

Partitioned Data Security on Outsourced Sensitive and Non-sensitive Data

Sharad Mehrotra¹, Shantanu Sharma¹, Jeffrey D. Ullman², and Anurag Mishra¹

¹University of California, Irvine. ²Stanford University, USA.

Abstract—Despite extensive research on cryptography, secure and efficient query processing over outsourced data remains an open challenge. This paper continues along the emerging trend in secure data processing that recognizes that the entire dataset may not be sensitive, and hence, non-sensitivity of data can be exploited to overcome limitations of existing encryption-based approaches. We propose a new secure approach, entitled query binning (QB) that allows non-sensitive parts of the data to be outsourced in clear-text while guaranteeing that no information is leaked by the joint processing of non-sensitive data (in clear-text) and sensitive data (in encrypted form). QB maps a query to a set of queries over the sensitive and non-sensitive data in a way that no leakage will occur due to the joint processing over sensitive and non-sensitive data. Interestingly, in addition to improve performance, we show that QB actually strengthens the security of the underlying cryptographic technique by preventing size, frequency-count, and workload-skew attacks.

I. INTRODUCTION

The last two decades have witnessed the development of secure and privacy-preserving encryption-based [1]–[3] or secret-sharing-based [4]–[7] techniques to realize the database as a service model. Despite significant progress, a cryptographic approach that is both *secure* (i.e., no leakage of sensitive data to the adversary) and *efficient* (in terms of time) simultaneously has proved to be very challenging. Broadly, work on cryptography to support secure outsourcing has taken the following directions:

Techniques that support strong security guarantees. The leading example of which is fully homomorphic encryption [1], which when mixed with oblivious-RAM (ORAM) [3], offers possibly amongst the most secure mechanisms. However, such mechanisms incur high computational overhead.

Techniques that do not depend on the data encryption but provide strong security, especially, information-theoretic security, by distributing a value in the form of the secret-shares to non-colluding clouds. Shamir’s secret-sharing [4],

distributed point functions (DPF) [6], and function secret-sharing (FSS) [7] are a few examples of such techniques. Such methods often limit the type of operations one can perform while imposing high overhead in terms of communication.

Techniques that try to support a wide range of operations including index-based retrieval or joins, such as CryptDB [8], Arx [9], searchable encryption [2], and cryptographic indexes [10]. Such techniques often trade security for performance; for instance, techniques that depend on deterministic and order-preserving encryptions, traversal of the index by the cloud, or leakage of the searching token do not offer strong security. Papers [11], [12] show that order-preserving and deterministic encryption techniques when used together, on a dataset in which the entropy of the values is not high enough can leak the entire data in clear-text to an attacker through frequency analysis on the encrypted data.

Techniques/systems that exploit secure hardware (Intel Software Guard Extensions (SGX) [13]), e.g., M2R [14], VC3 [15], Opaque [16], and EnclaveDB [17]. Such techniques also leak information during a query execution due to different attacks on SGX (e.g., cache-line, branch shadowing, and page-fault attacks [18], [19]) and are significantly slow when overcoming these attacks using ORAM-based computations or emerging architectures such as T-SGX [20] or Sanctum [21].

Given the state of the research, this paper explores a radically different approach to secure outsourcing that scales cryptographic mechanisms using database techniques while providing strong security guarantees. Our work is motivated by recent works on the hybrid cloud that has exploited the fact that for a large class of application contexts, data can be partitioned into sensitive and non-sensitive components [22]–[24]. Such a classification, which is common in industries for secure computing [25], [26] and done via appropriately using existing techniques surveyed in [27]; for example, (i) inference detection using graph-based semantic data modeling [28], (ii) user-defined relationships between sensitive and non-sensitive data [29], (iii) constraints-based mechanisms, (iv) sensitive patterns hiding using sanitization matrix [30], and (v) common knowledge-based association rules [31]. However, it is important to mention here that non-sensitive data can, over time, become sensitive and/or lead to inferences about sensitive data. This is an inevitable risk of the approaches that exploit sensitive data classification. Note that all the above-mentioned work based on sensitive/non-sensitive classification make a similar assumption. Indeed, another way to view this assumption is that today, cloud solutions, already outsource

Accepted in IEEE International Conference on Data Engineering (ICDE), 2019. For the final version, please refer to the conference proceeding. Anurag Mishra contributed in implementing the proposed algorithms. This material is based on research sponsored by DARPA under agreement number FA8750-16-2-0021. The U.S. Government is authorized to reproduce and distribute reprints for Governmental purposes notwithstanding any copyright notation thereon. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of DARPA or the U.S. Government. This work is partially supported by NSF grants 1527536 and 1545071.

The authors are thankful to the PC chairs, vicechairs, and reviewers for helping in improving the presentation of the paper.

databases without encryption and are risking the loss of not just non-sensitive data but also sensitive data.

Based on data classification into sensitive and non-sensitive data, secure solutions for hybrid cloud have been developed [22]–[24]. These solutions outsource only non-sensitive data and enjoy both the benefits of the public cloud as well as strong security guarantees (without revealing sensitive data to an adversary). While these techniques provide an effective and secure solution, they are, however, requiring data owners to maintain potentially unbounded storage locally and also suffer from significant inter-cloud communication overheads.

