

Leakage-Abuse Attacks against Order-Revealing Encryption

Paul Grubbs^{1,4}, Kevin Sekniqi¹, Vincent Bindschaedler², Muhammad Naveed³, Thomas Ristenpart^{1,4}

¹ Cornell University

² UIUC

³ USC

⁴ Cornell Tech

Abstract—Order-preserving encryption and its generalization order-revealing encryption (OPE/ORE) are used in a variety of settings in practice in order to allow sorting, performing range queries, and filtering data — all while only having access to ciphertexts. But OPE and ORE ciphertexts necessarily leak information about plaintexts, and what level of security they provide has been unclear.

In this work, we introduce new leakage-abuse attacks that show how to recover plaintexts from OPE/ORE-encrypted databases. Underlying our new attacks against practically-used schemes is a framework in which we cast the adversary’s challenge as a non-crossing bipartite matching problem. This allows easy tailoring of attacks to a specific scheme’s leakage profile. In a case study of customer records, we show attacks that recover 99% of first names, 97% of last names, and 90% of birthdates held in a database, despite all values being encrypted with the OPE scheme most widely used in practice.

We also show the first attack against the recent frequency-hiding Kerschbaum scheme, to which no prior attacks have been demonstrated. Our attack recovers frequently occurring plaintexts most of the time.

I. INTRODUCTION

Due to frequent data breaches and broad interest in outsourcing data to the cloud, companies increasingly want to encrypt sensitive information before storing it in local databases or uploading to services operated by third parties. Standard encryption mechanisms would, however, reduce the value of these databases and services by preventing them from doing useful operations on the data. A seeming solution is to use so-called property-revealing encryption (PRE) schemes that allow limited operations over ciphertexts by making public specific properties of plaintexts. Systems based on PREs are already used in industry [1, 3, 5, 9] and increasingly studied in the academic literature [9, 23, 26, 29, 37, 38].

A widely desired property to preserve is order. This allows sorting, as well as range and prefix search queries, which are needed to support other server-side operations. Order-preserving encryption (OPE) [8, 12] ensures that $E_k(m_1) < E_k(m_2)$ for $m_1 < m_2$. By allowing standard comparison operators directly on ciphertexts one can easily integrate OPE into systems, and many systems have done so [1, 3, 5, 9, 37]. A generalization of OPE is order-revealing encryption (ORE) [14] which reveals ordering relations by way of a public comparison function that operates on pairs of plaintexts.

What level of security OPE and ORE schemes achieve is poorly understood. Boldyreva, et al. [12, 13] explicitly warn about the uncertain security of OPE, and subsequently a recent

study by Naveed, Kamara, and Wright (NKW) [35] empirically explore the security of OPE and ORE. They exploit inference attacks: using some auxiliary information about the distribution of plaintexts, the adversary abuses the order and frequency of plaintexts preserved within a sequence of ciphertexts to recover as many plaintexts as possible. Their best performing attack, called the cumulative attack, worked well to recover data that lies within small domains, such as days of the year (365 possible values).

Intuitively, attacks on such small plaintext domains are easier because most of the plaintexts in the domain are likely to appear in any given target dataset, and (deterministic) schemes reveal all plaintexts trivially in this case. The NKW cumulative attack also scales poorly, requiring $\mathcal{O}(n^3)$, where n is the max of the number of unique ciphertexts and the number of unique items in the auxiliary data. All this leaves open the possibility that OPE and ORE schemes remain secure in practice for plaintext data drawn from larger domains, and practitioners could simply avoid using OPE for small-domain data. Indeed researchers have suggested exactly that approach as a hedge against known attacks [21].

Our contributions. We develop new attacks that target, in aggregate, all suggested OPE/ORE schemes, and show their efficacy against plaintexts drawn from large domains. Our attacks come in two flavors: generic, like the NKW attack, and scheme-specific. Generic attacks work against any scheme that leaks frequency and order information. We also introduce the first scheme-specific attacks that exploit additional leakage about plaintexts. These latter attacks cover the OPE scheme most widely used in practice, due to Boldyreva, Chenette, Lee, and O’Neill (BCLO) [12]. Finally, we provide an attack against the recent frequency-hiding scheme of Kerschbaum [27].

We use as a running case study an encrypted customer record database, a frequent scenario for these encryption technologies in industry. We utilize a variety of public datasets as stand-ins for the kind of data targeted in an attack, including such customer attributes as first names, last names, ZIP codes, and birthdates. These attribute types come from domains much larger than those considered in the prior work; for example there are hundreds of thousands of unique first names, tens of thousands of birthdates (across a given century span), and tens of thousands of unique ZIP codes in the United States. We identify distinct public data sets that can be used as auxiliary information in inference attacks.

Generic attacks and non-crossing matching. We first revisit the NKW attack, which is generic in that it works against

Scheme(s)	First names	Last names
Kerschbaum [27]	26%	6%
Popa et al. [36], Kerschbaum [28]	84%	38%
BCLO [12, 13]	99%	97%
CLWW [18]	98%	75%
BCLO + CLWW [18]	85%	44%
Baseline Guessing	4%	1%

Fig. 1: Average recovery rate of plaintexts (in percent of all target ciphertexts) across all first and last name datasets for our best attacks against the indicated OPE/ORE scheme(s). The last row indicates the recovery rate given by just guessing the most likely message for each ciphertext. We have bolded those attacks that fully recover more than half the plaintexts, on average.

any scheme that reveals plaintext frequency and order. Its accuracy in our customer record case study is relatively poor. For example, it recovers on average 44% of encrypted first name records but this accounts for only 5% of the unique first names in the data set, meaning the attack only recovers a small fraction of the high-frequency elements of the distribution. The attack is also relatively slow, requiring six hours for larger target datasets on a well-provisioned cluster.

We explore their attack in more detail, and point out some limitations of their approach. Their most accurate attack, called the cumulative attack, casts the adversarial task as solving a linear sum assignment problem. This approach, however, typically yields solutions that violate adversarially known ordering constraints, outputting a guess that ciphertext c_1 encrypts plaintext p_2 and ciphertext c_2 encrypts plaintext p_1 , yet $c_1 < c_2$ and $p_2 > p_1$.

We observe that the adversary’s goal can be cast as a classic graph problem: a min-weight non-crossing bipartite matching. One set of nodes consists of ciphertexts and the other possible plaintexts, and the edges are weighted using frequency information. Non-crossing refers to the fact that we want the solution to abide by the adversarially known ordering constraints. Unlike NKW attack, our non-crossing attack takes full advantage of frequency and order information. We use a classic algorithm [32] to solve the min-weight non-crossing matching problem, and its runtime is both asymptotically and concretely better than the NKW approach. We call this the non-crossing attack, and show that it achieves significantly better accuracy, for example it recovers on average 84% of first names in our target encrypted datasets, about a 2x improvement over NKW. High-frequency plaintexts are particularly at risk: our attack recovers on average 95% of the 20 most frequent first name plaintexts, and 88% of the 20 most frequent last names. The non-crossing attack runs in only a few hours, even for the largest target dataset, where the induced matching problem has over 17 billion edges.

Exploiting further leakage. The attacks so far are generic, and do not take advantage of the additional leakage exhibited by the OPE schemes used in practice. The most widely used scheme is due to Boldyreva, Chenette, Lee, and O’Neill

(BCLO) [12], because it is fast, easy to deploy, and has ciphertexts that are only a few bits longer than plaintexts. It’s been proved secure in the sense of indistinguishability from a random order-preserving function, but what level of security this offers is not clear: prior work has shown that for uniformly chosen messages, ciphertexts leak almost the entire first half of the plaintext [13]. Despite widespread deployment, the implications of this leakage for real datasets has not been studied.

We rectify this, showing that in practical scenarios using the BCLO scheme to encrypt a set of first names, for example, allows an attacker to recover almost half the data set. The leakage is even worse for last names, with almost 97% of last names trivially recoverable. A key issue making this leakage so damaging is that varying-length plaintexts must be padded to the maximum length plaintext, ensuring that shorter messages completely reside in the leaked first half of the padded plaintext. We emphasize that exploiting this leakage does not require mounting an inference attack, rather an adversary simply can inspect ciphertexts, perform a few elementary calculations, and produce (most) plaintexts.

That said, our graph-based viewpoint on inference attacks allows us to easily combine this leakage with inference to improve recovery rates further. We adapt our non-crossing attack to first compute the leakage implied by [13], use it to exclude matchings that cannot occur, and then solve the resulting, narrower non-crossing matching problem. The resulting attack allows us to recover on average 99% of first names, more than doubling the recovery rate over using the BCLO leakage alone without inference.

We also apply our approach to exploit the leakage of the more recent Chenette, Lewi, Weis, and Wu (CLWW) [18] ORE scheme. The leakage of CLWW is different than that of BCLO, and they argue that their scheme may be more secure: CLWW prove that for uniform randomly distributed messages, their scheme’s leakage is asymptotically less than the BCLO leakage. They also propose that security may be further improved by composing an OPE scheme with their ORE scheme. We apply our framework to both the CLWW scheme and the composition of BCLO with CLWW, and show that, unfortunately, CLWW provides an equivalently poor level of security as compared to BCLO in the settings we consider. Interestingly, some plaintext distributions are particularly bad for CLWW leakage: for example, our attack against BCLO does not perform well on ZIP codes (achieving 12% recovery rate), but on CLWW our attack recovers 97% of encrypted ZIP codes. Composition of the two schemes does decrease attack accuracy, but is still far from providing an acceptable level of security.

Exploiting known plaintexts. Our attacks, as well as the NKW attack, assumes a relatively weak adversary that only obtains ciphertexts, but has no information about any specific plaintext values. As we discuss in detail in the body, in real-world deployments many systems relying on OPE/ORE offer attackers the ability to mount known- or chosen-plaintext attacks. We show how to extend our attack framework to take advantage of known plaintexts by simply partitioning the

attack problem based on the known values and running an attack against each resulting sub-problem independently. As the BCLO and CLWW attacks already recover most plaintexts, we see how well this partitioning approach would help our generic non-crossing attack. Knowing some small percentage of plaintexts provides a modest improvement for first names and last names, but a huge boost for birthdates and ZIP codes, attesting to the value to attackers of known plaintexts for some settings.

Attacking Kerschbaum’s scheme. The attacks mentioned above work against schemes that at least leak both frequency and order, but Kerschbaum [27] recently introduced a scheme that hides frequency information. Here there is no prior work, as the NKW attacks do not apply to frequency-hiding schemes, and our non-crossing attack framework also does not apply. We propose a new “bucketing” attack that performs reasonably well, recovering on average 26% of first names. Notably, it recovers the majority of high-frequency plaintexts (despite not having frequency information leaked), suggesting these plaintexts are particularly poorly protected by order-revealing schemes.

