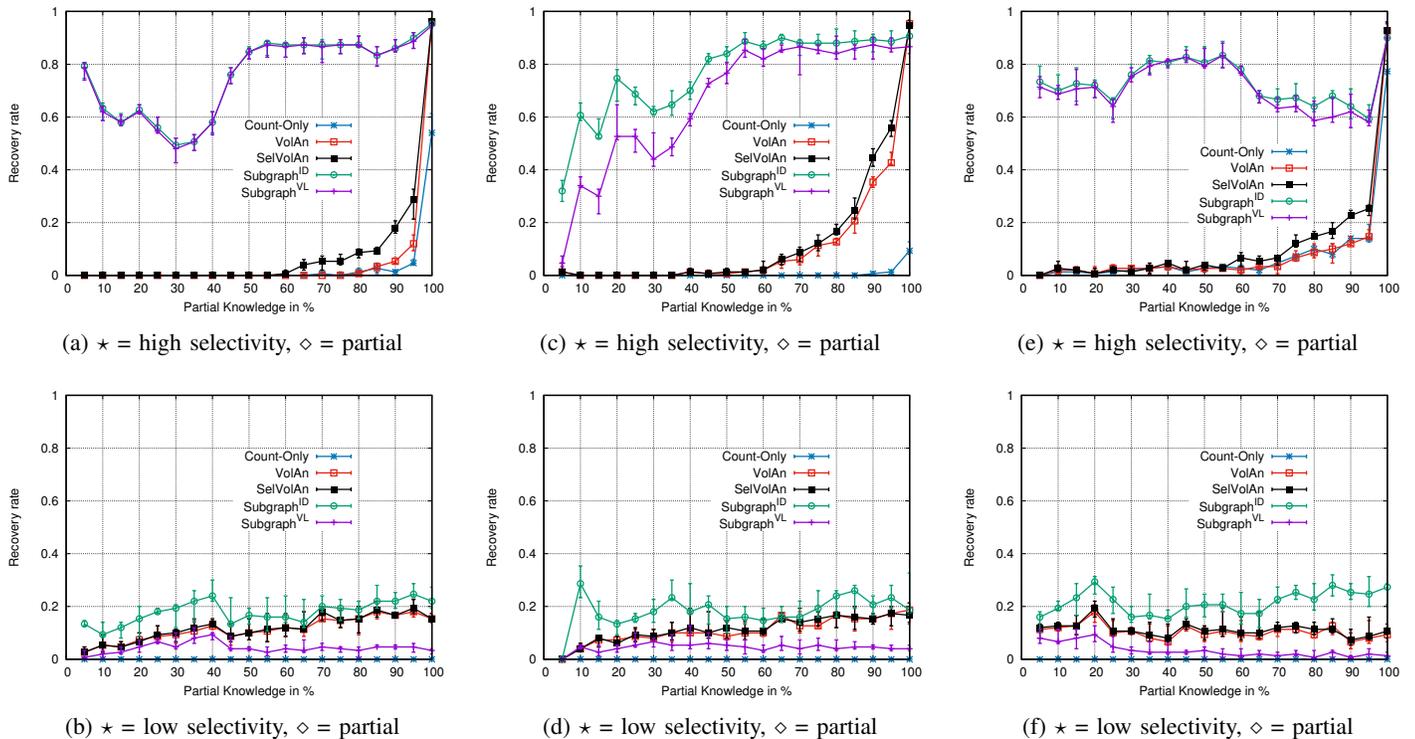


Figure 1: SU dataset (Figures 1a and 1b: 150 keywords queried u.a.r. from 500 \star keywords in the \diamond dataset. SU dataset (Figures 1c and 1d): 150 keywords queried u.a.r. from 5000 \star keywords in the \diamond dataset. M-MU dataset (Figures 1e and 1f): 150 keywords queried u.a.r. from 500 \star keywords in the \diamond dataset.

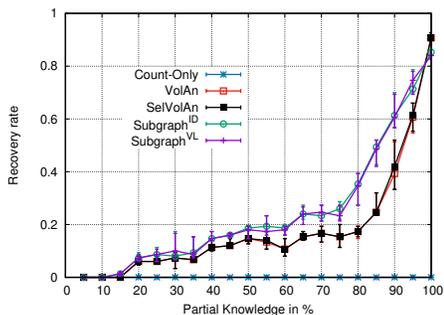


Figure 2: 150 keywords queried u.a.r. from 500 pseudo-low selectivity keywords in SU.

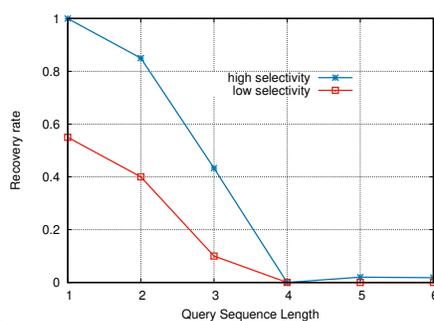


Figure 3: PBS brute-force attack.

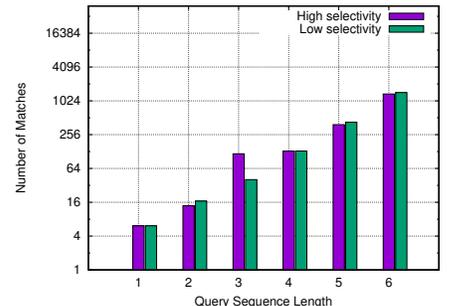


Figure 4: Number of matches in PBS brute-force attack.

VI. TAKEAWAYS

In this work we revisited leakage abuse and injection attacks against ESAs. In particular, we argued that the oft-cited IKK and Count attacks are mostly of theoretical interest due to the following limitations and assumptions:

- *high known-data rates*: both the IKK and Count attack require high known-data rates to achieve reasonable recovery rates and it is not clear whether such rates are realistic;
- *known queries*: in addition to relying on known-data, the Count v.1 attack also relies on known queries;
- *suppressable leakage*: the IKK and Count attacks rely on the co-occurrence pattern which can be easily hidden at the cost of additional storage using our OPQ construction.
- *experimental evaluation*: the experimental evaluations of the IKK and Count attack were not conducted in all settings of interest. This includes, for example, querying

low-selectivity keywords or keywords that are not in the adversary's known dataset.