Our goal, in this paper, is to explore how sensitive and non-sensitive classification can be exploited by secure data processing techniques that store the entire data in the public cloud to bring new efficiencies to secure data processing. In particular, in the envisioned public cloud model, data is stored in a partitioned way – sensitive data is secured using any existing cryptographic technique (unlike the hybrid cloud solution where the owner stores the sensitive data) and non-sensitive data resides in plaintext. Query processing is also split into encrypted and plaintext query processing. We refer to this as *partitioned computing*. Unlike the case of the hybrid cloud, when implementing partitioned computing at a public cloud, data processing performed on the sensitive and non-sensitive parts of the data may reveal exact encrypted tuples and cleartext tuples that satisfy the query to the adversary. Consequently, this leads to inferences about sensitive data, which will be explained in detail in §II.

We define a security model (§III) that formally states what it means to be secure in partitioned computing. We then develop a query binning (QB) approach that realizes secure partitioned computing for selection queries. We focus on selection queries for several reasons. First, selection queries are important in their own right. For instance, several key-value stores (*e.g.*, Amazon Dynamo) and document stores (*e.g.*, MongoDB) focus exclusively on selection queries (with limited support for joins). Furthermore, most cryptographic research has also focused on selection queries [1]–[3], [6]. Since our goal is to speed up existing cryptographic techniques (and not to extend their functionality and make them resilient against attacks, such as order-revealing, inferences from deterministic encryptions, leakages from SGX, and different side-channel attacks [11], [12], [19]), we focus on selection queries. Nonetheless, there are recent work on cryptographic joins [32] and also on joins using SGX [16]. These approaches, however, are not yet practical, *e.g.*, from the efficiency perspective, Opaque [16] takes 89 seconds to execute a selection query on a dataset of size 700MB. The same query takes about 0.2 milliseconds over cleartext processing. Also, systems, *e.g.*, Opaque, support limited operations (only primary-to-foreign key joins) and, furthermore, leaks information due to cache-line, page table-based, branch shadowing, and output-size attacks [18], [19]. Many of these attacks can be overcome with expensive ORAM techniques, and the QB approach alongside such approaches can be exploited to improve efficiency.

We show two interesting effects of using QB: (i) By avoid-

ing cryptographic processing on non-sensitive data, *the joint cost of communication and computation of QB is significantly less than the computation cost of a strongly secure cryptographic technique*¹ (*e.g.*, homomorphic encryptions, DPF [6], or secret-sharing-based technique [5]) on the entire encrypted data; and hence, QB improves the performance of strong cryptographic techniques over a large-scale dataset (§V). (ii) *QB provides an enhanced security by preventing several attacks such as output size, frequency-count, and workload-skew attacks, even when the underlying cryptographic technique is susceptible to such attacks (§VI).*

Contributions. The primary contributions of this paper are: (i) A formal definition of *partitioned data security* when jointly processing sensitive and non-sensitive data (§III). (ii) An efficient QB approach (§IV) that guarantees partitioned data security, supporting cloud-side-indexes, and that can be built on top of any cryptographic technique. (iii) An analytical model and experimental validation to show the effectiveness of QB over a strong secure cryptographic technique (§V). (iv) A weak cryptographic technique (*e.g.*, cloud-side indexable techniques [9], [10]) becomes secure and efficient when mixed with QB (§VI).

Full version. [34] provides the full version of this paper. The full version provides: (i) formal security and computational complexity proofs of QB, (ii) extensions of QB to deal with non-identical searchable attribute-based column-level sensitivity, join, and range queries, (iii) some additional experiments to show insert and the use of indexable cryptographic techniques, and (iv) an analytical formal security model to compare QB with a pure cryptographic technique under different conditions and different security levels such as preventing size, frequency-count, and workload-skew attacks. QB can also be extended to support group-by aggregation queries as well; however, extending it to support nested queries is more complex and will need a significant extension.

Related work on secure selection queries. Broadly, existing research on secure selection query execution techniques can be classified into four categories, as follows: (i) *Encryption-based techniques* examples of which include order-preserving encryption, deterministic encryption, homomorphic encryption [1], searchable encryption [2], and ORAM [3]. (ii) *Secret-sharing [4] based techniques* that include DPF [6], FSS [7], and [5]. (iii) *Trusted-hardware-based techniques* that include [14]–[17]. (iv) *Sensitivity-based techniques*. MapReduce [22], [23] and SQL data processing [24]. Both MapReduce and SQL

¹QB trades off increased communication costs for executing queries, while reducing very significantly cryptographic operations. This tradeoff significantly improves performance, especially, when using cryptographic mechanisms, *e.g.*, fully homomorphic encryption that takes several seconds to compute a single operation [33], secret-sharing-based techniques that take a few seconds [5], or techniques such as bilinear maps that take over 1.5 hours to perform joins on a dataset of size less than 10MB [32]. When considering such cryptography, increased communication overheads are fully compensated by the savings. A similar observation, albeit in a very different context was also observed in [23] in the context of MapReduce, where overshuffling to prevent the adversary to infer sensitive keys in the context of hybrid cloud was shown to be significantly better compared to private side operations.

execution solutions work on the principle of sensitivity-based data partitioning over the hybrid cloud.

Each of the above strategies has resulted in corresponding systems that support secure data processing; *e.g.*, CryptDB [8], Arx [9], and Opaque [16] are some novel encryption-based systems. Likewise, Microsoft Always Encrypted, Oracle 12c, Amazon Aurora, and MariaDB are industrial secure encrypted databases. DSSE-based SDB [35] is a secret-sharing and encryption-based system while Arx [9] and Opaque [16] work on the data sensitivity principle.