Summary. We are the first to explore the security of OPE and ORE schemes when used with plaintext data types that were, before our work, not known to admit attacks. What’s more, our case study of customer data is representative of common industry practice. Underlying our new results is a framework for constructing attacks based on min-weight non-crossing bipartite matching, which allow for easy extensibility in the face of leakage beyond frequency and order. Most importantly for current industry practice, our results show, for the first time, how the leakage of the BCLO scheme would enable recovery of essentially all plaintexts encrypted in typical customer record databases. Suggested practical alternatives such as the CLWW scheme, or the composition of it with BCLO, do not fair much better. See Figure 1 for a high level summary of our quantitative results for first and last names.

Our results offer guidance to practitioners about the security level offered by OPE and ORE schemes. While obviously using property-revealing encryption is better than leaving data in the clear (in some settings the only viable alternative currently), our work indicates that the security benefits of deployed schemes is quite marginal.

In terms of countermeasures, an obvious suggestion is to move towards less leaky schemes, such as those that only reveal order, including Kerschbaum’s scheme and the more recent Lewi et al. scheme based on multilinear maps [14, 30]. Unfortunately in most settings there exists inherent challenges to deployment of these schemes. Kerschbaum’s scheme is relatively efficient, but requires client-side state which impedes scaling. The Lewi et al. scheme has ciphertexts larger by 10 orders of magnitude than BCLO ciphertexts and requires tens of minutes to compute encryptions. Even so, our attack against such frequency hiding schemes shows that for common use cases the high frequency plaintexts may nevertheless be revealed to attackers.

II. PRELIMINARIES

Basic notation. Let \mathcal{D} be a set and let its size be denoted by $s = |\mathcal{D}|$. We assume some total ordering of elements in \mathcal{D} . Let D be a sequence of elements (d_1, \dots, d_n) , where each $d_i \in \mathcal{D}$. For $1 \leq i \leq s$, we define the histogram function for D to be the function $H_D(i)$ that outputs the number of times the i^{th} element of \mathcal{D} appears in D , divided by s . For $1 \leq j \leq s$, we define the cumulative distribution function (CDF) for D to be the function $F_D(j) = (\sum_{i=1}^j h_i)/s$. Observe that one can represent histograms and CDFs linearly in the number of unique elements in D .

Order-preserving and revealing encryption. An order-preserving encryption (OPE) [8, 12] E allows encrypting data under a secret key k such that $E_k(m_1) < E_k(m_2)$ for any $m_1 < m_2$. The benefit of OPE is that existing comparison operations will work transparently on ciphertexts, thus providing a drop-in way to replace plaintexts with ciphertexts while preserving the ability to perform order comparisons, range queries, etc. Order-revealing encryption (ORE) [14] does not allow an existing comparison operator to work on ciphertexts, but instead comes with a public procedure that can determine the order of two plaintexts given only their two ciphertexts.

PRE-based systems. OPE and ORE are examples of what we refer to as property-revealing encryption (PRE) schemes. PRE schemes encrypt data while allowing some limited amount of computation over ciphertexts. Schemes facilitate such computation efficiently by revealing specific properties of plaintexts, in our cases the relative order of two plaintexts. (Looking ahead, in-use PRE schemes for the order property leak a lot more than just order.)

A motivating scenario is in client-server systems where one wants to perform encryption on the client side but take advantage of the server performing some operations on the client’s behalf. In a PRE-based encrypted database, for example, OPE is often used to enable certain kinds of expressive SQL-like queries on encrypted data. The standard database functionality in such a system is extended and modified by a database proxy, which rewrites queries and performs encryption/decryption on behalf of database clients. The clients are often application servers (acting on behalf of human users) which address the database proxy as though it were the database. Example PRE-based databases include IQrypt [2], CryptDB [37, 39], and Cipherbase [9].

Another concrete application arises with security-conscious businesses that use network middleware to encrypt and decrypt data as users interact with cloud software-as-a-service (SaaS) applications. Such “encryption proxies” are similar to the database proxy described above, except in most cases it cannot change the way the SaaS application works. PRE schemes including OPE and ORE become useful to ensure stored data does not break (some) useful cloud functionality. Commercial products from Skyhigh Networks [5], CipherCloud [1], and Perspecsys [3] are examples.

In all these deployment settings, a minimal security requirement is the confidentiality of plaintext data in the face of

attackers that obtain access to a PRE-encrypted database, e.g., by compromising a server or obtaining insider access to it.

Types of ORE/OPE schemes. We note OPE and ORE schemes come in several flavors that affect their deployability in the above application scenarios. Some schemes [27, 28, 36] require state to be stored with the client beyond the secret key. This state can potentially be held by proxies on behalf of clients, but regardless complicates scaling to large numbers of clients. A OPE/ORE scheme is mutable if the ciphertexts might change as more values are added to a database. Examples of such schemes include [27, 28, 36]. Finally some schemes are interactive [27, 28, 36], requiring multiple rounds of communication between the client and server to store or retrieve ciphertexts.

Statefulness, mutability, and interactivity all hinder deployability. The reason to consider such schemes is that the more deployable schemes [12], which are the only ones currently used in practice, leak more information about plaintexts than just ordering information.

III. OVERVIEW AND METHODOLOGY

We later introduce new attacks against OPE and ORE schemes and explore them experimentally. The prior work by NKW focused on medical settings, where a database of patient data was outsourced in encrypted form. The attributes (columns) of the databases they considered had plaintext values falling within small domains, the largest being the days of the year (365 possible values). We want to explore security for larger domains where it is not known if effective attacks can be mounted.

Customer records as case study. We fix a running case study of an outsourced database of customer information. In industry currently, OPE is used to encrypt customer records before uploading to cloud services such as Salesforce. A client (sometimes an encryption proxy acting on its behalf) takes as input plaintext records, encrypts their attributes independently, and uploads them to a cloud service using (often) an existing API. While customer data can take on a number of forms, we focus particularly on a subset: (1) first name, (2) last name, (3) US ZIP code, and (4) date of birth. All such attributes benefit from server-side processing that takes advantage of the ability to compare plaintext ordering, and OPE in industry is used to encrypt such data for exactly this reason.

We restrict our attention to male first names, because research on inference attacks has shown that the binary “gender” attribute cannot be hidden by any PRE scheme. Partitioning database records based on gender would be a trivial preprocessing step to remove some uncertainty about the underlying plaintexts of encrypted values. To model real adversaries as closely as possible, we perform this preprocessing as well.

In these contexts, revealing plaintext order to servers enables a variety of server-side operations. Sorting is the most obvious, and for names this allows sorting alphabetically. It also allows range queries, such as finding all names starting with “A”, “B”, or “C”. Perhaps more subtly, OPE enables prefix searches

over ciphertexts, so one could search for “Dav*” and retrieve Dave, David, Davik, etc. Date of birth and ZIP codes also benefit from such range and prefix searches because prefixes have structural meaning. For example, querying “606*” gives all ZIP codes associated with the city of Chicago. Together, all this allows preserving functionality on the server-side, in some cases allowing drop-in replacement of plaintext data with encrypted without modifying server implementations at all.

Threat model. We investigate the security of OPE/ORE encrypted attributes when an adversary obtains a one-time snapshot of the encrypted database. Our attacks will target each attribute (column) independently (i.e., we won’t use information about the encrypted last name column to help improve recovering first names). Exploring the benefit of such cross-column attacks in this setting remains an interesting topic for future work. So our attacks will consider an adversary that obtains a sequence $C = (c_1, \dots, c_n)$ of encryptions of one column of data. Here $c_i = \mathcal{E}_K(m_i)$ for some unknown plaintext m_i and encryption scheme \mathcal{E} . The key K remains unknown to the attacker, and we assume it is intractable to recover it. The passive adversary can take advantage of *auxiliary information* about the distribution from which plaintexts are drawn. This is the same attack setting as considered by NKW. In contrast, the cross-column setting would give the adversary a set of encrypted columns C_1, \dots, C_k along with some auxiliary information about the distribution of each individual column as well as the joint distribution of all the columns together.

We will also look at stronger adversarial models. In Section VII we consider non-adaptive known-plaintext attacks. In these, a subset of the n plaintexts are known by the adversary. Despite not being considered in prior attacks on OPE/ORE, we believe non-adaptive known-plaintext attacks are likely to be prevalent risks in practice. As we discuss in Section VII, for example, attackers in many contexts can obtain (perhaps indirect) access to an encryption oracle before compromising the encrypted database. We will not examine chosen-plaintext attacks in great detail, but our techniques apply to chosen-plaintext attacks as well.

In all attacks the adversary finishes by outputting a list of n ciphertext, plaintext pairs indicating the attacker’s guesses. We measure success in several ways. The raw recovery rate is the fraction of n records that are mapped to the correct plaintext. For deterministic encryption schemes, this means that if we can correctly infer the mapping between the ciphertext for “Michael” and the correct plaintext, and this ciphertext accounts for (say) 4% of the database, then our raw recovery rate will be at least 4%. We therefore also report on unique recovery rate, which is the fraction of unique plaintext values recovered correctly. Finally, neither of these measures indicate how much partial information may arise, since the attacker gets no credit for mapping an encryption of the birthdate “19620105” to “19620106” despite the fact that this leaks a lot of information. We report on prefix recovery rate as the average, over the n data items, of the length of the prefix that matches between the correct plaintext and the one output by the adversary. Note that prefix accuracy for birthdates has

Dataset	# 1st names	# Last Names	Total Records
FOP (FOP)	3,862	116,677	621,662
California Muni (CALC)	3,777	59,935	255,956
Washington (WA)	3,525	67,206	228,934
Texas Compt. (TXCOM)	2,416	33,802	149,678
Florida (FL)	2,091	32,986	112,566
Maryland (MD)	2,551	36,698	111,183
Connecticut (CT)	2,016	30,623	77,613
New Jersey (NJ)	1,964	29,094	73,119
Iowa (IA)	1,734	22,616	60,035
Ohio (OH)	1,440	21,034	58,792
Texas A&M U. (TXAMU)	1,466	11,437	25,192
North Carolina (NCAR)	696	3,688	6,976
Illinois (IL)	243	1,021	1,259

Fig. 2: Unique and overall record counts for our target datasets that include first and last names.

an important caveat: all our target birth dates occurred in the twentieth century, so the baseline prefix accuracy is 25% (because the first two characters are always “19”).