New attacks. To address these limitations, we introduced four new known-data attacks and two new injection attacks. We believe our known-data attacks are of practical interest since they work with low known-data rates and do not rely on any known queries. Most surprisingly, our attacks make use of only volumetric leakage and therefore apply, not only to structured ESAs, but also to oblivious ESAs.

We implemented our attacks and evaluated them empirically in various settings and using different kinds of queries. We hope that our study provides useful insights that may help the community better understand the real-world impact of leakage abuse attacks. In the following, we list our main takeaways, sometimes referencing the constructions described in Section II and Appendix A:

- None of the attacks worked against low-selectivity or pseudo-low-selectivity keywords;
- When querying high-selectivity keywords, the rid and vol patterns (which are respectively leaked by the BSL and SMI constructions) can be exploited by our Subgraph attacks with very low known-data rates (as low as 5%). Note that high recovery rates are maintained across different settings. In addition, we believe our attacks would do even better on larger data collections. For example, on collections on the order of gigabytes we estimate that relatively high recovery rates could still be achieved with known-data rates as low as 1%. We conclude that, for high-selectivity keywords, there are *practical* attacks against leakage profiles that include rid and vol.
- When querying high-selectivity keywords, the total volume pattern (which is leaked by OPQ and FLL) seems resistant to our attacks for $\delta \leq .8$.
- The total volume pattern tvol can be successfully attacked with our injection attacks (though in the case of the Binary Search attack only if the target query has a unique total volume).
- Structured and oblivious ESAs seem to provide the same level of security against both our known-data (i.e., leakage abuse) and chosen-data (i.e., injection) attacks.

VII. COUNTERMEASURES

Our study revealed two settings in which our attacks could be practical. The first is using our Subgraph attacks to exploit the rid and vol patterns on high-selectivity keywords and the second is using our volumetric injection attacks to exploit the total volume pattern. For all other settings, we do not believe any countermeasures are required though they are certainly available.

High-selectivity keywords. For high-selectivity keywords, one should simply use a scheme that does not leak rid or vol like the PBS construction of [45] (see Appendix A for a brief overview of PBS) or the OPQ or FLL constructions described in Appendix A. These schemes have the following leakage profiles

$$\Lambda_{\text{PBS}} = (\star, (\text{req}, \text{svol})) \quad \text{and} \quad \Lambda_{\text{OPQ}} = (\star, (\text{keq}, \text{tvol}))$$

and,

$$\Lambda_{\text{FLL}} = (\star, (\text{rlen}, \text{tvol})),$$

where svol is the *sequence volume pattern*,

$$\text{svol}(\mathbf{D}, w_1, \dots, w_t) = \sum_{i=1}^t \sum_{D \in \mathbf{D}(w)} |D|_{\mathbf{w}}.$$

Our experiments suggest that for $\delta \leq .8$ either OPQ or FLL can be used but that for $\delta > .8$ one should use PBS. Note that the sequence volume pattern seems to be a very “low leakage” pattern in the sense that even a “brute-force” attack that simply tries to match keywords to the sequence volume leakage does not work on our dataset. We provide more details below.

Brute force. We assume the adversary has full knowledge of the client’s data, i.e., $\delta = 1$. Given the sequence volume pattern of a query sequence of length λ drawn uniformly at random from a set of 500 either low- or high-selectivity keywords, the attack finds all possible sequences of

λ keywords that have sequence volume leakage equal to the given/observed leakage. If there is only a single such sequence, the attack returns it as its output otherwise it fails.

Note that this is the best possible attack against the sequence volume pattern (not taking efficiency into account). We define the attack’s success rate as the fraction of its output that is correct; that is, the number of keywords in its output sequence (assuming it outputs a sequence) that are indeed in the client sequence over the size of the sequence. We ran the attack for high- and low-selectivity keywords, with λ ranging from 1 to 6 and found that when $\lambda \geq 4$ the attack stopped working. More precisely, in the case of low-selectivity keywords its success rate was 0 and in the case of high-selectivity keywords it was 0.02. Figure 3 describes these results in detail.

Note that the success rate of the brute-force attack does not capture partial knowledge since it only accounts for the case where the attack finds a single “matching” sequence. For example, the success rate could be 0 even though the attack found just 2 matching sequences. To address this, we ran an additional experiment that computes the number of matching sequences found by the attack. The results are described in Figure 4 which shows that the average number of matches grows exponentially as a function of the sequence length (this holds independently of the selectivity of the keywords in the sequence). For example, for $\lambda = 5$, there are 385 matching sequences even when the keyword space is as small as 100.¹¹ Notice that increasing the number of keywords will significantly increase the number of matches.

On the cost of OPQ, FLL, and PBS. The cost of OPQ and FLL is described in Section II so we focus on PBS [45]. This construction uses a response-hiding multi-map encryption scheme so the query overhead is comparable to BSL and OPQ. However, it does introduce additional latency which is not incurred by previous constructions. That is, while PBS’s query algorithm is very efficient (for the server) the user might need to wait until it makes additional queries to retrieve the complete response of a query. The latency can be tuned to be smaller but this requires some knowledge about the underlying dataset. Note however that for applications where the user issues a batch of queries all at once, PBS does not incur much latency. In addition, we showed in the previous paragraph that batch sizes as small as five can be safe. Additionally, the user could make dummy queries instead of batching queries, in which case latency will not be a factor.

Volumetric attacks. Recently, Kamara and Moataz [44] proposed the first volume-hiding encrypted multi-maps that do not rely on naïve padding. The two constructions, VLH and AVLH, achieve different trade-offs between storage efficiency, query efficiency and lossiness and can be used to protect against volumetric attacks. We refer the reader to Appendix A for more details.

Note that VLH is lossy and can return false negatives. The lossiness/false negatives, however, are tunable and can be traded-off for additional storage. Depending on the choice of parameters, VLH can be as efficient as BSL and OPQ.

¹¹For this experiment we had to reduce the size of the keyword space from 500 to 100 keywords, because the former results in an extremely large number of sequences to check, i.e., $\binom{500}{6} \approx 2 \cdot 10^{13}$

AVLH, on the other hand, has no false negatives but its query complexity is the maximum response length (over the keyword space). The storage overhead of AVLH is comparable to the overhead of BSL. Recent volume-hiding results by Patel, Persiano, Yeo and Yung [61] show how to get even better query complexity and storage overhead.