These systems/techniques are unable to prevent one or more of the following attacks: (i) size attack, *i.e.*, an adversary having some background knowledge can deduce the full/partial outputs by simply observing the output sizes [16]; (ii) frequency attack, *i.e.*, an adversary can deduce how many tuples have an identical value [11]; (iii) workload-skew attack, *i.e.*, an adversary, having the knowledge of frequent selection queries by observing many queries, can estimate which encrypted tuples potentially satisfy the frequent section selection queries; (iv) access-pattern attack, *i.e.*, addresses of encrypted tuples that satisfy the query. Note that computationally expensive and access-pattern-hiding cryptographic techniques (*e.g.*, PIR, ORAM, DSSE, and secret-sharing) can prevent the size, frequency-count, and workload-skew attacks *only* on *non-skewed and non-deterministically encrypted* datasets. To the best of our knowledge, there is no cryptographic technique that prevents all the four attacks on a *skewed dataset*. Table I shows notations used in this paper.

II. PARTITIONED COMPUTATION

In this section, we first define more precisely what we mean by partitioned computing, illustrate how such a computation can leak information due to the joint processing of sensitive and non-sensitive data, discuss the corresponding security definition, and finally discuss system and adversarial models under which we will develop our solutions.

The Partition Computation Model. We assume the following two entities in our model:

A trusted database (DB) owner who divides a relation R having attributes, say A_1, A_2, \dots, A_n , into the following two relations based on row-level data sensitivity: R_s and R_{ns} containing all sensitive and non-sensitive tuples, respectively. The DB owner outsources the relation R_{ns} to a public cloud. The tuples of the relation R_s are encrypted using any existing non-deterministic encryption mechanism before outsourcing to the same public cloud. In our setting, the DB owner has to store metadata such as searchable values and their frequency counts, which will be used for appropriate query formulation. The DB owner is assumed to have sufficient storage for such metadata, and also computational capabilities to perform encryption and decryption. The size of metadata is smaller than the size of the original data.

The untrusted public cloud that stores the databases, executes queries, and provides answers.

Let us consider a query q over the relation R , denoted by $q(R)$. A partitioned computation strategy splits the exe-

Notations	Meaning
$ S $ (or $ NS $)	Number of sensitive (or non-sensitive) data values
R_s (or R_{ns})	Sensitive (or non-sensitive) parts of a relation R
s_i (or ns_j)	i^{th} sensitive (or j^{th} non-sensitive) value
SB (or NSB)	The number of sensitive (or non-sensitive) bins
SB_i (or NSB_i)	i^{th} sensitive (or non-sensitive) bin
$ SB = y$ (or $ NSB = x$)	Sensitive (or non-sensitive) values in a sensitive (or non-sensitive) bin or the size of a sensitive (or non-sensitive) bin
$q(w)$	A query, q , for a predicate w
$q(W_{ns})(R_{ns})$	A query, q , for a set, W_{ns} , of predicates in clear-text over R_{ns}
$q(W_s)(R_s)$	A query, q , for a set, W_s , of predicates in encrypted form over R_s
$q(W)(R_s, R_{ns})[A]$	A query, q , for a set, W , of values, searching on the attribute, A , of the relations R_s and R_{ns} , where $W = W_s \cup W_{ns}$
$E(t_i)$	i^{th} encrypted tuple

Table I: Notations used in the paper.

	EId	FirstName	LastName	SSN	Office	Dept
t_1	E101	Adam	Smith	111	1	Defense
t_2	E259	John	Williams	222	2	Design
t_3	E199	Eve	Smith	333	2	Design
t_4	E259	John	Williams	222	6	Defense
t_5	E152	Clark	Cook	444	1	Defense
t_6	E254	David	Watts	555	4	Design
t_7	E159	Lisa	Ross	666	2	Defense
t_8	E152	Clark	Cook	444	3	Design

Figure 1: A relation: *Employee*.

cution of q into two independent subqueries: $q(R_s)$: a query to be executed on the encrypted sensitive relation R_s , and $q(R_{ns})$: a query to be executed on the non-sensitive relation R_{ns} . The final result is computed (using a query q_{merge}) by appropriately merging the results of the two subqueries at the DB owner side. In particular, the query q on a relation R is partitioned, as follows: $q(R) = q_{merge}(q(R_s), q(R_{ns}))$.

EId	SSN
E101	111
E259	222
E199	333
E152	444
E254	555
E159	666

(a) A sensitive relation: *Employee1*.

	EId	FirstName	LastName	Office	Dept
t_1	E101	Adam	Smith	1	Defense
t_4	E259	John	Williams	6	Defense
t_5	E152	Clark	Cook	1	Defense
t_7	E159	Lisa	Ross	2	Defense

(b) A sensitive relation: *Employee2*.

	EId	FirstName	LastName	Office	Dept
t_2	E259	John	Williams	2	Design
t_3	E199	Eve	Smith	2	Design
t_6	E254	David	Watts	4	Design
t_8	E152	Clark	Cook	3	Design

(c) A non-sensitive relation: *Employee3*.

Figure 2: Three relations obtained from *Employee* relation.