Data sets. We use a number of datasets to drive our simulated attacks. We have two kinds of datasets, target datasets and, when used, auxiliary data sets. The latter is given to the adversary in the clear to provide it an empirical estimate of the target dataset’s distribution.

For target first and last name data sets, we used a mixture of municipal, state, and public university government employees, publicly published by the government. We additionally use the database dump from the Fraternal Order of Police (FOP) breach¹ See Figure 2. This includes 623,372 records on police officers for which all rows contain both first and last names, 237,392 rows contain birthdates, and 617,280 contain ZIP codes. There are 22,485 unique birthdates and 26,914 unique ZIP codes.

The distributions observed in these datasets are non-uniform. The most common first name appeared in 4.06% of records on average across all the datasets (with variance 0.29% across data sets). The most common last name appeared in 0.8% of records on average across all datasets (with variance of 0.04%). This translates to an empirical min-entropy of 4.63 bits for first names and 7 bits for last names, on average. For the FOP data set, the most common ZIP code appeared in 0.09% of the records (7.6 bits), and the most common birth date in 0.01% of records (11 bits). For reference, password leaks often indicate min-entropies of about 6–7 bits [15]. The maximum accuracy of the baseline guessing attack, in which the most frequent element of the auxiliary data is matched to every ciphertext, was 4.8% for first names and 1.2% for last names. Below we will consider first and last name datasets as subsets of the set of all alphabetical strings less than or equal to some fixed length. The effect of this is that our first and last name datasets are quite sparse.

Most of our attack simulations require auxiliary data. We restrict ourselves to publicly available data that would be easy for any attacker to obtain, and do not consider scenarios where an attacker obtains, say, an earlier version of the same

¹For privacy reasons we will not include links to these datasets, but they are available on request.

database. Our auxiliary data for experiments on first names is statistics for baby names gathered by the US Social Security Administration [7]. For our experiments, we used a year-by-year tally of the most popular American male first names for the years 1945 to 1993. For experiments on last names, we used statistics gathered by the US Census Bureau during the 2000 census on the exact frequency of last names for every person who filled out the census that year [6]. The census data included 5,023 and 151,672 distinct first and last names, respectively.

Our auxiliary data for birth dates came from the American Community Survey (ACS) 2013, a yearly survey conducted by the US Census Bureau. First, since our target dataset is a database of law enforcement union members, we used the 2013 American Community Survey (ACS) to compute a histogram for the *ages* of all respondents who marked “law enforcement” as their current line of work. Using known statistics on birth month frequency we synthesized an accurate distribution for birth dates with per-day granularity for all days between 1920 and 1999, then re-sampled it to have the same number of records as the target FOP data.

Our auxiliary data for ZIP codes is a list of the reported population of each assigned ZIP code according to the 2010 census. The frequency of a ZIP code in the auxiliary data was computed as the proportion of the total US population living there. Intuitively, it seems like the distribution of ZIP codes in a nationwide database like the FOP dump will be similar to the distribution of people into ZIP codes because more populous areas will likely have more police officers, and therefore more FOP members. However, we note that the two distributions are not particularly well-correlated and that our attacks below would probably be more effective with better auxiliary data.

For large, sparse domains like the ones considered here, it is possible for a target datum to be “un-recoverable” according to our accuracy metrics for attacks with auxiliary data, meaning its underlying plaintext does not exist in the auxiliary data. For birth dates and ZIP codes this is not the case, but it is for first and last names. On average for our datasets, above 99% of first name records are recoverable but only 91% of unique first names are. For last names, only 89% of records and 71% of unique values are recoverable on average. Below, we will not correct our results to account for this artificial cap on attack accuracy, since a real adversary would likely face this same problem. This is in line with our very conservative approach to auxiliary data overall. However, in Section V we will evaluate an attack that does not require auxiliary data, so this limitation does not apply there.

IV. THE NON-CROSSING ATTACK ON IDEAL LEAKAGE

In this section we recall the NKW cumulative attack which can work against OPE and ORE schemes that leak both frequency and order. This includes all deterministic OPE schemes, including [12, 18, 28, 31, 36].

We discuss some limitations both in terms of efficiency and accuracy of their attack, and suggest a new attack that performs significantly better. We call this the non-crossing attack. It works for any scheme for which the original NKW cumulative

attack works, and it also will be what we build off in later sections when we take advantage of more leakage and stronger adversarial models.

Attack setting. In this section, we follow NKW and consider known-ciphertext attacks with auxiliary information, but no knowledge of any plaintexts. More precisely, an attacker obtains a sequence of ciphertexts $C = (c_1, \dots, c_n)$ for $c_i = \mathcal{E}_K(m_i)$. Plaintexts may repeat and are drawn according to some (typically unknown) distribution p_m over a message space \mathcal{M} . The attacker additionally has *auxiliary information* about p_m , which in practice is simply a sequence $Z = (z_1, \dots, z_\psi)$ of plaintexts that is believed to be drawn from \mathcal{M} using a distribution close to p_m . For our datasets, it’s always the case that $n < \psi$. The attacker outputs a guess of each ciphertexts’ plaintexts, and we measure success as per the three recovery rate types defined in the previous section.

The NKW attacks. NKW describes two attacks against schemes which leak order and frequency: the sorting attack and the cumulative attack. The sorting attack simply sorts C and Z and matches the i^{th} largest element of C to the i^{th} largest element of Z . When the target data is sparse, or when $|C| \neq |Z|$ (as is the case for our datasets), this attack performs poorly, so we will not consider it further.

In the cumulative attack, NKW models the inference task as a linear sum assignment problem [16]. We find it conceptually simpler to cast the problem NKW solves in the language of graph algorithms. Let $G = (U, V, E)$ be a bipartite graph where every vertex in U corresponds to a distinct ciphertext and every vertex V corresponds to a plaintext value in the auxiliary data. Thus $n = |U|$ and $\psi = |V|$. For $u \in U$ and $v \in V$, the edge $(u, v) \in E$ is labeled according to a cost (or weight) of mapping ciphertext u to auxiliary datum v . A matching for G is a subset of edges for which no two share a common vertex. For our context, any matching is a potential decryption of some or all of the target ciphertexts. Finding a minimum cost (equivalently, maximum weight) matching should help an attacker obtain a good solution.

NKW explore the following approach to labeling edges. Given a ciphertext sequence C and a side information sequence Z , their attack lets edge $(i, j) \in U \times V$ have label given by

$$w(i, j) = |H_C(i) - H_Z(j)|^2 + |F_C(i) - F_Z(j)|^2 \quad (1)$$

The attacker seeks a minimal-cost bipartite matching of the graph G , one that contains an edge for each node in U . To do so they use the Hungarian algorithm, padding if necessary U with dummy nodes until $|U| = |V|$. The algorithm runs in time $\mathcal{O}(\psi^3)$.

NKW also briefly discuss labeling edges with only the square distance of CDFs, but they conclude that using CDF and frequency as in Equation 1 performs substantially better. We therefore only use the latter in our experiments.

Limitations of the cumulative attack. NKW show that the cumulative attack can effectively recover plaintexts, albeit only ones from relatively small domains. The largest message space they consider is 365 elements (days of the year). This has left open the question of whether plaintexts drawn from larger

domains might still be secure when encrypted under OPE. Unfortunately the performance of the attack is prohibitive for large datasets such as the last name, birthdate, and ZIP code datasets we consider in this paper.

We observe that there exists a classic greedy heuristic that works to find min-cost bipartite matchings efficiently [10, 16]. It just takes the minimum cost for each ciphertext individually, and gives a 2-approximate solution in $\mathcal{O}(n\psi)$. In experiments we have observed that this heuristic often produces fairly accurate approximations. Below we will refer to the attack that finds the matching with the greedy heuristic as “NKW greedy”.

The NKW attack, when using either the exact Hungarian algorithm or greedy approximation to find the matching, does not necessarily enforce adversarially-known ordering constraints on plaintexts. The algorithm can obtain solutions from the Hungarian algorithm that include edges that “cross”, i.e. ciphertext c is mapped to p' and ciphertext c' is mapped to p but $c < c'$ and $p' > p$. In fact, the NKW attack violated ordering constraints in every one of our experiments on real datasets. One can avoid such crossings by labeling edges with just the square distance of CDFs, but as mentioned this performs poorly. We therefore seek a way to avoid violating ordering constraints, while still taking advantage of the available frequency information.

The non-crossing attack. We introduce what we refer to as the non-crossing attack. An attacker can avoid crossings in their solution by solving a max-weight non-crossing bipartite matching problem on the graph G . Since ordering constraints are encoded into the matching algorithm, we will not include any ordering information in our edge labels. Rather, we will use the L1 distance of frequencies. For ciphertext i and auxiliary datum j , the labeling function is

$$w(i, j) = \alpha - |H_C(i) - H_Z(j)|$$

where α is a fixed constant parameter that converts a min-cost problem into a max-weight problem. In our context, any choice of $\alpha > 1$ will not change which matchings are maximum-weight, so the solution to the inference problem will be the same for any $\alpha > 1$. Different choices of α will change what that maximum weight is, but for our purposes this is inconsequential.

This approach also significantly improves computational performance over NKW: there is a well-known dynamic-programming approach that runs in time $\mathcal{O}(n\psi)$ to find the optimal non-crossing matching [32].