Another approach to protecting against volumetric attacks is to use padding techniques. Naïve padding (adding dummy values to ensure the volume of every query response is of the same size) will protect against volumetric attacks but incurs a large storage overhead. More efficient padding techniques were proposed by Bost and Fouque [12]. While these techniques seem to make attacks harder, it is not clear if they can completely mitigate them.

ACKNOWLEDGMENT

We would like to thank the authors of [13] for answering our questions about the Count attack, for sharing the code of the Count attack with us, and for useful feedback on this work.

REFERENCES

- [1] “Apache lucene,” 1999, <http://lucene.apache.org>.
- [2] R. Agrawal, J. Kiernan, R. Srikant, and Y. Xu, “Order preserving encryption for numeric data,” in *ACM SIGMOD International Conference on Management of Data*, 2004, pp. 563–574.
- [3] G. Amjad, S. Kamara, and T. Moataz, “Breach-resistant structured encryption,” *Proceedings on Privacy Enhancing Technologies*, vol. 2019, no. 1, pp. 245–265, 2019.
- [4] G. Asharov, M. Naor, G. Segev, and I. Shahaf, “Searchable symmetric encryption: Optimal locality in linear space via two-dimensional balanced allocations,” in *ACM Symposium on Theory of Computing (STOC ’16)*, ser. STOC ’16. New York, NY, USA: ACM, 2016, pp. 1101–1114. [Online]. Available: <http://doi.acm.org/10.1145/2897518.2897562>
- [5] G. Asharov, T. H. Chan, K. Nayak, R. Pass, L. Ren, and E. Shi, “Oblivious computation with data locality,” *ePrint IACR*, 2017.
- [6] J.-P. Aumasson, “Cryptanalysis vs. reality,” in *Black Hat (Abu Dhabi)*, 2011.
- [7] M. Bellare, A. Boldyreva, and A. O’Neill, “Deterministic and efficiently searchable encryption,” in *Advances in Cryptology – CRYPTO ’07*, ser. Lecture Notes in Computer Science, A. Menezes, Ed. Springer, 2007, pp. 535–552.
- [8] A. Boldyreva, N. Chenette, Y. Lee, and A. O’Neill, “Order-preserving symmetric encryption,” in *Advances in Cryptology - EUROCRYPT 2009*, 2009, pp. 224–241.
- [9] D. Boneh, A. Sahai, and B. Waters, “Functional encryption: Definitions and challenges,” in *Theory of Cryptography Conference (TCC ’11)*, ser. Lecture Notes in Computer Science, vol. 6597. Springer, 2011, pp. 253–273.
- [10] R. Bost, “Sophos - forward secure searchable encryption,” in *ACM Conference on Computer and Communications Security (CCS ’16)*, 2016.
- [11] R. Bost, B. Minaud, and O. Ohrimenko, “Forward and backward private searchable encryption from constrained cryptographic primitives,” in *ACM Conference on Computer and Communications Security (CCS ’17)*, 2017.
- [12] R. Bost and P.-A. Fouque, “Thwarting leakage abuse attacks against searchable encryption – a formal approach and applications to database padding,” *IACR Cryptology ePrint Archive*, Tech. Rep. 2017/1060, 2017.
- [13] D. Cash, P. Grubbs, J. Perry, and T. Ristenpart, “Leakage-abuse attacks against searchable encryption,” in *ACM Conference on Communications and Computer Security (CCS ’15)*. ACM, 2015, pp. 668–679.
- [14] D. Cash, S. Jarecki, C. Jutla, H. Krawczyk, M. Rosu, and M. Steiner, “Highly-scalable searchable symmetric encryption with support for boolean queries,” in *Advances in Cryptology - CRYPTO ’13*. Springer, 2013.
- [15] D. Cash and S. Tessaro, “The locality of searchable symmetric encryption,” in *Advances in Cryptology - EUROCRYPT 2014*, 2014.
- [16] D. Cash, P. Grubbs, J. Perry, and T. Ristenpart, “Leakage-abuse attacks against searchable encryption,” *IACR Cryptology ePrint Archive*, vol. 2016, p. 718, 2016. [Online]. Available: <http://eprint.iacr.org/2016/718>
- [17] D. Cash, J. Jaeger, S. Jarecki, C. Jutla, H. Krawczyk, M. Rosu, and M. Steiner, “Dynamic searchable encryption in very-large databases: Data structures and implementation,” in *Network and Distributed System Security Symposium (NDSS ’14)*, 2014.
- [18] M. Chase and S. Kamara, “Structured encryption and controlled disclosure,” in *Advances in Cryptology - ASIACRYPT ’10*, ser. Lecture Notes in Computer Science, vol. 6477. Springer, 2010, pp. 577–594.
- [19] R. Curtmola, J. Garay, S. Kamara, and R. Ostrovsky, “Searchable symmetric encryption: Improved definitions and efficient constructions,” in *ACM Conference on Computer and Communications Security (CCS ’06)*. ACM, 2006, pp. 79–88.
- [20] J. L. Dautrich Jr, E. Stefanov, and E. Shi, “Burst oram: Minimizing oram response times for bursty access patterns,” in *USENIX Security Symposium*, 2014, pp. 749–764.
- [21] I. Demertzis and C. Papamanthou, “Fast searchable encryption with tunable locality,” in *ACM International Conference on Management of Data (SIGMOD ’17)*, ser. SIGMOD ’17. New York, NY, USA: ACM, 2017, pp. 1053–1067. [Online]. Available: <http://doi.acm.org/10.1145/3035918.3064057>
- [22] 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, Vienna, Austria, October 24-28, 2016*, E. R. Weippl, S. Katzenbeisser, C. Kruegel, A. C. Myers, and S. Halevi, Eds. ACM, 2016, pp. 1155–1166. [Online]. Available: <http://doi.acm.org/10.1145/2976749.2978379>
- [23] M. Etemad, A. K p c , C. Papamanthou, and D. Evans, “Efficient dynamic searchable encryption with forward privacy,” *Proceedings on Privacy Enhancing Technologies*, vol. 2018, no. 1, pp. 5–20, 2018.
- [24] B. A. Fisch, B. Vo, F. Krell, A. Kumarasubramanian, V. Kolesnikov, T. Malkin, and S. M. Bellovin, “Malicious-client security in blind seer: A scalable private DBMS,” in *2015 IEEE Symposium on Security and Privacy, SP 2015, San Jose, CA, USA, May 17-21, 2015*. IEEE Computer Society, 2015, pp. 395–410. [Online]. Available: <https://doi.org/10.1109/SP.2015.31>
- [25] C. W. Fletcher, L. Ren, A. Kwon, M. van Dijk, and S. Devadas, “Freecursive oram:[nearly] free recursion and integrity verification for position-based oblivious ram,” in *ACM SIGARCH Computer Architecture News*, vol. 43, no. 1. ACM, 2015, pp. 103–116.
- [26] C. W. Fletcher, L. Ren, A. Kwon, M. Van Dijk, E. Stefanov, D. Serpanos, and S. Devadas, “A low-latency, low-area hardware oblivious ram controller,” in *Field-Programmable Custom Computing Machines (FCCM), 2015 IEEE 23rd Annual International Symposium on*. IEEE, 2015, pp. 215–222.
- [27] S. Garg, P. Mohassel, and C. Papamanthou, “TWRAM: efficient oblivious RAM in two rounds with applications to searchable encryption,” in *Advances in Cryptology - CRYPTO 2016*, 2016, pp. 563–592. [Online]. Available: https://doi.org/10.1007/978-3-662-53015-3_20
- [28] C. Gentry, “Fully homomorphic encryption using ideal lattices,” in *ACM Symposium on Theory of Computing (STOC ’09)*. ACM Press, 2009, pp. 169–178.
- [29] C. Gentry, S. Halevi, C. Jutla, and M. Raykova, “Private database access with he-over-oram architecture,” in *International Conference on Applied Cryptography and Network Security*. Springer, 2015, pp. 172–191.
- [30] E.-J. Goh, “Secure indexes,” *IACR ePrint Cryptography Archive*, Tech. Rep. 2003/216, 2003. See <http://eprint.iacr.org/2003/216>.
- [31] O. Goldreich and R. Ostrovsky, “Software protection and simulation on oblivious RAMs,” *Journal of the ACM*, vol. 43, no. 3, pp. 431–473, 1996.
- [32] P. Grubbs, T. Ristenpart, and V. Shmatikov, “Why your encrypted database is not secure,” in *Workshop on Hot Topics in Operating*