Let us illustrate partitioned computations through an example. **Example 1.** Consider an *Employee* relation, see Figure 1. Note that the notation t_i ($1 \leq i \leq 8$) is not an attribute of the relation; we used this to indicate the i^{th} tuple. In this relation, the attribute *SSN* is sensitive, and furthermore, all tuples of employees for the *Dept* = “Defense” are sensitive. In such a case, the *Employee* relation may be stored as the following three relations: (i) *Employee1* with attributes *EId* and *SSN* (see Figure 2a); (ii) *Employee2* with attributes *EId*, *FirstName*, *LastName*, *Office*, and *Dept*, where *Dept* = “Defense” (see Figure 2b); and (iii) *Employee3* with attributes *EId*, *FirstName*, *LastName*, *Office*, and *Dept*, where *Dept* <> “Defense” (see Figure 2c). Since the relations *Employee1* and *Employee2* (Figures 2a and 2b) contain only sensitive data,

these two relations are encrypted before outsourcing, while *Employee3* (Figure 2c), which contains only non-sensitive data, is outsourced in clear-text. We assume that the sensitive data is strongly encrypted such that the property of *ciphertext indistinguishability* is achieved. Thus, the two occurrences of E152 have two different ciphertexts.

Consider a query q : *SELECT FirstName, LastName, Office, Dept from Employee where FirstName = John*. In the partitioned computation, the query q is partitioned into two subqueries: q_s that executes on *Employee2*, and q_{ns} that executes on *Employee3*. q_s will retrieve the tuple t_4 while q_{ns} will retrieve the tuple t_2 . q_{merge} in this example is simply a union operator. Note that the execution of the query q will also retrieve the same tuples. However, such a partitioned computation, if performed naively, leads to inferences about sensitive data from non-sensitive data. Before discussing inference attacks, we first present the adversarial model.

Adversarial Model. We assume an honest-but-curious (HBC) adversary [36], which is considered in the standard setting for security in the public cloud that is *not trustworthy*. An HBC adversarial public cloud stores an outsourced dataset without tampering, correctly computes assigned tasks, and returns answers; however, it may exploit side knowledge (e.g., query execution, background knowledge, and the output size) to gain as much information as possible about the sensitive data. Furthermore, the HBC adversary can eavesdrop on the communication channels between the cloud and the DB owner and that may help in gaining knowledge about sensitive data, queries, or results; hence, a secure channel is assumed. In our setting, the adversary has full access to the following:

All the non-sensitive data. For example, for the *Employee* relation in Example 1, an adversary knows the complete *Employee3* relation (refer to Figure 2c).

Auxiliary/background information of the sensitive data. The auxiliary information [11], [12] may contain metadata, schema of the relation, and the number of tuples in the relation (note that having an adversary with the auxiliary information is also considered in literature). In Example 1, the adversary knows that there are two sensitive relations, one of them containing six tuples and the other one containing four tuples, in the *Employee1* and the *Employee2* relations; Figures 2a and 2b. In contrast, the adversary is not aware of the following information before the query execution: how many people work in a specific sensitive department, is a specific person working only in a sensitive department, only in a non-sensitive department, or both.

Adversarial view. When executing a query, an adversary knows which encrypted sensitive tuples and cleartext non-sensitive tuples are sent in response to a query. We refer this as the adversarial view, denoted by AV : $AV = In_c \cup Op_c$, where In_c refers to the query arrives at the cloud and Op_c refers to the encrypted and non-encrypted tuples, transmitted in response to In_c . For example, the first row of Table II shows an adversarial view that shows that $Op_c = t_2$ tuples from the non-sensitive relation and encrypted $Op_c = t_4$ tuples from the sensitive

relation are returned to answer the query for $In_c = E259$. *Some frequent query values.* The adversary observes query predicates on the non-sensitive data, and hence, can deduce the most frequent query predicates by observing many queries.

Inference Attacks in Partitioned Computations. To see the inference attack on the sensitive data while jointly processing sensitive and non-sensitive data, consider following three queries on the *Employee2* and *Employee3* relations; refer to Figures 2b and 2c.

Example 2. ($Q1$) retrieve tuples of employee E259, ($Q2$) retrieve tuples of employee E101, and ($Q3$) retrieve tuples of employee E199.² When answering a query, the adversary knows the tuple ids of retrieved encrypted tuples and the full information of the returned non-sensitive tuples. We refer to this information gain by the adversary as the *adversarial view*, see Table II, where $E(t_i)$ denotes an encrypted tuple t_i .

Query value	Returned tuples/Adversarial view	
	Employee2	Employee3
E259	$E(t_4)$	t_2
E101	$E(t_1)$	null
E199	null	t_3

Table II: Queries and returned tuples/adversarial view.

Outputs of the above three queries will reveal enough information to learn something about sensitive data. In $Q1$, the adversary learns that E259 works in both sensitive and non-sensitive departments, because the answers obtained from the two relations contribute to the final answer. Moreover, the adversary may learn which sensitive tuple has an *Eid* equals to E259. In $Q2$, the adversary learns that E101 works only in a sensitive department, because the query will not return any answer from the *Employee3* relation. In $Q3$, the adversary learns that E199 works only in a non-sensitive department.

The Query Binning (QB) Approach. To prevent the inference attack in a partitioned computation, a new security definition is needed. Before discussing the formal definition of partitioned data security (§III), we provide a possible solution to prevent inference attacks and then intuition for the security definition.

The query binning (QB) strategy stores a non-sensitive relation, say R_{ns} , in clear-text while it stores a sensitive relation, say R_s , using a cryptographically secure approach. QB prevents leakage such as in Example 2 by appropriately mapping a query for a predicate, say $q(w)$, to corresponding queries both over the non-sensitive relation, say $q(W_{ns})(R_{ns})$, and encrypted relation, say $q(W_s)(R_s)$, which represent a set of predicates (or selection queries) that are executed over the relation R_{ns} in plaintext and over the sensitive relation R_s , using the underlying cryptographic method, respectively. The set of predicates in $q(W_{ns})(R_{ns})$ (likewise in $q(W_s)(R_s)$) correspond to the non-sensitive (sensitive) *bins* including the predicate w , denoted by NSB (SB). The predicates in $q(W_s)(R_s)$ are encrypted before transmitting to the cloud.