Results. We compare the non-crossing and cumulative attacks for our customer record datasets detailed in Section III. We could not scale the NKW’s use of the Hungarian algorithm to complete in reasonable time for last name, birthdate, or ZIP code datasets. With the largest last name dataset, we estimated the NKW attack would have taken roughly one hundred days to complete a single experiment on the well-provisioned compute cluster we used for the first name experiments. The Hungarian algorithm is highly nontrivial to parallelize, compounding its scalability issues. For last name, birthdate, and ZIP codes, we will compare our non-crossing attack to the NKW greedy

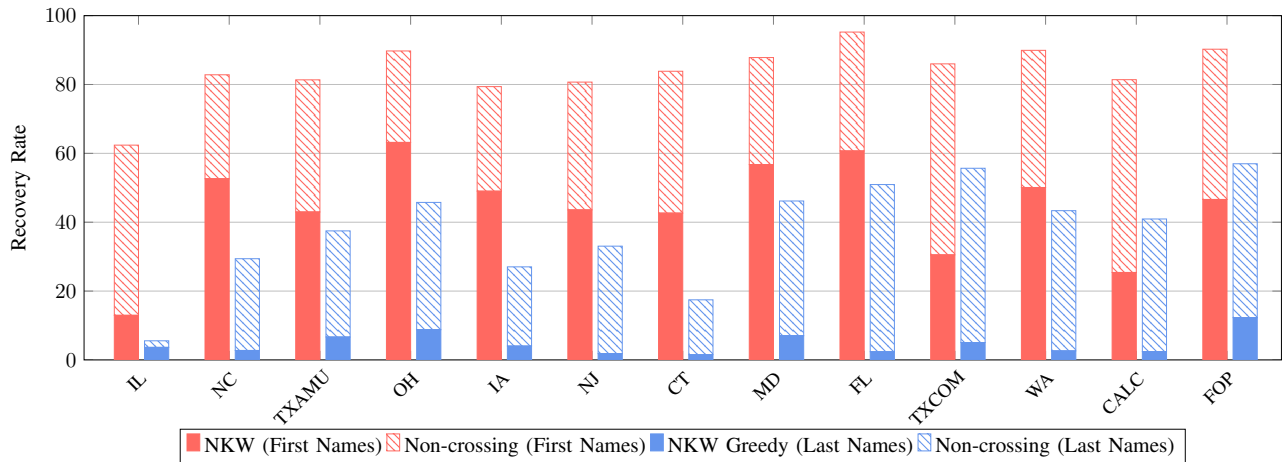


Fig. 3: Comparison of raw recovery rates of NKW and the non-crossing attack on first and last names. The non-crossing attack outperforms NKW for all target datasets.

attack. For first names, we will compare the non-crossing attack with the exact Hungarian algorithm.

Figure 3 compares the success of the two attacks for first names (left bars) and last names (right bars), showing raw recovery rates for each dataset. The non-crossing attack always performs strictly better than NKW’s cumulative attack, and some times substantially so, nearly doubling the average recovery rate (44% vs. 83%) for first names and septupling (5% vs. 38%) the average recovery rate for last names. When taken as percentage of the “recoverable” names (those plaintexts that also appear in the auxiliary data), on average we recover 84% of first names and 42% of last names.

The unusually low performance on the Illinois dataset is due to its small size — it is only about one-third of the size of the next smallest dataset. Its small size is problematic because it impacts the statistical quality of the sample. For example, the most frequent first name only occurs 34 times and the most frequent last name only occurs 10 times. Whether inference attacks on relatively small datasets can be improved is an interesting open question.

Figure 4 shows the average unique recovery rate for high-frequency first and last names. It is cumulative, so the x axis labels refer to the recovery rate for that number of top names. For example, the point (40, 95) means the attack recovered 38 of the top 40 most frequent names, on average. Our attack is especially accurate for these values. From this, we can conclude that for almost all datasets no real security guarantee can be made for high-frequency names. We do not include a separate graph for unique recovery rates for first and last names, but the overall trend is very similar.

Inference on birthdates and ZIP codes is much less accurate for both attacks: exact recovery rates were less than 2% across the board. This occurred for two reasons. First, the auxiliary information available for attackers is not as accurate a reflection of the target data distribution, compared to the names data. For example, our ZIP code data is not really a sample from the same distribution as our auxiliary data — the

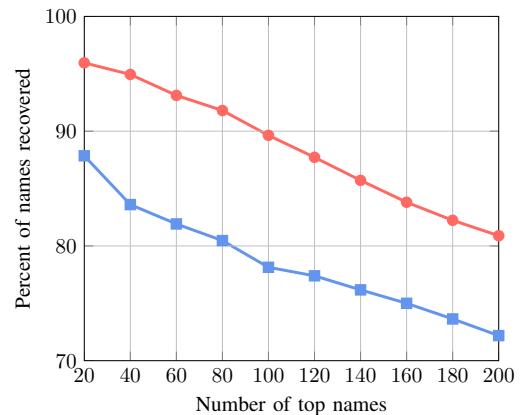


Fig. 4: Average unique recovery rates for high-frequency first and last names for non-crossing attack. The red line with circles is first names. The blue line with squares is last names.

ZIP code frequencies in our auxiliary data are proportional to the number of people living there, but the frequencies in our target dataset are proportional to FOP membership, which varies state-to-state. Second, the distributions themselves are flatter and do not have frequency “peaks” which can be easily recovered by inference attacks.

We do note that partial information is often leaked. The non-crossing attack’s prefix recovery rate was 34% for birthdates and 23% for ZIP codes. The corresponding prefix recovery rates for the NKW attack are similar. The non-crossing attack outperforms the baseline by about ten percent in overall prefix recovery for birth dates: we recovered the decade of birth for 75% of the unique birth dates in the database (16,847 out of 22,485). The non-crossing attack also recovered, on average, the first digit of a target ZIP code, giving the adversary the region of residence of the record (e.g. eastern seaboard,

midwest, southwest). In all cases, the NKW with greedy heuristic attack performs worse than non-crossing.

In terms of runtime performance, both the NKW with greedy heuristic and the non-crossing attack are reasonably fast. They require at most ten and twelve hours, respectively, to run to completion on the FOP last name dataset, the largest of any we examined. It includes $n \approx 116,000$ unique ciphertexts and $\psi \approx 151,000$ unique values in the auxiliary data. The induced bipartite graph has around 17.4 billion edges. The non-crossing attack is slower than the greedy heuristic due to the need for a backtracking step (similar to the one used in the classical edit distance algorithm) to recover the edges of the max-weight matching. It also writes to two large data structures at each iteration of the inner loop, which can be slow if the memory layout is not tuned. We optimized our implementation by exploiting the fact that the algorithm only needs to examine two rows of the dynamic programming table at a time: the one currently being written to, and the previous one. We reuse two fixed-width buffers for these rows rather than allocating an entire $n \times \psi$ matrix in memory. These two buffers are small enough to fit in the CPU’s cache, so the number of slow operations on DRAM is effectively halved. It also reduces the overall memory footprint by about 90%.

V. ATTACKING THE BCLO SCHEME

The attacks discussed in the previous section only take advantage of frequency and order information. While there exist schemes that leak only this information (or even less), they are not deployed in any real-world systems because all such schemes require either inefficient multilinear maps or client state, handling mutable ciphertexts (their value changes over time), and multiple rounds of communication to insert a ciphertext or perform a search. Instead, practitioners have widely been deploying more efficient schemes such as the classic one due to Boldyreva, Chenette, Lee, and O’Neill (BCLO) [12]. This scheme is known to leak more than just frequency and order [13], and we now show how to build highly damaging attacks that augment our non-crossing attack to exploit this additional leakage.

The BCLO scheme. The BCLO scheme realizes an OPE scheme by a recursive procedure that samples according to the hypergeometric distribution based on coins pseudorandomly derived from a secret key. The details of the construction are not important to our attacks, and so for brevity we refer readers to [12] for details. It is secure in the sense of being indistinguishable from a random order-preserving function (ROPF).

An ROPF, however, still may leak significant information about the plaintext. Boldyreva, Chenette, and O’Neill [13] analyzed ROPFs relative to a notion of security they call window one-wayness. Let $M = |\mathcal{M}|$ be the size of the domain. They show that an adversary, given the encryption of a uniformly chosen plaintext, can use the ciphertext to immediately infer with high probability that the hidden plaintext falls within a set of size $b \cdot \sqrt{M}$ for some small constant b . When b is small relative to \sqrt{M} , this means the attacker learns most of the first half of the plaintext.

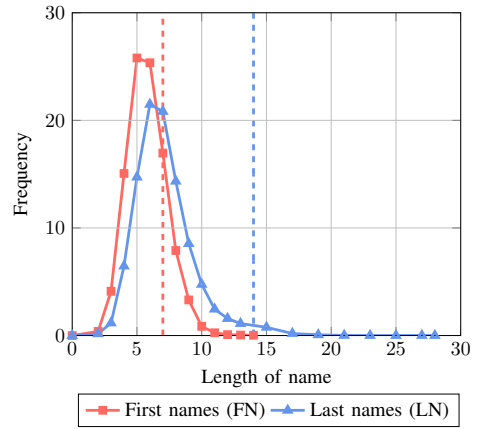


Fig. 5: Length distributions of first and last names. The right vertical line is half of the longest last name length, which roughly means every last name to the left of that line will be leaked fully. The left vertical line is the same boundary for first names.

We will show how to exploit this additional leakage in inference attacks by augmenting our non-crossing attack to take advantage of it. First, however, we discuss one security-critical issue that arises in practical use of the BCLO scheme.

The problem of padding. In previous work on OPE and ORE, the question of variable-length inputs is rarely discussed. Real systems that use OPE for variable-length plaintexts, though, must pad to preserve the functionality of being able to compare strings of different length.

For encryption schemes that reveal or preserve order, variable-length inputs represent a trickier problem than for other types of encryption, because different ways of ordering strings of characters handle variable lengths differently. For example, take the two strings “banana” and “zoo”. Viewed as English words, it is clear that “banana” is lexicographically less than “zoo”. However, some OPE algorithms (in particular [12]) only accept inputs that are integers in some set. Thus, the question becomes how to convert strings in a specified alphabet to integers *but preserve their alphabetical order*. The naive way to do this for “banana” and “zoo” is to treat them as big-endian numbers base 26, and convert to base 2 (or 10) before encrypting with OPE. It is not hard to see that this approach fails, because the base-2 number represented in base-26 as “banana” is *larger* than the corresponding base-2 number for “zoo”. Encrypted with OPE, their relative ordering would be reversed.

The solution is to right-pad all strings to a common input length (i.e., the length of the longest possible input) with the lexicographically smallest character (in our case, a space) before converting the string to an integer for encryption with OPE. This is painful when it comes to security of schemes like BCLO that leak high order bits: the padding pushes sensitive data into those bits that are leaked. Figure 5 gives histograms of the lengths of first names and last names across our datasets.

As can be seen, a bit more than half of first names and the vast majority of last names fall below half the maximal length. As we will see, this means that many or, for last names, most plaintexts will be immediately leaked when encrypted under BCLO.

In theory, one could pad to a shorter length and truncate values that are too long. Besides the obvious loss of strict order preservation for all plaintexts, this also removes the ability to decrypt all ciphertexts to the correct plaintext. In practice this truncation would necessitate appending another encryption of the plaintext to the end of the OPE ciphertext so that decryption could occur correctly, which would at least double the storage required for that column in the database. Thus, this method of handling variable-length inputs is at best much more expensive than the alternative, and at worst impossible because of, e.g., unchangeable constraints in SaaS applications.