- Systems (HotOS '17)*. New York, NY, USA: ACM, 2017, pp. 162–168. [Online]. Available: <http://doi.acm.org/10.1145/3102980.3103007>
- [33] P. Grubbs, M.-S. Lacharité, B. Minaud, and K. G. Paterson, “Pump up the volume: Practical database reconstruction from volume leakage on range queries,” in *Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security*. ACM, 2018, pp. 315–331.
- [34] 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 2017, San Jose, CA, USA, May 22-26, 2017*. IEEE Computer Society, 2017, pp. 655–672. [Online]. Available: <https://doi.org/10.1109/SP.2017.44>
- [35] A. Hamlin, A. Shelat, M. Weiss, and D. Wichs, “Multi-key searchable encryption, revisited,” in *Public-Key Cryptography - PKC 2018 - 21st IACR International Conference on Practice and Theory of Public-Key Cryptography, Rio de Janeiro, Brazil, March 25-29, 2018, Proceedings, Part I*, 2018, pp. 95–124.
- [36] W. He, D. Akhawe, S. Jain, E. Shi, and D. Song, “Shadowcrypt: Encrypted web applications for everyone,” in *ACM Conference on Computer and Communications Security (CCS '14)*, 2014.
- [37] T. Hoang, A. A. Yavuz, and J. Guajardo, “Practical and secure dynamic searchable encryption via oblivious access on distributed data structure,” in *Proceedings of the 32nd Annual Conference on Computer Security Applications*. ACM, 2016, pp. 302–313.
- [38] M. S. Islam, M. Kuzu, and M. Kantarcioglu, “Access pattern disclosure on searchable encryption: Ramification, attack and mitigation,” in *Network and Distributed System Security Symposium (NDSS '12)*, 2012.
- [39] S. Jarecki, C. Jutla, H. Krawczyk, M. Rosu, and M. Steiner, “Outsourced symmetric private information retrieval,” in *ACM Conference on Computer and Communications Security (CCS '13)*, 2013, pp. 875–888.
- [40] S. Kamara and T. Moataz, “Boolean searchable symmetric encryption with worst-case sub-linear complexity,” in *Advances in Cryptology - EUROCRYPT '17*, 2017.
- [41] S. Kamara and C. Papamanthou, “Parallel and dynamic searchable symmetric encryption,” in *Financial Cryptography and Data Security (FC '13)*, 2013.
- [42] S. Kamara, C. Papamanthou, and T. Roeder, “Dynamic searchable symmetric encryption,” in *ACM Conference on Computer and Communications Security (CCS '12)*. ACM Press, 2012.
- [43] S. Kamara and T. Moataz, “SQL on structurally-encrypted databases,” in *Advances in Cryptology - ASIACRYPT 2018 - 24th International Conference on the Theory and Application of Cryptology and Information Security, Brisbane, QLD, Australia, December 2-6, 2018, Proceedings, Part I*, 2018, pp. 149–180. [Online]. Available: https://doi.org/10.1007/978-3-030-03326-2_6
- [44] —, “Computationally volume-hiding structured encryption,” in *EUROCRYPT '19*, 2019.
- [45] S. Kamara, T. Moataz, and O. Ohrimenko, “Structured encryption and leakage suppression,” in *Advances in Cryptology - CRYPTO '18*, 2018.
- [46] G. Kellaris, G. Kollios, K. Nissim, and A. O. Neill, “Generic attacks on secure outsourced databases,” in *ACM Conference on Computer and Communications Security (CCS '16)*, 2016.
- [47] E. Kushilevitz, S. Lu, and R. Ostrovsky, “On the (in) security of hash-based oblivious ram and a new balancing scheme,” in *ACM-SIAM Symposium on Discrete Algorithms (SODA '12)*, 2012, pp. 143–156.
- [48] M.-S. Lacharité, 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)*. IEEE, 2018, pp. 297–314.
- [49] B. Lau, S. Chung, C. Song, Y. Jang, W. Lee, and A. Boldyreva, “Mimesis aegis: A mimicry privacy shield—a system’s approach to data privacy on public cloud,” in *USENIX Security Symposium (USENIX Security 14)*, 2014, pp. 33–48.
- [50] K. Lewi and D. Wu, “Order-revealing encryption: New constructions, applications, and lower bounds,” in *ACM Conference on Computer and Communications Security (CCS '16)*, 2016.
- [51] J. R. Lorch, J. W. Mickens, B. Parno, M. Raykova, and J. Schiffman, “Toward practical private access to data centers via parallel oram,” *IACR Cryptology ePrint Archive*, vol. 2012, p. 133, 2012.
- [52] J. R. Lorch, B. Parno, J. W. Mickens, M. Raykova, and J. Schiffman, “Shroud: ensuring private access to large-scale data in the data center,” in *FAST*, vol. 2013, 2013, pp. 199–213.