²We used random *Eids*, which is common in a real employee relation. In contrast, in sequential ids, the absence of an id from the non-sensitive relation directly informs the adversary that the given id exists in the sensitive relation.

Query value	Returned tuples/Adversarial view	
	Employee2	Employee3
E259	$E(t_4), E(t_1)$	t_2, t_6
E101	$E(t_4), E(t_1)$	t_3, t_8
E199	$E(t_4), E(t_1)$	t_3, t_8

Table III: The adversarial view when following QB.

The bins are selected such that: (i) $w \in q(W_{ns})(R_{ns}) \cap q(W_s)(R_s)$ to ensure that all the tuples containing the predicate w are retrieved, and, (ii) joint execution of the queries $q(W_{ns})(R_{ns})$ and $q(W_s)(R_s)$ (hereafter, denoted by $q(W)(R_s, R_{ns})$, where $W = W_s \cup W_{ns}$) does not leak the predicate w . Results from the execution of the queries $q(W_{ns})(R_{ns})$ and $q(W_s)(R_s)$ are decrypted, possibly filtered, and merged to generate the final answer. Note that *bins are created only once for all the values of a searching attribute before any query is executed*. The details of the bin formation will be discussed in §IV.

For answering the above-mentioned three queries, QB creates two bins on sensitive parts: {E101, E259}, {E152, E159}, and two sets on non-sensitive parts: {E259, E254}, {E199, E152}. Table III illustrates the generated adversarial view when QB is used to answer queries as shown in Example 2. In this example, row 1 of Table III shows that this instance of QB maps the query for E259 to $\langle E259, E254 \rangle$ over cleartext and to encrypted version of values for $\langle E259, E101 \rangle$ over sensitive data. Note that simply from the generated adversarial views, the adversary cannot determine the query value w (E259 in the example) or find a value that is shared between the two sets. Thus, while answering a query, the adversary cannot learn which employee works only in defense, design, or in both. The reason is that the desired query value, w , is encrypted with other encrypted values of W_s , and, furthermore, the query value, w , cannot be distinguished from many requested non-sensitive values of W_{ns} , which are in clear-text. Consequently, *the adversary is unable to find an intersection of the two sets, which is the exact value*.

III. PARTITIONED DATA SECURITY

This section formalizes the notion of *partitioned data security* that establishes when a partitioned computation over sensitive and non-sensitive data does not leak any sensitive information. We begin by first formalizing the concepts of: *associated values*, *associated tuples*, and *relationship between counts of sensitive values*.

Notations used in the definitions. Let t_1, t_2, \dots, t_m be tuples of a sensitive relation, say R_s . Thus, the relation R_s stores the encrypted tuples $E(t_1), E(t_2), \dots, E(t_m)$. Let $s_1, s_2, \dots, s_{m'}$ be values of an attribute, say A , that appears in one of the sensitive tuples of R_s . Note that $m' \leq m$, since several tuples may have an identical value. Furthermore, $s_i \in \text{Domain}(A)$, $i = 1, 2, \dots, m'$, where $\text{Domain}(A)$ represents the domain of values the attribute A can take. By $\#_s(s_i)$, we refer to the number of sensitive tuples that have s_i as the value for attribute A . We further define $\#_s(v) = 0, \forall v \in \text{Domain}(A)$,

$v \notin s_1, s_2, \dots, s_{m'}$. Let t_1, t_2, \dots, t_n be tuples of a non-sensitive relation, say R_{ns} . Let $ns_1, ns_2, \dots, ns_{n'}$ be values of the attribute A that appears in one of the non-sensitive tuples of R_{ns} . In analogy with the case where the relation is sensitive, $n' \leq n$, and $ns_i \in \text{Domain}(A)$, $i = 1, 2, \dots, n'$.

Associated values. Let $e_i = E(t_i)[A]$ be the encrypted representation of an attribute value of A in a sensitive tuple of the relation R_s , and ns_j be a value of the attribute A for some tuple of the relation R_{ns} . We say that e_i is *associated* with ns_j , (denoted by $\stackrel{a}{=}$), if the plaintext value of e_i is identical to the value ns_j . In Example 1, the value of the attribute `Empid` in tuple t_4 (of *Employee2*, see Figure 2b) is associated with the value of the attribute `Empid` in tuple t_2 (of *Employee3*, see Figure 2c), since both values correspond to E259.

Associated tuples. Let t_i be a sensitive tuple of the relation R_s (i.e., R_s stores encrypted representation of t_i) and t_j be a non-sensitive tuple of the relation R_{ns} . We state that t_i is associated with t_j (for an attribute, say A) iff the value of the attribute A in t_i is associated with the value of the attribute A in t_j (i.e., $t_i[A] \stackrel{a}{=} t_j[A]$). Note that this is the same as stating that the two values of attribute A are equal for both tuples.

Relationship between counts of sensitive values. Let v_i and v_j be two distinct values in $\text{Domain}(A)$. We denote the relationship between the counts of sensitive tuples with these A values (i.e., $\#_s(v_i)$ (or $\#_s(v_j)$)) by $v_i \stackrel{r}{\sim} v_j$. Note that $\stackrel{r}{\sim}$ can be one of $<, =, \text{ or } >$ relationships. For instance, in Example 1, the $E101 \stackrel{r}{\sim} E259$ corresponds to $=$, since both values have exactly one sensitive tuple (see Figure 2b), while $E101 \stackrel{r}{\sim} E199$ is $>$, since there is one sensitive tuple with value E101 while there is no sensitive tuple with E199.