In our experiments, we encoded first and last names as base-27 integers and padded plaintexts to the length of the longest possible plaintext with the space character, which we defined to be the lexicographically smallest character of our input alphabet. This method of encoding strings as integers may seem strange initially, but observe that the naive way of treating each byte as a base-256 digit will cause the parameters (and, by extension, the ciphertexts) to be much larger than they need to be, since only a small subset of the 256 possibilities are valid characters. As mentioned before, the longest first name appearing in our data sets is 14 characters, and the longest last name is 28 characters. Birth dates and ZIP codes are fixed-width, so padding is only required for first names and last names.

BCLO’s additional leakage. Since the BCLO scheme leaks additional information other than order and frequency, it is logical to ask how much information an adversary learns just by computing the BCLO leakage. This is not an “inference attack” as we have defined it above, because it does not use auxiliary data. It does require knowledge of the input alphabet and padding rules. Below, we will describe how to combine this attack with inference.

We first describe how to compute the leakage given a ciphertext as per [13]. Let \mathcal{M} be the message space with size $M = |\mathcal{M}|$ and let the ciphertext space be $\mathcal{C} = \{1, \dots, C\}$. Typically in implementations of BCLO one uses message and ciphertext spaces that are of size a power of two, and sets $\log C = 3 + \log M$ as suggested in [13]. Then let

$$m_c = \left\lceil \frac{M \cdot c}{C + 1} \right\rceil \quad \text{and} \quad \delta = \frac{3}{2\sqrt{M - 1}}.$$

The three in the numerator of δ is the parameter b in their lower-bound, which sets the width of the window. Roughly speaking, m_c is the ciphertext-specific point around which a window of size up to 2δ can be drawn that contains the plaintext with high probability. Of course we must truncate the window by the endpoints of the message space, meaning that m will be with high probability contained in the range

$$R_c = [\max\{1, m_c - \lfloor \delta M \rfloor\}, \min\{M, m_c + \lfloor \delta M \rfloor\}]$$

The analysis of Boldyreva et al. shows that the probability of landing in this window is negligibly far from one for

Plaintext	Ciphertext	m_c
michael	cyrzjipnouushzh	michaekypfbkfr
david	aenpse cevvpkmr	david jwbvhec
robert	emlqrynycvblqqnd	robert lweorr
john	ccnnczzzpruvjhd	johmzzzysfbunn
james	bzkxrq gzortby	james zyovtq
daniel	aelfspocabjdvjc	daniel jgaginu
richard	ekrzjmjhjxykbba	richardkfmfnwx
jose	ccqrlzzziozokby	josedzzzxfvruqq
mark	cwmlfzzzjxhklh	marjzzzzzydudv
christopher	zokwwbrbibyouo	christotnqfolw

Fig. 6: The value m_c computed for encryptions of ten first names in the California dataset.

uniformly sampled plaintexts [13]. For non-uniform message distributions, it is likely that more careful analysis could shrink the window and obtain more leakage, but we will be conservative and use the larger window for our first and last name attacks. We will use $b = 1$ for our attacks on ZIP codes and birth dates. Making the window smaller slightly increases the leakage and helps reduce the overlap in leaked ranges for nearby values, though it increases the likelihood that the true plaintext falls outside the range. Our birthdate and ZIP code datasets contained a substantial fraction of all the possible plaintexts², so there is a great deal of overlap in the leakage which, intuitively, gives the non-crossing attack more chances to match a value incorrectly. Having a smaller window mitigates this.

Abusing just the BCLO leakage. An adversary can attempt to immediately recover the plaintext for each ciphertext in a database using just the leaked range and knowledge of the input alphabet and padding rules. This approach does *not* require any auxiliary data, and requires just a few elementary computations for each ciphertext. In fact, a human attacker can easily just read off partial or even full plaintexts from the m_c values trivially computable from a ciphertext. Some examples are shown in Figure 6.

To automate this, we fix a simple heuristic that an attacker can use to guess a message given the ciphertext c . First, compute m_c and the range R_c . Let $[m_l, m_u] = R_c$. Check if any of m_l, m_c, m_u contains two consecutive³ spaces in the first half of their string representations. If none do, then forego outputting a guess. Otherwise, for (an arbitrary) one of the strings that does contain two consecutive spaces, simply take the prefix preceding these spaces as the plaintext guess. Observe that this is not *guaranteed* to be the correct plaintext as there could be other validly padded plaintexts in the range R_c .

This heuristic performs quite well, particularly for last names. On average across all our datasets, the raw recovery rate is 45% for first names and 97% for last names. Referring

²This fraction is about 56% for ZIP codes and 66% for birthdates, precluding effective use of the sorting attack from NKW, which requires almost all plaintexts to be present.

³In our datasets some names contain a single space, so we can only confirm padding after seeing two consecutive spaces.

back to Figure 5, our heuristic is able to do almost as well as predicted by simply halving the maximum length and observing the fraction of plaintexts that lie below that length: on average across datasets, 67% of first names are less than 7 characters long, and 99% of last names are less than 14 characters long. While this leakage has been known in the academic literature, we believe its severity was not understood for practical scenarios before our work.

We have described above how to directly exploit BCLO’s leakage for first and last names, but this approach is readily generalizable to any plaintext distribution containing messages of different lengths. In the case of fixed input length plaintexts, such as ZIP codes and birth dates, there is no need for padding, and so it is impossible to exactly recover any plaintexts using the heuristic above. Nevertheless significant partial information leaks that we will exploit next.

Inference attack with BCLO leakage. The heuristic above recovers less than half of first names on average, and as just mentioned cannot recover full plaintexts for fixed-input-length domains. We can however integrate the BCLO leakage into our non-crossing inference attack.

Let $G = (U, V, E)$ be a bipartite graph where every vertex in U corresponds to a unique ciphertext in C and vertices in V correspond to unique auxiliary data from Z . The sets U and V are sorted, so that the i th largest unique ciphertext is vertex u_i and likewise for v . For each ciphertext $u \in U$ where the ciphertext corresponding to u is c , the adversary computes R_c and then adds an edge (u, v) for each v whose corresponding value p falls within R_c . This excludes edges that fall outside the window for c . Each edge is weighted as before, by $\alpha - |H_C(i) - H_Z(j)|$. The adversary outputs the mappings implied by the solution to the max-weight non-crossing bipartite matching problem the graph defines.

For fixed-length inputs such as ZIP codes and birth dates, we run the inference as described above. For variable-length data such as first names and last names, we first run our heuristic attack, and then run the inference as described above only on the ciphertexts for which the heuristic fails to make a guess. The reason to do this is that many plaintexts that are not in the auxiliary data are nevertheless fully recovered by the heuristic.

Inference attack results. Combining the heuristic with inference increases first name recovery rate from 45% to 99% on average across datasets. The smallest recovery rate in any single dataset was 97%. For last names the increase from inference was negligible, as the heuristic alone already obtained 97% recovery on average. The average unique recovery rate was 90% for first names and 94% for last names, with standard deviation less than 2% for both. Our attack here recovers the vast majority of plaintext records, as well as most unique plaintexts. The few unrecovered plaintexts are ones for which the heuristic fails to retrieve them fully and they additionally do not appear in the auxiliary data (i.e., names that are both longer than average and quite rare). Of course, even for these long and rare names, a prefix of the name is nevertheless apparent to attackers.

Figure 7 summarizes the inference attack’s recovery rates for birthdates and ZIP codes, for $b = 3$ and $b = 1$. The

	Birthdates	ZIP codes
Raw, $b = 1$	91	12
Unique, $b = 1$	70	9
Raw, $b = 3$	0	4
Unique, $b = 3$	0	3

Fig. 7: Raw and unique recovery rates for birthdates and ZIP codes encrypted with BCLO. The value b refers to the window width discussed above.

attacks are much more accurate with $b = 1$. With $b = 3$ we only recover about 1.5% more ZIP codes than with the non-crossing attack. Our accuracy for birthdates with $b = 3$ does not improve compared to the non-crossing attack. With $b = 1$ the attack performed well for birthdates, recovering more than 90% of records and nearly 70% of unique values. This attack performed more modestly on ZIP codes, but we still recovered about 12% of the ZIP code records and 8% of the unique values. The BCLO leakage by itself revealed, on average, 36% of the plaintext for ZIP codes and 31% for birth dates.

That our attack performs well with $b = 1$ is surprising because the probability that the correct plaintext is in the computed window is (roughly) an inverse exponential of b . When $b = 1$ the success probability of their lower-bound attack should ostensibly be only about 0.3 for each ciphertext, but this was clearly not the case in our experiments. The probability of success for the attack in [13] is, however, only analyzed for uniformly-sampled messages, and so they may not be predictive for the non-uniform birthdate and ZIP code distributions.

VI. ATTACKING THE CLWW SCHEMES

Our results show that the BCLO scheme’s additional leakage represents a significant threat to plaintext confidentiality for real datasets. We now turn to suggestions by Chenette, Lewi, Weis, and Wu (CLWW) to provide ORE and OPE that they prove leak less than the BCLO scheme, yet remains practical [18]. We will test this empirically.

The CLWW ORE scheme. In [18] the authors construct a new ORE scheme. As with BCLO, we will for brevity omit the details of the CLWW scheme; our attacks will only abuse its leakage profile. Towards that profile, for two equal-length bit strings x, y let $\text{ind}_{\text{diff}}(x, y)$ be the index of the first bit that differs between x and y . If $x = y$ then $\text{ind}_{\text{diff}}(x, y)$ outputs $|x| + 1$. Then, a sequence of ciphertexts $C = (c_1, \dots, c_n)$ for which $c_i = \mathcal{E}_K(m_i)$ leaks order as well as, for every pair $1 \leq i < j \leq n$, the value $\text{ind}_{\text{diff}}(m_i, m_j)$. (Note that frequency of plaintexts is leaked by the latter.)

The CLWW authors argue that their scheme leaks less than BCLO for uniform messages: for only one ciphertext their scheme is semantically secure, and for n revealed ciphertexts of length k , the probability that more than $\log k$ bits leak is negligible. The BCLO scheme, in contrast, always reveals approximately one-half of the bits of a plaintext. While this is true, it may not matter in practice, as for certain plaintext

distributions the CLWW leakage may be worse (or better) than BCLO. Indeed, we will see that for ZIP codes in particular the CLWW leakage reveals nearly every bit of every plaintext.

The CLWW attack. Because the scheme leaks the index of the first bit that differs between two ciphertexts and the ordering of their underlying plaintexts, it also leaks the *value* of the plaintext bit at that index.⁴ We now detail how to take advantage of this leakage. Like the attack against BCLO in the last section, we first compute leakage, modify the bipartite graph to include only viable solutions relative to this leakage, and then apply the non-crossing attack.