- [53] M. Maffei, G. Malavolta, M. Reinert, and D. Schröder, “Privacy and access control for outsourced personal records,” in *Security and Privacy (SP), 2015 IEEE Symposium on*. IEEE, 2015, pp. 341–358.
- [54] X. Meng, S. Kamara, K. Nissim, and G. Kollios, “GreCs: Graph encryption for approximate shortest distance queries,” in *ACM Conference on Computer and Communications Security (CCS 15)*, 2015.
- [55] I. Miers and P. Mohassel, “IO-DSSE: scaling dynamic searchable encryption to millions of indexes by improving locality,” in *24th Annual Network and Distributed System Security Symposium, NDSS 2017, San Diego, California, USA, February 26 - March 1, 2017*. The Internet Society, 2017. [Online]. Available: <https://www.ndss-symposium.org/ndss2017/>
- [56] M. Naveed, S. Kamara, and C. V. Wright, “Inference attacks on property-preserving encrypted databases,” in *ACM Conference on Computer and Communications Security (CCS)*, ser. CCS '15. ACM, 2015, pp. 644–655. [Online]. Available: <http://doi.acm.org/10.1145/2810103.2813651>
- [57] NIST, “Trec 2007 public corpus,” <https://plg.uwaterloo.ca/~gvcormac/trec corpus07/>, 2018.
- [58] O. Ohrimenko, M. T. Goodrich, R. Tamassia, and E. Upfal, “The melbourne shuffle: Improving oblivious storage in the cloud,” in *International Colloquium on Automata, Languages, and Programming*. Springer, 2014, pp. 556–567.
- [59] V. Pappas, F. Krell, B. Vo, V. Kolesnikov, T. Malkin, S.-G. Choi, W. George, A. Keromytis, and S. Bellovin, “Blind seer: A scalable private dbms,” in *Security and Privacy (SP), 2014 IEEE Symposium on*. IEEE, 2014, pp. 359–374.
- [60] S. Patel, G. Persiano, and K. Yeo, “Symmetric searchable encryption with sharing and unsharing,” *IACR Cryptology ePrint Archive*, vol. 2017, p. 973, 2017.
- [61] S. Patel, G. Persiano, K. Yeo, and M. Yung, “Mitigating leakage in secure cloud-hosted data structures: Volume-hiding for multi-maps via hashing,” in *Proceedings of the 2019 ACM SIGSAC Conference on Computer and Communications Security*. ACM, 2019, pp. 79–93.
- [62] M. Porter, “Stop word - snowball,” <http://snowball.tartarus.org/>, 2018.
- [63] C. Project, “Enron email dataset,” <https://www.cs.cmu.edu/~enron>, 2018.
- [64] L. Ren, C. W. Fletcher, A. Kwon, E. Stefanov, E. Shi, M. Van Dijk, and S. Devadas, “Constants count: Practical improvements to oblivious ram,” in *USENIX Security Symposium*, 2015, pp. 415–430.
- [65] L. Ren, C. W. Fletcher, X. Yu, M. Van Dijk, and S. Devadas, “Integrity verification for path oblivious-ram,” in *High Performance Extreme Computing Conference (HPEC), 2013 IEEE*. IEEE, 2013, pp. 1–6.
- [66] D. S. Roche, A. Aviv, and S. G. Choi, “A practical oblivious map data structure with secure deletion and history independence,” in *Security and Privacy (SP), 2016 IEEE Symposium on*. IEEE, 2016, pp. 178–197.
- [67] C. Sahin, V. Zakhary, A. El Abbadi, H. Lin, and S. Tessaro, “Taostore: Overcoming asynchronicity in oblivious data storage,” in *Security and Privacy (SP), 2016 IEEE Symposium on*. IEEE, 2016, pp. 198–217.
- [68] E. Shi, T.-H. Chan, E. Stefanov, and M. Li, “Oblivious RAM with $O((\log N)^3)$ Worst-Case Cost,” in *Advances in Cryptology - ASIACRYPT '11*. Springer-Verlag, 2011, pp. 197–214. [Online]. Available: http://dx.doi.org/10.1007/978-3-642-25385-0_11
- [69] D. Song, D. Wagner, and A. Perrig, “Practical techniques for searching on encrypted data,” in *IEEE Symposium on Research in Security and Privacy*. IEEE Computer Society, 2000, pp. 44–55.
- [70] E. Stefanov, C. Papamanthou, and E. Shi, “Practical dynamic searchable encryption with small leakage,” in *Network and Distributed System Security Symposium (NDSS '14)*, 2014.
- [71] 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 *ACM Conference on Computer and Communications Security (CCS '13)*, 2013.
- [72] E. Stefanov and E. Shi, “Multi-cloud oblivious storage,” in *2013 ACM SIGSAC Conference on Computer and Communications Security, CCS'13, Berlin, Germany, November 4-8, 2013*, A. Sadeghi, V. D. Gligor, and M. Yung, Eds. ACM, 2013, pp. 247–258. [Online]. Available: <http://doi.acm.org/10.1145/2508859.2516673>