Given the above definitions, we can now formally state the security requirement that ensures that simultaneous execution of queries over sensitive (encrypted) and non-sensitive (plaintext) data does not leak any information. Before that, we wish to mention the need of a new security definition in our context. The inference attack in the partitioned computing can be considered to be related to the known-plaintext attack (KPA) wherein the adversary knows some plaintext data which is hidden in a set of ciphertext. In KPA, the adversary's goal is to determine which ciphertext data is related to a given plaintext, i.e., determining a mapping between ciphertext and the corresponding plaintext data representing the same value. In our setup, non-sensitive values are visible to the adversary in plaintext. However, the attacks are different since, unlike the case of KPA, in our setup, the ciphertext data might not contain any data value that is the same as some non-sensitive data visible to the adversary in plaintext.³

Definition: Partitioned Data Security. Let R be a relation containing sensitive and non-sensitive tuples. Let R_s and R_{ns} be the sensitive and non-sensitive relations, respectively. Let AV be an adversarial view generated for a query

³The HBC adversary cannot launch the chosen-plaintext attack (CPA) and the chosen-ciphertext attack (CCA). Since the sensitive data is non-deterministically encrypted (by our assumption), it is not prone to the ciphertext only attack (COA).

$q(w)(R_s, R_{ns})[A]$, where the query, q , for a value w in the attribute A of the R_s and R_{ns} relations. Let X be the auxiliary information about the sensitive data, and Pr_{Adv} be the probability of the adversary knowing any information. A query execution mechanism ensures the partitioned data security if the following two properties hold:

- (1) $Pr_{Adv}[e_i \stackrel{a}{=} ns_j | X] = Pr_{Adv}[e_i \stackrel{a}{=} ns_j | X, AV]$, where $e_i = E(t_i)[A]$ is the encrypted representation for the attribute value A for any tuple t_i of the relation R_s and ns_j is a value for the attribute A for any tuple of the relation R_{ns} .
- (2) $Pr_{Adv}[v_i \stackrel{r}{\sim} v_j | X] = Pr_{Adv}[v_i \stackrel{r}{\sim} v_j | X, AV]$, for all $v_i, v_j \in Domain(A)$.

Equation (1) captures the fact that an initial probability of *associating* a sensitive tuple with a non-sensitive tuple will be identical after executing a query on the relations, *i.e.*, an adversary cannot learn anything from an adversarial view generated after the query execution. Satisfying this condition also prevents an adversary to have success against KPA. Equation (2) states that the probability of an adversary gaining information about the relative frequency of sensitive values does not increase after the query execution. In Example 2, an execution of any three queries (for values E101, E199, or E259) without using QB does not satisfy Equation (1). For example, the query for E199 retrieves the only tuple from non-sensitive relation, and that changes the probability of estimating whether E199 is sensitive or non-sensitive to 0 than an initial probability of the same estimation, which was 1/4. Hence, an execution of the three queries violates partitioned data security. However, the query execution for E259 and E101 satisfies Equation (2), since the count of returned tuples from *Employee2* is equal. Hence, the adversary cannot distinguish between the count of the values (E259 and E101) in the domain of *Eid* of *Employee2* relation.

IV. QUERY BINNING TECHNIQUE

We develop our strategy initially under the assumption that queries are only on a single attribute, say A . QB approach takes as inputs: (i) the set of data values (of the attribute A) that are sensitive, along with their counts, and (ii) the set of data values (of the attribute A) that are non-sensitive, along with their counts. QB returns partitions of attribute values that form the query bins for both the sensitive and for the non-sensitive parts of the query. We begin in §IV-A by developing the approach for the case when a sensitive tuple is associated with at most one non-sensitive tuple (Algorithm 1). Finally, we provide a general strategy to create bins when a sensitive tuple is associated with several non-sensitive tuples, in §IV-B.

Informally, QB distributes attribute values in a matrix, where rows are sensitive bins, and columns are non-sensitive bins. For example, suppose there are 16 values, say $0, 1, \dots, 15$, and assume all the values have sensitive and associated non-sensitive tuples. Now, the DB owner arranges 16 values in a 4×4 matrix, as follows:

	NSB_0	NSB_1	NSB_2	NSB_3
SB_0	11	2	5	14
SB_1	10	3	8	7
SB_2	0	15	6	4
SB_3	13	1	12	9

Here, we have four sensitive bins: $SB_0 \{11, 2, 5, 14\}$, $SB_1 \{10, 3, 8, 7\}$, $SB_2 \{0, 15, 6, 4\}$, $SB_3 \{13, 1, 12, 9\}$, and four non-sensitive bins: $NSB_0 \{11, 10, 0, 13\}$, $NSB_1 \{2, 3, 15, 1\}$, $NSB_2 \{5, 8, 6, 12\}$, $NSB_3 \{14, 7, 4, 9\}$. When a query arrives for a value 1, the DB owner searches for tuples containing values 2, 3, 15, 1 (*viz.* NSB_1) on the non-sensitive data and values in SB_3 (*viz.*, 13, 1, 12, 9) on the sensitive data using a cryptographic mechanism integrated into QB. We will show that in the proposed approach, while the adversary learns that a query corresponds to one of the four values in NSB_1 , since query values in SB_3 are encrypted, the adversary does not learn the actual sensitive value or the actual non-sensitive value that is identical to a clear-text sensitive value.