In more detail, for each ciphertext, compute the leaked bit relative to every other ciphertext in the database. This gives the values of bits at certain positions in the plaintext, which we can represent as a list of pairs $\mathcal{L}(c) = ((p_0, b_0), \dots, (p_\ell, b_\ell))$. Then, initialize a bipartite graph $G = (U, V, E)$ as per the BCLO attack. The edge set E is initially empty. An edge (i, j) corresponding to ciphertext c_i and auxiliary datum z_j is added only if for each (p_k, b_k) pair in $\mathcal{L}(c_i)$, the value of the p_k bit position of z_j equals b_k . Such edges are given weight as in the non-crossing attack. Finally, the adversary solves the non-crossing bipartite matching problem and outputs the resulting guesses.

CLWW attack results. Recovery rates for our CLWW are presented in Figure 8. For first and last names, the average raw recovery rate across all datasets for the attack against CLWW is 98% and 75%. When these percentages are taken as a fraction of the “maximum” recovery rate, meaning names that cannot be matched are not included, the recovery rate for first names is 98% on average, and 85% on average for last names. Our attack against CLWW is nearly as accurate as the attack against BCLO, meaning that despite the difference in leakage profiles (and proofs by CLWW that it leaks less for uniform plaintexts), the leakage is, in effect, the same for our non-uniform plaintext distributions. For first and last names, the ORE leakage by itself removes most of the uncertainty from the ciphertexts. We can measure this explicitly by looking at the number of auxiliary data elements that have the same bits in the same positions as those that are leaked from a ciphertext. These are the auxiliary elements that will have edges to this ciphertext in the graph G above. For first names on average, there were only four such elements in the auxiliary data. For last names on average there were nineteen elements. The variance in this number was quite high for last names — there was one dataset whose average was only 1.3 elements, but another dataset whose average was 148 elements.

The recovery rates for birthdates and ZIP codes show that the CLWW attack is quite devastating, recovering almost 100% of values. This is better than the attack against BCLO, suggesting that empirically at least, there are realistic distributions for which CLWW performs poorly. The reason is that the birthdates and ZIP codes target datasets contain a large fraction of the possible plaintexts, for example the ZIP codes dataset has about 56% of the approximately 40,000 possible values.

⁴The lesser plaintext has a 0 at that position, and the greater plaintext has a 1.

In such a situation, the leakage of CLWW is very bad: 85% of all plaintext bits were immediately leaked by ind_{diff} . The non-crossing step has little trouble recovering the remaining unknown bits.

Composed OPE and ORE. CLWW also suggest that one might improve security by first encrypting with an OPE scheme and then double-encrypting using an ORE scheme. They argue that this at least inherits the security of the OPE scheme, and may do even better. We refer to this as the composed scheme, and investigate it for the case of applying first the BCLO scheme and then the CLWW scheme.

The decomposition attack. We can combine the CLWW and BCLO attacks to recover a smaller prefix of the message (than would be revealed directly by BCLO ciphertexts). At a high level the attack will simply compute the individual bits leaked by CLWW, which are now intermediate BCLO OPE ciphertext bits. We can nevertheless still use these bits to define a window within which the plaintext falls with high probability. The window now will be larger, because we must conservatively span the windows for all possibilities of the unknown intermediate ciphertext bits. In turn, larger windows mean a smaller prefix being leaked. In more detail the attack works as follows, given target ciphertext $C = \{c_0, \dots, c_n\}$ and auxiliary data $Z = \{z_0, \dots, z_\psi\}$.

In the first step, the adversary computes the CLWW leakage for each ciphertext, resulting in a list of pairs for each ciphertext $\mathcal{L}(c_i) = ((p_0, b_0), \dots, (p_\ell, b_\ell))$. Note that the bit values b_0, \dots, b_ℓ are not plaintext values, but rather the intermediate OPE ciphertext bits. Form a bit string for each ciphertext c_i by setting bit p_i to b_i and all other bits to zero. Call this bit string c_i^L . Form another bit string c_i^R similarly, except now setting bits whose positions do not appear in the list to one. The true intermediate ciphertext must lie in the range $[c_i^L, c_i^R]$.

Now the attacker computes a BCLO leakage window that contains the leakage windows for both c_i^L and c_i^R . Do so by computing the leakage for each of those two ciphertexts, and then taking the lower bound on the window for c_i^L and the upper bound on the window for c_i^R . Call this range R_{c_i} . The probability (for uniformly-sampled messages) that the message lies in this range is at least the same probability for the BCLO leakage function, since the range is computed in the same way but is larger.

Initialize a bipartite graph $G = (U, V, E)$ as before, and add edge (c_i, z_j) only if $z_j \in R_{c_i}$. Weight edges as before, by the distance in frequencies. Finally solve the non-crossing bipartite matching instance and output the associated guesses.

Decomposition attack results. Recovery rates for the decomposition attack are also presented in Figure 8. Composition does reduce attack efficacy, suggesting the CLWW intuition about composition improving security is sound, with average raw recovery rates of 85% and 44% for first and last names. This is about two percent and seven percent better for first and last names than the attacks with just order and frequency leakage discussed in Section IV. When taken as a percentage of the “maximum”, the recovery rates are, on average, 86% for

		Birthdates	ZIP codes	First names	Last names
CLWW ORE	Raw RR	98	97	98	75
	Unique RR	90	77	79	47
Composed ORE	Raw RR	86	11	85	44
	Unique RR	61	8	48	18

Fig. 8: Results of CLWW and decomposition attacks.

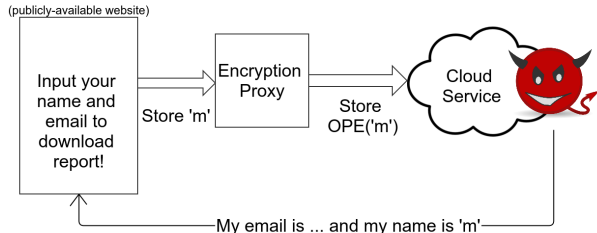


Fig. 9: A chosen-plaintext attack against cloud-hosted marketing automation

first names and 50% for last names. Composition also seems to improve security for birthdates and ZIP codes, reducing the raw recovery rates by 13% and 86% respectively. These are still much better than the corresponding numbers from Section IV. Unfortunately, this reduction in attack efficacy is probably not enough to conclude security, and it may be that one can improve on our composition attack somehow.

VII. KNOWN-PLAINTEXT ATTACKS

The threat models considered in the OPE literature thus far have excluded known- and chosen-plaintext attacks. In some use-cases of OPE/ORE, untrusted parties may not be able to know or add values. However, this assumption seems unfounded in many practical scenarios, as we now discuss.

Adversarial capabilities in practice. Recall from Section II that ORE and OPE are often deployed to encrypt customer data ultimately stored in customer relationship management systems such as Salesforce. A typical practice of companies is to allow any website visitor to download technical reports if they provide some personal information in a contact form. This technique, sometimes referred to as marketing automation, is widely used. For those companies that use encryption tools, the information entered in these forms is automatically encrypted and then stored in a cloud service used for marketing analytics.

This immediately gives rise to the ability to mount chosen-plaintext attacks. The adversary can fill out the form as desired, and later see the encryptions of this data in the database. It will be apparent to the adversary what encryptions relate to the adversarially chosen inputs, if for no other reason than the timing of submission of data and receipt of new encrypted database items. A diagram of the attack is shown in Figure 9.

There exists a folklore chosen-plaintext attack against any OPE construction that allows an adversary to learn every bit

of a plaintext from a ciphertext given only $\log(M)$ encryption oracle queries, where M is the size of the ciphertext space. Call the ciphertext the adversary wants to decrypt c_t . The attack is a simple binary search, as follows: first the adversary queries the plaintext $\frac{M}{2}$ and gets a ciphertext c_1 . If $c_1 < c_t$ it queries $\frac{M}{4}$, else it queries $\frac{3M}{4}$. It repeatedly halves the range until it finds the point p_t whose ciphertext is c_t .

Since this attack can trivially recover every plaintext from any message distribution, we will not present results for it in this work. We will instead look at the weaker assumption where the plaintexts are known but not chosen, which may make the attack even easier to mount in practice. As we show below, just a handful of known values combined with inference is enough to drastically improve success rates. It may be interesting to explore weaker versions of the chosen-plaintext setting, such as one where the adversary has a limited number of queries or must choose its queries non-adaptively. We expect these attacks to perform even better than the ones presented below, but we defer analysis and results of these attacks to future work.

The partitioning attack. We now sketch a simple, generic approach to taking advantage of known or chosen plaintexts. It can be used against any scheme that leaks at least frequency and order, and in conjunction with any of the chosen-ciphertext inference attacks against such schemes.

Assume the attacker is given not only the ciphertext sequence $C = (c_1, \dots, c_n)$ and auxiliary information $Z = (z_1, \dots, z_\psi)$, but also the plaintexts for some $q < n$ of these ciphertexts. The adversary knows the positions of these q ciphertexts, let those positions be p_1, \dots, p_q . For notational simplicity assume that $p_1 > 1$ and $p_q < n$.

The adversary can then partition the inference problem for the n ciphertexts into (at most) $q + 1$ sub-problems by splitting at each location p_i . In more detail, for each $1 \leq i \leq q$ and letting $p_0 = 1$ and $p_{q+1} = n$, define the new problem instances to be $C^i = (c_{p_{i-1}+1}, \dots, c_{p_i-1})$ and $Z^i = (z_{p_{i-1}+1}, \dots, z_{p_i-1})$. Each (C^i, Z^i) pair we can then run independently using whichever ciphertext-only inference attack we prefer. The adversary then takes the union of the resulting guesses, adds the q known plaintexts to this solution appropriately, and outputs the result as its guess.

Results. We perform experiments using the partitioning attack together with the non-crossing attack using just frequency and order leakage. While the attack increases success for the BCLO, CLWW, and decomposition attacks as well, the relative gain will be more modest as those attacks (without known plaintexts) already perform well. To enable comparison across the various names datasets, we set the number of

	Birth dates			ZIP codes		
	0.25%	2.75%	5.25%	0.25%	2.75%	5.25%
Raw RR	90	95	96	22	51	61
Unique RR	68	78	81	14	35	41

Fig. 10: Known-plaintext attacks on Birth dates and ZIP codes with order and frequency leakage. Numbers in the second row refer to the percentage of known plaintexts.

known plaintexts to be a fraction of the total number of target plaintexts. We experiment with 0.25%, 2.75%, and 5.25% of randomly-sampled plaintexts being known by the adversary. These correspond to somewhere between 5 and 300 plaintexts for 0.25% and 150 and 6,000 plaintexts for 5.25%, depending on the dataset. Recovery rates are averaged over five trials. Because the distribution is uniform over *unique* names and first and last names are long-tailed in terms of frequency, the raw frequency of the randomly-sampled known plaintexts was very low in all experiments. Intuitively, this is because only a small fraction of unique names have high frequency; most occur only once or twice. Birthdates and ZIP codes have a flatter distribution, so the sampled plaintexts had very low frequency there as well.