- [73] —, “Oblivstore: High performance oblivious distributed cloud data store,” in *NDSS*, 2013.
- [74] X. S. Wang, K. Nayak, C. Liu, T. Chan, E. Shi, E. Stefanov, and Y. Huang, “Oblivious data structures,” in *Proceedings of the 2014 ACM SIGSAC Conference on Computer and Communications Security*. ACM, 2014, pp. 215–226.
- [75] Y. Zhang, J. Katz, and C. Papamanthou, “All your queries are belong to us: The power of file-injection attacks on searchable encryption,” in *USENIX Security Symposium*, 2016.

APPENDIX A OVERVIEW OF ESA CONSTRUCTIONS

We recall some common ESA constructions based on both STE and ORAM.

Baseline (BSL). Let $\Sigma_{\text{mm}} = (\text{Setup}, \text{Get})$ be a response-revealing multi-map encryption scheme and $\text{SKE} = (\text{Gen}, \text{Enc}, \text{Dec})$ be a symmetric-key encryption scheme. Consider the ESA scheme $\text{BSL} = (\text{Setup}, \text{Search})$ where each algorithm works as follows:

- $\text{BSL.Setup}(1^k, \mathbf{D})$: it builds a multi-map MM that maps each keyword $w \in \mathbb{W}$ to the tuple $\mathbf{D}(w)$. It then computes $(K_1, \text{EMM}) \leftarrow \Sigma_{\text{mm}}.\text{Setup}(1^k, \text{MM})$ and, for all $i \in [n]$, $\text{ct}_i \leftarrow \text{Enc}(K_2, \text{D}_i)$, where $K_2 \leftarrow \text{SKE.Gen}(1^k)$. It outputs (K, \mathbf{ED}) , where $K = (K_1, K_2)$ and $\mathbf{ED} = (\text{EMM}, \text{ct}_1, \dots, \text{ct}_n)$.
- $\text{BSL.Search}(K, w; \mathbf{ED})$: the parties execute $(\perp; \mathbf{I}) \leftarrow \Sigma_{\text{mm}}.\text{Get}(K_1, w; \text{EMM})$ after which the server returns $(\text{ct}_i)_{i \in \mathbf{I}}$ to the client. For all $i \in \mathbf{I}$, the client computes $\text{D}_i := \text{SKE.Dec}(K_2, \text{ct}_i)$.

If the multi-map encryption scheme Σ_{mm} is instantiated with any of the standard constructions [19], [18], [17], [10], [11], [3], the SSE scheme will have leakage profile

$$\Lambda_{\text{BSL}} = (\star, (\text{req}, \text{rid}, \text{vol})).$$

Its storage complexity will be

$$O\left(\sum_{w \in \mathbb{W}} \#\mathbf{D}(w) + \sum_{i=1}^n |\text{D}_i|_w\right),$$

and its search and communication complexity will be

$$O\left(\#\mathbf{D}(w) + \sum_{\mathbf{D} \in \mathbf{D}(w)} |\mathbf{D}|_w\right),$$

which is optimal.

Opaque (OPQ). As far as we know, the construction we now describe has not appeared in prior work. It has a relatively low leakage profile and optimal search and communication complexity at the cost of additional storage. Let $\Sigma_{\text{mm}} = (\text{Setup}, \text{Get})$ be a response-hiding multi-map encryption scheme and let $\text{SKE} = (\text{Gen}, \text{Enc}, \text{Dec})$ be a symmetric-key encryption scheme. Consider the structured ESA scheme $\text{OPQ} = (\text{Setup}, \text{Search})$ where each algorithm works as follows:

- $\text{OPQ.Setup}(1^k, \mathbf{D})$: it builds a multi-map MM that maps each keyword $w \in \mathbb{W}$ to a tuple $\mathbf{t} = (t_1, \dots, t_{a/B})$ composed of B -sized blocks, where t_i is the i th block of

the concatenation of the documents (not identifiers) that contain w , and where

$$a = \sum_{\mathbf{D} \in \mathbf{D}(w)} |\mathbf{D}|_w.$$

The algorithm then computes $(K, \text{EMM}) \leftarrow \Sigma_{\text{mm}}.\text{Setup}(1^k, \text{MM})$ and outputs (K, \mathbf{ED}) where $\mathbf{ED} = \text{EMM}$.

- $\text{OPQ.Search}(K, w; \mathbf{ED})$: the parties execute

$$(\mathbf{t}; \perp) \leftarrow \Sigma_{\text{mm}}.\text{Get}(K, w; \text{EMM})$$

and the client parses \mathbf{t} as $(\mathbf{D})_{\mathbf{D} \in \mathbf{D}(w)}$.

If the multi-map encryption scheme Σ_{mm} is instantiated with a standard response-hiding encrypted multi-map [19], [18], [17], [10], [11]¹² the ESA will have leakage profile

$$\Lambda_{\text{OPQ}} = (\mathcal{L}_S, \mathcal{L}_Q) = (\star, (\text{req}, \text{tvol})).$$

Its storage complexity will be

$$O\left(\sum_{w \in \mathbb{W}} \sum_{\mathbf{D} \in \mathbf{D}(w)} |\mathbf{D}|_w\right),$$

and the search and communication complexity will be

$$O\left(\#\mathbf{D}(w) + \sum_{\mathbf{D} \in \mathbf{D}(w)} |\mathbf{D}|_w\right),$$

which is optimal.

Semi-ORAM (SMI). Let $\text{ORAM} = (\text{Setup}, \text{Read})$ be an ORAM scheme and $\text{SKE} = (\text{Gen}, \text{Enc}, \text{Dec})$ be a symmetric-key encryption scheme. Consider the scheme $\text{SMI} = (\text{Setup}, \text{Search})$ where each algorithm works as follows:

- $\text{SMI.Setup}(1^k, \mathbf{D})$: it builds a multi-map MM that maps each keyword $w \in \mathbb{W}$ to the tuple $\mathbf{D}(w)$. It then computes $(K_1, \text{OMM}) \leftarrow \text{ORAM.Setup}(1^k, \text{MM})$ and, for all $i \in [n]$, $\text{ct}_i \leftarrow \text{Enc}(K_2, \text{D}_i)$, where $K_2 \leftarrow \text{SKE.Gen}(1^k)$. It outputs (K, \mathbf{ED}) , where $K = (K_1, K_2)$ and $\mathbf{OD} = (\text{OMM}, \text{ct}_1, \dots, \text{ct}_n)$.
- $\text{SMI.Search}(K, w; \mathbf{OD})$: the client uses ORAM to simulate an execution of $\text{Get}(\text{MM}, w)$; that is, it runs $\text{Get}(\text{MM}, w)$ locally but replaces every read operation to location i with an execution of $\text{ORAM.Read}(K_1, i; \text{OMM})$. At the end of this simulation, the client holds a set of indices \mathbf{I} which it sends to the server. The server returns $(\text{ct}_i)_{i \in \mathbf{I}}$ which the client decrypts.