A. The Base Case

QB consists of two steps. First, query bins are created (information about which will reside at the DB owner) using which queries will be rewritten. The second step consists of rewriting the query based on the binning. Here, QB is explained for the base case, where a sensitive tuple, t_s , is associated with at most a single non-sensitive tuple, t_{ns} , and vice versa (*i.e.*, $\stackrel{a}{=}$ is a 1:1 relationship). Thus, if the value has two tuples, then one of them must be sensitive and the other one must be non-sensitive, but both the tuples cannot be sensitive or non-sensitive. A value can also have only one tuple, either sensitive or non-sensitive. Note that t_1, t_2, \dots, t_l are sensitive tuples, with values of an attribute A being s_1, s_2, \dots, s_n , $s_i \neq s_j$ if $i \neq j$. Thus, in the remainder of the section, we will refer to association between encrypted value $E(t_i)[A]$ and a non-sensitive value ns_j simply as an association between values s_i and ns_j , where s_i is the cleartext representation of $E(t_i)[A]$ and ns_j is a value in the attribute A of a non-sensitive relation; *i.e.*, $s_i \stackrel{a}{=} ns_j$ represents $E(t_i)[A] \stackrel{a}{=} ns_j$.

The scenario in Example 1 satisfies the base case. The *Eid* attribute values corresponding to sensitive tuples include $\langle E101, E259, E152, E159 \rangle$ and corresponding to non-sensitive tuples are $\langle E199, E259, E254, E152 \rangle$ for which $\stackrel{a}{=}$ is 1:1. We discuss QB under the above assumption, but relax the assumption in §IV-B. Before describing QB, we first define the concept of *approximately square factors of a number*.

Approximately square factors. We say two numbers, say x and y , are approximately square factors of a number, say $n > 0$, if $x \times y = n$, and x and y are equal or close to each other such that the difference between x and y is less than the difference between any two factors, say x' and y' , of n such that $x' \times y' = n$.

Step 1: Bin-creation. QB, described in Algorithm 1, finds two approximately square factors of $|NS|$, say x and y , where $x \geq y$. QB creates $SB = x$ sensitive bins, where each sensitive bin contains at most y values. Thus, we assume $|S| \geq x$. QB,

Algorithm 1: Bin-creation algorithm, the base case.

Inputs: $|NS|$: the number of values in the non-sensitive data,
 $|S|$: the number of values in the sensitive data.
Outputs: SB : sensitive bins; NSB : non-sensitive bins

```
1 Function create_bins( $S, NS$ ) begin
2   Permute all sensitive values
3    $x, y \leftarrow \text{approx\_sq\_factors}(|NS|): x \geq y$ 
4    $|NSB| \leftarrow x, NSB \leftarrow \lceil |NS|/x \rceil, SB \leftarrow x, |SB| \leftarrow y$ 
5   for  $i \in (1, |S|)$  do  $SB[i \bmod x][*] \leftarrow S[i]$ ;
6   for  $(i, j) \in (0, SB - 1), (0, |SB| - 1)$  do
    $NSB[j][i] \leftarrow \text{allocateNS}(SB[i][j])$ ;
7   for  $i \in (0, NSB - 1)$  do  $NSB[i][*] \leftarrow$  fill the bin if
   empty with the size limit to  $x$ ;
8   return  $SB$  and  $NSB$ 
9 Function allocateNS( $SB[i][j]$ ) begin
   find a non-sensitive value associated with the  $j^{\text{th}}$ 
   sensitive value of the  $i^{\text{th}}$  sensitive bin
```

further, creates $NSB = \lceil |NS|/x \rceil$ non-sensitive bins, where each non-sensitive bin contains at most $|NSB| = x$ values. Note that we are assuming that $|S| \leq |NS|$. (QB can also handle the case of $|S| > |NS|$ by applying Algorithm 1 in a reverse way, *i.e.*, factorizing $|S|$.)

Assignment of sensitive values. We number the sensitive bins from 0 to $x-1$ and the values therein from 0 to $y-1$. To assign a value to sensitive bins, QB first permutes the set of sensitive values. This permutation is kept secret from the adversary by the DB owner.⁴ To assign sensitive values to sensitive bins, QB takes the i^{th} sensitive value and assigns it to the $(i \bmod x)^{\text{th}}$ sensitive bin (see Lines 2 and 5 of Algorithm 1).

Assignment of non-sensitive values. We number the non-sensitive bins from 0 to $\lceil |NS| \rceil / x - 1$ and values therein from 0 to $x-1$. To assign non-sensitive values, QB takes a sensitive bin, say j , and its i^{th} sensitive value. Assign the non-sensitive value associated with the i^{th} sensitive value to the j^{th} position of the i^{th} non-sensitive bin. Here, if each value of a sensitive bin has an associated non-sensitive value and $|S| = |NS|$, then QB has assigned all the non-sensitive values to their bins (Line 6 of Algorithm 1). Note that it may be the case that only a few sensitive values have their associated non-sensitive values and $|S| \leq |NS|$. In this case, we assign the sensitive and their associated non-sensitive values to bins like we did in the previous case. However, we need to assign the non-sensitive values that are not associated with a sensitive value, by filling all the non-sensitive bins to size x (Line 7 of Algorithm 1).

Aside. Note that QB assigned at least as many values in a non-sensitive bin as it assigned to a sensitive bin. QB may form the non-sensitive and sensitive bins in such a way that the number of values in sensitive bins is higher than the non-sensitive bins.