Figure 10 shows the results of our known-plaintext partitioning attack for 0.25%, 2.75%, and 5.25% of the unique birthdates and ZIP codes. The partitioning attack performs extremely well for birthdates. With only 0.25% of values known, the raw recovery rate jumps from less than one percent (in our attack with no known plaintexts) to nearly 90%. This jump can be attributed to the density of birth months and days for high-frequency years. Since the known plaintexts will, with high probability, reveal the year of most birth dates by upper- and lower-bounding unknown values, the non-crossing attack simply matches the days of the year in sequence. Another way of looking at this is that once the partitioning occurs, the non-crossing attack just performs a kind of sorting attack on the days in each partition. This “density” property is not specific to our datasets — any real birthdate dataset of comparable size to ours will have this density property. The partitioning attack also increases accuracy for ZIP codes substantially.

For first names, the increase in average recovery rates was modest. For 0.25% the average was 84%, only about 0.5% higher than the attack with no known plaintexts. With 5.25% the average was 87%, again a modest gain. The non-crossing attack with no known values already does quite well for first names, so having known plaintexts can only aid us in recovering very low-frequency values. For last names the known plaintexts had a bigger effect. Compared to a 38% average with no known plaintexts, having 0.25% of values known gives a 40% average, and having 5.25% of values known gives a large increase to 49% on average. The standard deviation of attacks on first names was around 7%, and for last names it was between 11% and 14%.

VIII. ATTACKING FREQUENCY-HIDING SCHEMES

The best-possible security notion, one which leaks only order and not frequency, was defined for ORE in [14]. Their

multilinear-maps-based scheme provably meets that definition, but is quite inefficient. The corresponding notion for OPE which does not allow frequency leakage was proposed by Kerschbaum [27], alongside a scheme that meets this definition. The scheme preserves order by storing state (in the form of a binary search tree) on the client. When a value is added, the tree is traversed as it would be in a standard binary tree insertion operation. If the value to be inserted is already present in the tree, randomness is used to choose a new ciphertext for the value, while preserving order.

This scheme is a stateful, mutable, and interactive protocol, and the client must store state whose size is proportional to the number of elements in the database. All these issues are inherent hurdles to deployment. Nevertheless, there may in the future be settings in which deployment is feasible, and so we’d like to understand the concrete security of the scheme. It also is the trickiest case for attackers in terms of having the least leakage. To that end, we will describe a simple attack on the scheme that targets high-frequency elements of a distribution by estimating where the ciphertexts of those elements are relative to the others.

The bucketing attack. At a high level, our bucketing attack will take advantage of the fact that on average and with a large enough sample, plaintexts with a common prefix will always appear in approximately the same “location”. For example, imagine sampling uniformly from n -bit strings. As the number of samples (call it f) increases, the probability that none of the samples less than the median begin with a zero is exponentially small in f , by a simple coin-flipping argument. If our “samples” are ciphertexts, an adversary can guess that all the underlying plaintexts of ciphertexts less than the median begin with a zero. The adversary may make some incorrect guesses but it should get many of them right.

We now generalize this intuition to arbitrary distributions by using the auxiliary data to guess these ranges. Our bucketing attack has parameters β and ℓ . The parameter β we call a “bucketing parameter”, and the parameter ℓ is the number of samples the attack will generate. Let the number of target ciphertexts be n . We start by generating ℓ independent samples of size n from the multinomial distribution defined by the probabilities of the elements in the auxiliary data. For each sample, we gather the elements into buckets based on whether they share a prefix of length β , then compute upper and lower bounds on the rank of the elements in each bucket. Essentially, we are computing a range where we expect the elements to have a shared prefix. We’re constructing the ranges based on the assumption that the auxiliary data and the target data are drawn from the same underlying distribution, so we can use random multinomial samples as guesses for the true target data. After ℓ samples, we use the mean of these bounds to define corresponding ranges for the target data. For each range in the target data we guess that the correct mapping for each ciphertext in the range is the auxiliary element of highest frequency in the corresponding bucket in the auxiliary data.

We will describe the attack more formally. Below, the notation $z[a : b]$ means the substring of z starting at index a and ending at index b , exclusive. First, the adversary is

First names				Last names			
Rank	RR	Rank	RR	Rank	RR	Rank	RR
1	95	6	0	1	69	6	100
2	100	7	0	2	100	7	0
3	92	8	99	3	100	8	76
4	94	9	92	4	90	9	100
5	100	10	0	5	100	10	0

Fig. 11: Average recovery rates (RR) for top ten first and last names in the auxiliary data with our bucketing attack. Rank refers to its position in the histogram sorted descending by frequency.

given some list of ciphertexts $C = \{c_0, c_1, \dots, c_n\}$ encrypted with a frequency-hiding OPE scheme. The adversary also has some auxiliary data $Z = \{z_0, \dots, z_\psi\}$. Call our bucketing parameter β . Next, generate ℓ multinomial samples of size n from Z . For every sample $\kappa \in \{1, \dots, \ell\}$, compute the number of distinct length- β prefixes of the elements in this sample. Call this number d and initialize d empty sets (the “buckets”) labeled B_1, \dots, B_d , each corresponding to one distinct length- β prefix. For $z_{j\kappa}$, $j \in \{1, \dots, n\}$ from sample κ , put it in the bucket corresponding to the prefix $z_{j\kappa}[0, \beta]$. The result for sample κ is $B_\kappa = \{B_1, \dots, B_d\}$. Let $\{B_1^i, \dots, B_\ell^i\}$ be buckets corresponding to the i th prefix for all samples. Call the mean of the lower bounds LB_e and the mean of the upper bounds UB_e . The adversary computes the empirical range for the i th prefix as $[LB_e, UB_e]$. Finally, let z_i^* be the highest-frequency element in bucket $B_i = [LB, UB]$. For each ciphertext c whose rank is in this range, output the mapping (z_i^*, c) .

Instantiating the attack. The parameter used to define bucket sizes in our experiments is obviously very important. If the parameter is too big, fewer elements will be recovered, because there is only one correct match possible per bucket. In the extreme case, there is only one bucket, and you can only recover the most frequent auxiliary element. If the parameter is too small, however, it is likely that many buckets will not recover any element correctly. Intuitively, this is because small differences in the auxiliary and target distribution can cause ciphertexts to be shifted into other buckets. The extreme case here is only having one element per bucket, which corresponds to the sorting attack from [35]. Therefore, recovery rate is maximized by some parameter “in the middle” between the two extremes.

Results. The prefix recovery was computed differently in this experiment than in the others. Because frequency is not leaked, the adversary may make different guesses for different ciphertexts of the same plaintext. We computed the total number of characters in the database by summing the product of the length and the count of each element, rather than simply summing the length of each element as in the other experiments.

For all experiments we chose $\ell = 100$. We chose $\beta = 2$ for first and last names, and $\beta = 1$ for birth dates and ZIP codes. For the latter two attributes, a smaller β was chosen because the distributions are closer to uniform, so our buckets must

be larger. Our raw and unique recovery rates for birthdates and ZIP codes were low. However, our prefix recovery rate was 38% for birthdates, one full character over the baseline on average. This means our hypothetical attacker can (on average) learn the decade of birth for some records in the database. Our prefix recovery rate for ZIP codes was 8%. Our attack did not perform particularly well for ZIP codes, which is unsurprising — its distribution is closer to uniform than the others. Additionally, the auxiliary and target data sets for ZIP codes are not well-correlated. Despite these obstacles, we note that we recovered the first digit of the ZIP code for about 160,000 of the 617,000 records in the database. Since the first digit of a ZIP code designates a region of the United States, the adversary recovers the region of residence for a database record about two and a half times better than baseline.

For first names, on average we recovered 26% of plaintexts. For last names, the average recovery rate was 6%. Our overall results are aided by the shape of the first name and last name distributions. Specifically, there are a few high-frequency values and a long tail of lower-frequency values. Nonetheless, our attack outperforms the baseline accuracy for first and last names by roughly six and eight times, respectively.

Recovery rates for the high-frequency names are presented in Figure 11. For first names, we recovered the name “michael” with 95% accuracy on average. For last names, we recovered the name “johnson” with 100% accuracy in all datasets. The two last names and three first names with 0% recovery are due to collisions in the “bucket” for their prefix. For example, if $\beta = 1$ the name “mark” will never be recovered, because the higher-frequency name “michael” is in the same bucket. The overall recovery rate for the top 10 most frequent first names was 71%, and for the top 10 most frequent last names was 78%.

This attack having high accuracy on frequent plaintexts makes sense, as the bucketing attack is designed in part to recover high-frequency elements. Stepping back, though, we should reflect on the implication of these results for the security of *any* OPE or ORE. Our results show that high-frequency elements encrypted with OPE or ORE can be reliably recovered (with probability many times better than the baseline) *even if their frequency is not leaked*. This raises questions about the fitness of OPE and ORE for the highly non-uniform distributions that arise in practice.

IX. RELATED WORK

Property-revealing encryption schemes. The study of encryption schemes that reveal order was initiated by Agrawal et al., who constructed a scheme that preserved the order property for numeric inputs, but lacked provable guarantees. [8] Property-revealing encryption methods with formal reductionist security guarantees for outsourced data originate with Bellare, Boldyreva, and O’Neill [11]. Their work introduced deterministic encryption (DTE), which is a fundamental cryptographic building block for such systems. Order-preserving encryption (OPE) was introduced to support SQL queries that require ordering of entries, and first studied formally by

BCLO [12] and its window one-wayness security was subsequently analyzed by Boldyreva, Chenette, and O’Neill [13]. We use their analysis in our attacks against the BCLO scheme discussed in Section V.

In [13], it was shown that any OPE scheme provably leaking only frequency and order must have exponentially large ciphertexts. Popa et al. [36] extended this negative result to cover stateful, interactive schemes that have immutable ciphertexts. They also introduced an interactive OPE scheme with mutable ciphertexts that leaks at most order and frequency information. Kerschbaum and Schröpfer propose a more efficient interactive OPE scheme [28]. Our attacks in Section IV work against these schemes. Later, Kerschbaum proposed a frequency-hiding, interactive OPE scheme [27], our in Section VIII applies to this scheme.