If the ORAM scheme is instantiated with any standard construction [31], [47], [71], the search scheme will have leakage profile

$$\Lambda_{\text{SMI}} = (\star, (\text{rlen}, \text{rid}, \text{vol})).$$

Using Path ORAM [71], the storage complexity is

$$O\left(\sum_{w \in \mathbb{W}} \#\mathbf{D}(w) + \sum_{i=1}^n |\text{D}_i|_w\right),$$

¹²We note that while most of these constructions are described as response-revealing constructions it is trivial to convert them to response-hiding schemes.

and the search and communication complexity are

$$O\left(\frac{\#\mathbf{D}(w)}{B} \cdot \log^2\left(\sum_{w \in \mathbb{W}} \frac{\#\mathbf{D}(w)}{B}\right) + \sum_{\mathbf{D} \in \mathbf{D}(w)} |\mathbf{D}|_w\right),$$

where B is the block size in bits.

Full ORAM (FLL). Let $\text{ORAM} = (\text{Setup}, \text{Read})$ be an ORAM scheme and consider the scheme $\text{FLL} = (\text{Setup}, \text{Search})$ where each algorithm works as follows:

- $\text{FLL.Setup}(1^k, \mathbf{D})$: it builds an array RAM that stores all the documents in \mathbf{D} . It then builds a multi-map MM that maps each keyword $w \in \mathbb{W}$ to the locations of the blocks in RAM that store the documents in $\mathbf{D}(w)$. It then computes $(K_1, \text{OMM}) \leftarrow \text{ORAM.Setup}(1^k, \text{MM})$ and $(K_2, \text{ORAM}) \leftarrow \text{ORAM.Setup}(1^k, \text{RAM})$. It outputs (K, \mathbf{OD}) , where $K = (K_1, K_2)$ and $\mathbf{OD} = (\text{OMM}, \text{ORAM})$.
- $\text{FLL.Search}(K, w; \mathbf{OD})$: the client uses ORAM to simulate an execution of $\text{Get}(\text{MM}, w)$. At the end of this simulation, the client holds a set of indices \mathbf{I} . It then uses ORAM again to simulate, for all $i \in \mathbf{I}$, an execution of $\text{Read}(\text{RAM}, i)$ to recover the documents.

If the ORAM scheme is instantiated with any standard construction [31], [47], [71], the search scheme will have leakage profile

$$\Lambda_{\text{FLL}} = (\star, (\text{rlen}, \text{tvol})).$$

The storage complexity is

$$O\left(\sum_{w \in \mathbb{W}} \#\mathbf{D}(w) + \sum_{\mathbf{D} \in \mathbf{D}(w)} |\mathbf{D}|_w\right),$$

and the search and communication complexity are

$$O\left(\frac{\#\mathbf{D}(w)}{B_1} \cdot \log^2\left(\sum_{w \in \mathbb{W}} \frac{\#\mathbf{D}(w)}{B_1}\right) + \sum_{i \in \mathbf{D}(w)} \frac{|\mathbf{D}_i|_w}{B_2} \cdot \log^2\left(\sum_{i=1}^n \frac{|\mathbf{D}_i|_w}{B_2}\right)\right)$$

where B_1 and B_2 are the block sizes in bits of the first and second ORAM, respectively. Note that this construction has leakage profile Λ_{FLL} only if the client retrieves all of the matching documents from the second ORAM at once. If, on the other hand, the client retrieves them one by one then it will have leakage profile $(\star, (\text{rlen}, \text{vol}))$.

Additional ORAM-based constructions. We note that there are alternative ORAM-based designs in addition to the ones we described above. One could, for example, merge the two ORAMs used in the full ORAM simulation into a single ORAM with the same block size. This would have leakage pattern (\star, tvol) .

The Piggyback scheme (PBS). PBS is an STE scheme recently introduced in [45] that partially hides the volume pattern. It comes in two variants. The first reveals only the sequence volume pattern (i.e., the sum of the volume associated to a query sequence) on non-repeating query sequences. The second variant reveals nothing (beyond a public parameter independent of the volume) on non-repeating query sequences.

At a high level, the scheme leverages a new trade-off in STE design; specifically, it trades latency for an improved leakage profile. At a high level, the scheme processes the input data structure such that the query responses are divided into smaller chunks of equal size. These chunks are then stored and encrypted so that, on each query, the client only retrieves a fixed number of chunks. If the whole response is not retrieved at that moment, then the query is added to a queue and the remaining chunks are retrieved on the next query. The responses can therefore be delayed but the authors show that the delay can be minimal for standard query distributions.

Volume-hiding constructions. VLH and AVLH are volume-hiding encrypted multi-map constructions recently introduced by Kamara and Moataz [44]. These schemes are the first volume-hiding STE constructions that do not rely on naïve padding. VLH makes use of a pseudo-random function F and an optimal multi-map encryption scheme. It is parameterized with a public parameter $\lambda \geq 1$ that affects correctness. Given a multi-map MM, the scheme determines a new response length for each label ℓ in MM which is computed by evaluating F on ℓ 's original response length and adding λ . If the new response length is larger than the original, then ℓ 's tuple is padded. If the new response length is smaller than the original, then ℓ 's tuple is truncated. AVLH is a more advanced construction based on a new design paradigm based on bi-partite graphs. More precisely, AVLH transforms its input multi-map as a bi-partite graph where top vertices correspond to the multi-map's labels and the bottom vertices correspond to bins. Each label's tuple values are then stored in its associated bin in a specific way. The bins are then padded to have the same size. At query time, the user always retrieves the same number of bins. AVLH does not improve on the query complexity of encrypted multi-map schemes but does improve on the storage efficiency of naïve padding. In [44] it is then shown that the storage can be further reduced by relying on the conjectured hardness of the planted densest subgraph problem.

APPENDIX B

COUNT v.1 WITH $\delta < 1$

The Count v.1 attack was shown in [13] to have high recovery rate when $\delta = 1$; that is, when the adversary has full knowledge of the data. For $\delta < 1$, however, the attack seems to only work if $\delta \leq .8$. We found that the experimental results for $\delta < 1$ that are reported in [13], however, are for an unpublished variant of the count attack that relies on knowledge of client queries. To better understand how known queries impact the recovery rate of Count v.1, we evaluated the attack with a varying fraction of known queries. The results are shown in Figure 5.¹³ When the adversary knows 5% of the queries, recovery rates are similar to the ones reported in [13]. When the adversary knows 2% known queries, however, the attack ceases to work even with $\delta = .9$.

APPENDIX C

KEYWORD SELECTIVITY

Our empirical evaluation (see Section V) clearly shows that the selectivity of the queries is by far the most important

¹³This experiment was performed using the implementation and dataset of [13]. We thank the authors for promptly sharing their implementation and data with us.

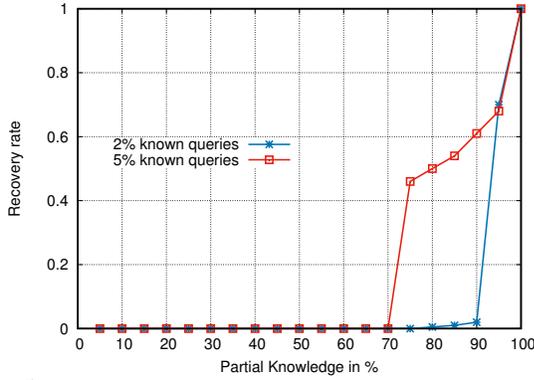
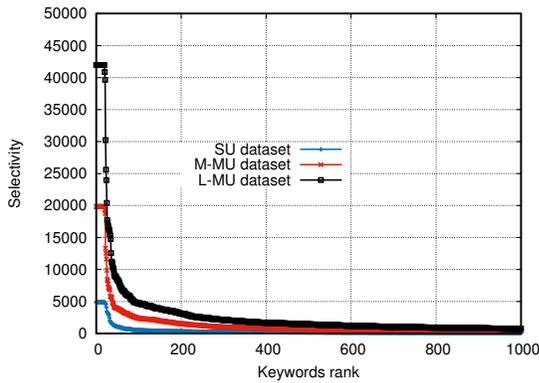
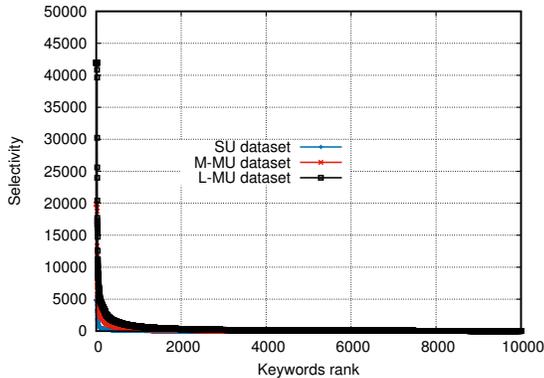


Figure 5: Count v.1 with varying fractions of known queries (on 150 queries out of a keywords space of size 500).

factor on the recovery rate of all the attacks. Understanding the selectivity of keywords in our dataset is therefore important. In Figures 6 (a) and (b) we plot the selectivity of 1000 and 10,000 most selective keywords, respectively, in our datasets after stemming and removing stop words. We can see in these Figures that keyword selectivity in Enron is power law distributed. In other words, only a few keywords have high and unique selectivities whereas the overwhelming majority of keywords have low and common selectivities (at most 3).



(a) Distribution of the most selective 1000 keywords.

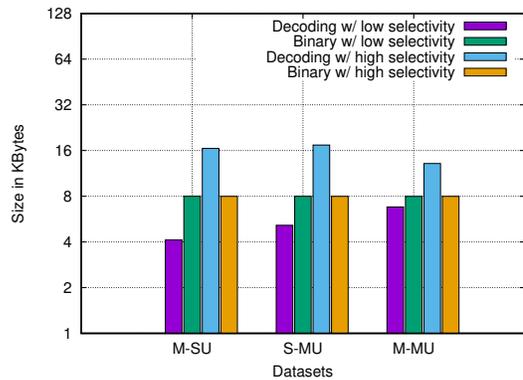


(b) Distribution of the most selective 10,000 keywords.

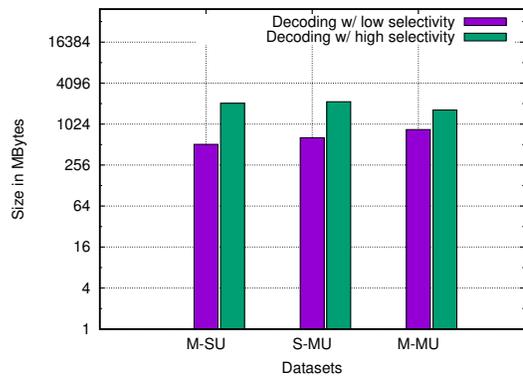
Figure 6: Keyword selectivity.

APPENDIX D QUANTIFYING THE OFFSET FOR INJECTION

The total number of files injected by both the Decoding and Binary Search attacks depend on an offset γ which is determined by characteristics of the data collection. Here, we study the values of these offsets on three different collections: SU, S-MU and M-MU as defined in Section V. Our results are described in Figure 7. We found that querying on high- or low-selectivity keywords did not have any impact on the Binary Search attack. However, as can be seen from its description, the size of the keyword space did have an impact. For the Decoding attack, the amount of injected data did depend on the selectivity of the queries: the amount for high-selectivity queries was about twice as much as for low-selectivity queries. This held for both the SU and S-MU datasets. We believe that this is inherent to the way the offset is computed. In fact, on high-selectivity queries, we noticed that the total volumes tend to have a higher gap between them. This is not the case for low-selectivity queries.



(a) Amount of injected data to recover one keyword.



(b) Amount of injected data to recover 500 keywords.

Figure 7: Amount of injected data for both the Decoding and Binary Search attacks (with $\#\mathbb{W} = 500$).