ORE was introduced by Boneh et al. [14] to provide better security, while still allowing order comparison. Their scheme relies on multilinear maps [22], making them currently inefficient [30] and potentially insecure given doubt about the validity of cryptographic hardness assumptions related to multilinear maps [24,33,34]. Our attack in Section VIII would affect these schemes. CLWW provide a practical ORE scheme that also leaks frequency [18], our attacks in Section VI affect this scheme, as well as their suggestion of composing OPE with ORE.

Systems using ORE/OPE. Several companies and services provide claims of encryption that preserves functionality, including CipherCloud and Skyhigh Networks [1–3,5]. However, CryptDB [37,39] was the first system in the academic literature that introduced a scheme for running a large subset of the SQL language on top of an encrypted database (EDB) using DTE and OPE encryption, as well as standard encryption. Since then, popularity of EDBs has increased, with many companies such as IQrypt [2] and SAP [4] producing their own CryptDB-inspired solutions.

Attacks against property-revealing encryption. Inference attacks were first considered against searchable symmetric encryption [19,40] by Islam, Kuzu, and Kantarcioglu [25] with subsequent improvements and investigation of active attacks by Cash et al. [17]. Naveed et al. [35] were the first to consider inference attacks against CryptDB-style EDBs; we discussed their attacks in detail in Section IV.

Concurrent work. In a concurrent and independent work, Durak, DuBuisson, Cash [20] show how the sorting attack from NKW [35] can be extended to the multi-column case, when two or more columns contain correlated data. They demonstrate that multi-column attacks are more devastating than separate attacks on individual columns. They also show attacks on non-ideal OPE/ORE schemes that leak more than the order and frequency of the ciphertexts, including the BCLO scheme. Their primary focus is sparse datasets such as GPS coordinates, and do not look at names, birthdates, or ZIP codes as we consider in our running case study in this work.

Durak et al. do not exploit auxiliary information, and instead rely only on leakage, whereas our attacks show how to both exploit leakage by itself and to augment an auxiliary-

information-using inference attack. Their attacks therefore only provide approximate recovery of plaintexts (e.g., determining that a plaintext lies within some 10km radius), whereas our attacks recover entire plaintexts, and in many cases fully recover most records in a database.

REFERENCES

- [1] “Ciphercloud,” <http://www.ciphercloud.com/>.
- [2] “IQrypt: encrypt and query your database,” <http://www.iqrypt.com/>.
- [3] “Perspecsys: A blue coat company,” <http://perspecsys.com/>.
- [4] “Sap seced project,” <https://www.sics.se/sites/default/files/pub/andreasschaad.pdf>.
- [5] “Skyhigh Networks,” <https://www.skyhighnetworks.com/>.
- [6] “US census bureau name statistics,” <https://www.ssa.gov/OACT/babynames/>.
- [7] “US social security name statistics,” <https://www.ssa.gov/OACT/babynames/>.
- [8] R. Agrawal, J. Kiernan, R. Srikant, and Y. Xu, “Order preserving encryption for numeric data,” in *Proceedings of the 2004 ACM SIGMOD international conference on Management of data*. ACM, 2004, pp. 563–574.
- [9] A. Arasu, S. Blanas, K. Eguro, R. Kaushik, D. Kossmann, R. Ramamurthy, and R. Venkatesan, “Orthogonal security with cipherbase.” in *CIDR*. Citeseer, 2013.
- [10] D. Avis, “A survey of heuristics for the weighted matching problem,” *Networks*, vol. 13, no. 4, pp. 475–493, 1983. [Online]. Available: <http://dx.doi.org/10.1002/net.3230130404>
- [11] M. Bellare, A. Boldyreva, and A. O’Neill, “Deterministic and efficiently searchable encryption,” in *Advances in Cryptology-CRYPTO 2007*. Springer, 2007, pp. 535–552.
- [12] A. Boldyreva, N. Chenette, Y. Lee, and A. O’Neill, “Order-preserving symmetric encryption,” in *Advances in Cryptology-EUROCRYPT 2009*. Springer, 2009, pp. 224–241.
- [13] A. Boldyreva, N. Chenette, and A. O’Neill, “Order-preserving encryption revisited: Improved security analysis and alternative solutions,” in *Advances in Cryptology-CRYPTO 2011*. Springer, 2011, pp. 578–595.
- [14] D. Boneh, K. Lewi, M. Raykova, A. Sahai, M. Zhandry, and J. Zimmerman, “Semantically secure order-revealing encryption: Multi-input functional encryption without obfuscation,” in *Advances in Cryptology-EUROCRYPT 2015*. Springer, 2015, pp. 563–594.
- [15] J. Bonneau, “The science of guessing: analyzing an anonymized corpus of 70 million passwords,” in *2012 IEEE Symposium on Security and Privacy*. IEEE, 2012, pp. 538–552.
- [16] R. Burkard, M. Dell’Amico, and S. Martello, *Assignment Problems*. Society for Industrial and Applied Mathematics, 2012. [Online]. Available: <http://epubs.siam.org/doi/abs/10.1137/1.9781611972238>
- [17] 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.
- [18] N. Chenette, K. Lewi, S. A. Weis, and D. J. Wu, “Practical order-revealing encryption with limited leakage,” in *23rd International Conference on Fast Software Encryption (FSE)*, 2016, <http://eprint.iacr.org/2015/1125>.
- [19] R. Curtmola, J. Garay, S. Kamara, and R. Ostrovsky, “Searchable symmetric encryption: improved definitions and efficient constructions,” *Journal of Computer Security*, vol. 19, no. 5, pp. 895–934, 2011.
- [20] F. B. Durak, T. M. DuBuisson, and D. Cash, “What else is revealed by order-revealing encryption?” *Cryptology ePrint Archive*, Report 2016/786, 2016, <http://eprint.iacr.org/2016/786>.
- [21] S. Gallagher, “MS researchers claim to crack encrypted database with old simple trick,” <http://arstechnica.com/security/2015/09/ms-researchers-claim-to-crack-encrypted-database-with-old-simple-trick/>.

- [22] S. Garg, C. Gentry, and S. Halevi, "Candidate multilinear maps from ideal lattices," in *Annual International Conference on the Theory and Applications of Cryptographic Techniques*. Springer, 2013, pp. 1–17.
- [23] W. He, D. Akhawe, S. Jain, E. Shi, and D. Song, "Shadowcrypt: Encrypted web applications for everyone," in *Proceedings of the 2014 ACM SIGSAC Conference on Computer and Communications Security*. ACM, 2014, pp. 1028–1039.
- [24] Y. Hu and H. Jia, "Cryptanalysis of ggh map," *Cryptology ePrint Archive*, Report 2015/301, 2015, <http://eprint.iacr.org/2015/301>.
- [25] M. Islam, M. Kuzu, and M. Kantarcioglu, "Access pattern disclosure on searchable encryption: Ramification, attack and mitigation," *Network and Distributed System Security Symposium (NDSS'12)*, 2012.
- [26] S. Keelveedhi, M. Bellare, and T. Ristenpart, "DupLESS: server-aided encryption for deduplicated storage," in *Presented as part of the 22nd USENIX Security Symposium (USENIX Security 13)*, 2013, pp. 179–194.
- [27] F. Kerschbaum, "Frequency-hiding order-preserving encryption," in *Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security*. ACM, 2015, pp. 656–667.
- [28] F. Kerschbaum and A. Schröpfer, "Optimal average-complexity ideal-security order-preserving encryption," in *Proceedings of the 2014 ACM SIGSAC Conference on Computer and Communications Security*. ACM, 2014, pp. 275–286.
- [29] 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 *23rd USENIX Security Symposium (USENIX Security 14)*, 2014, pp. 33–48.
- [30] K. Lewi, A. J. Malozemoff, D. Apon, B. Carmer, A. Foltzer, D. Wagner, D. W. Archer, D. Boneh, J. Katz, and M. Raykova, "5Gen: A framework for prototyping applications using multilinear maps and matrix branching programs," *Cryptology ePrint Archive*, Report 2016/619, 2016, <http://eprint.iacr.org/2016/619>.
- [31] T. Malkin, I. Teranishi, and M. Yung, "Order-preserving encryption secure beyond one-wayness," *Advances in Cryptology – ASIACRYPT 2014: 20th International Conference on the Theory and Application of Cryptology and Information Security*, vol. 2014.
- [32] F. Malucelli, T. Ottmann, and D. Pretolani, "Efficient labelling algorithms for the maximum noncrossing matching problem," *Discrete Applied Mathematics*, vol. 47, no. 2, pp. 175–179, 1993.
- [33] L. D. Martin Albrecht, Shi Bai, "A subfield lattice attack on overstretched ntru assumptions: Cryptanalysis of some fhe and graded encoding schemes," *Cryptology ePrint Archive*, Report 2016/127, 2016, <http://eprint.iacr.org/2016/127>.
- [34] E. Miles, A. Sahai, and M. Zhandry, "Annihilation attacks for multilinear maps: Cryptanalysis of indistinguishability obfuscation over ggh13," *Cryptology ePrint Archive*, Report 2016/147, Tech. Rep., 2016.
- [35] M. Naveed, S. Kamara, and C. V. Wright, "Inference attacks on property-preserving encrypted databases," in *Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security*. ACM, 2015, pp. 644–655.
- [36] R. A. Popa, F. H. Li, and N. Zeldovich, "An ideal-security protocol for order-preserving encoding," in *Security and Privacy (SP), 2013 IEEE Symposium on*. IEEE, 2013, pp. 463–477.
- [37] R. A. Popa, C. Redfield, N. Zeldovich, and H. Balakrishnan, "Cryptdb: protecting confidentiality with encrypted query processing," in *Proceedings of the Twenty-Third ACM Symposium on Operating Systems Principles*. ACM, 2011, pp. 85–100.
- [38] R. A. Popa, E. Stark, S. Valdez, J. Helfer, N. Zeldovich, and H. Balakrishnan, "Building web applications on top of encrypted data using mylar," in *11th USENIX Symposium on Networked Systems Design and Implementation (NSDI 14)*, 2014, pp. 157–172.
- [39] R. A. Popa, N. Zeldovich, and H. Balakrishnan, "Guidelines for using the cryptdb system securely," *Cryptology ePrint Archive*, Report 2015/979, 2015, <http://eprint.iacr.org/2015/979>.
- [40] D. X. Song, D. Wagner, and A. Perrig, "Practical techniques for searches on encrypted data," in *Security and Privacy, 2000. S&P 2000. Proceedings. 2000 IEEE Symposium on*. IEEE, 2000, pp. 44–55.