⁴The DB owner permutes sensitive values to prevent the adversary to create bins at her end; *e.g.*, if the adversary knows that employee ids are ordered, she can also create bins by knowing the number of resultant tuples to a query. For simplicity, we do not show permuted sensitive values in any figure.

Algorithm 2: Bin-retrieval algorithm.

Inputs: w : the query value. **Outputs:** SB_a and NSB_b : one sensitive bin and one non-sensitive bin to be retrieved for answering w .
Variables: $found \leftarrow \text{false}$

```
1 Function retrieve_bins( $q(w)$ ) begin
2   for  $(i, j) \in (0, SB - 1), (0, |SB| - 1)$  do
   if  $w = SB_i[j]$  then
3     return  $SB_i$  and  $NSB_j$ ;  $found \leftarrow \text{true}$ ;
   break
4   if  $found \neq \text{true}$  then
5     for  $(i, j) \in (0, NSB - 1), (0, |NSB| - 1)$  do
6       if  $w = NSB_i[j]$  then
   return  $NSB_i$  and  $SB_j$ ; break
7   Retrieve the desired tuples from the cloud by sending
   encrypted values of the bin  $SB_i$  (or  $SB_j$ ) and
   clear-text values of the bin  $NSB_j$  (or  $NSB_i$ ) to the
   cloud
```

We chose sensitive bins to be smaller since the processing time on encrypted data is expected to be higher than clear-text data processing; hence, by searching and retrieving fewer sensitive tuples, we decrease the encrypted data-processing time.

Step 2: Bin-retrieval – answering queries. Algorithm 2 presents the pseudocode for the bin-retrieval algorithm. The algorithm, first, checks the existence of a query value in sensitive bins and/or non-sensitive bins (see Lines 2 and 4 of Algorithm 2). If the value exists in a sensitive bin and a non-sensitive bin, the DB owner retrieves the corresponding two bins (see Line 7). Note that here the adversarial view is not enough to leak the query value or to find a value that is shared between the two bins. The reason is that the desired query value is encrypted with a set of other encrypted values and, furthermore, the query value is obscured in many requested non-sensitive values, which are in clear-text. Consequently, the adversary is unable to find an intersection of the two bins, which is the exact value.

There are the following three other cases to consider: (i) Some sensitive values of a bin are not associated with any non-sensitive value. For example, in Figure 3, the sensitive values s_4, s_7, s_8, s_9 , and s_{10} are not associated with any non-sensitive value. (ii) A sensitive bin does not hold any value that is associated with any non-sensitive value. For example, the sensitive bin SB_4 in Figure 3 satisfies this clause. (iii) A non-sensitive bin containing no value that is associated with any sensitive value.

In all three cases, if the DB owner retrieves only either a sensitive or non-sensitive bin containing the value, it leads to information leakage similar to Example 2 (or incomplete answers). In order to prevent such leakage, Algorithm 2 follows two rules stated below (see Lines 3 and 6 of Algorithm 2):

Tuple retrieval rule R1. If the query value w is a sensitive value that is at the j^{th} position of the i^{th} sensitive bin (*i.e.*, $w = SB_i[j]$), then the DB owner will fetch the i^{th} sensitive and the j^{th} non-sensitive bins (see Line 3 of Algorithm 2). By

Line 2 of Algorithm 2, the DB owner knows that the value w is either sensitive or non-sensitive.

Tuple retrieval rule R2. If the query value w is a non-sensitive value that is at the j^{th} position of the i^{th} non-sensitive bin, then the DB owner will fetch the i^{th} non-sensitive and the j^{th} sensitive bins (see Line 6 of Algorithm 2).

Note that if query value w is in both sensitive and non-sensitive bins, then both the rules are applicable, and they retrieve *exactly the same* bins. In addition, if the value w is neither in a sensitive or a non-sensitive bin, then there is no need to retrieve any bin.

Aside. After knowing the bins, the DB owner sends all the sensitive values in the encrypted form and the non-sensitive values in clear-text to the cloud. The tuple retrieval based on the encrypted values reveals only the tuple addresses that satisfy the requested values. We can also hide the access-patterns by using PIR, ORAM, or DSSE on each required sensitive value. As mentioned in §I, access-pattern-hiding techniques are prone to size and workload-skew attacks. Nonetheless, the use of QB with access-pattern-hiding techniques makes them secure against these attacks, which is discussed in detail in the full version. QB is designed as a general mechanism that provides partitioned data security when coupled with any cryptographic technique. For special cryptographic techniques that hide access-patterns, it may be possible to design a different mechanism that may provide partitioned data security.

Associated bins. We say a sensitive bin is associated with a non-sensitive bin, if the two bins are retrieved for answering at least one query.

Our aim when answering queries for all the sensitive and non-sensitive values using Algorithm 2 is to associate each sensitive bin with each non-sensitive bin; resulting in the adversary being unable to predict which (if any) is the value shared between two bins.

Example 3: QB example Step 1: Bin Creation. We show the bin-creation algorithm for 10 sensitive values and 10 non-sensitive values. We assume that only five sensitive values, say s_1, s_2, s_3, s_5, s_6 , have their associated non-sensitive values, say $ns_1, ns_2, ns_3, ns_5, ns_6$, and the remaining 5 sensitive (say, $s_4, s_7, s_8, \dots, s_{10}$) and 5 non-sensitive values (say, $ns_{11}, ns_{12}, \dots, ns_{15}$) are not associated. For simplicity, we use different indexes for non-associated values.

QB creates 2 non-sensitive bins and 5 sensitive bins, and divides 10 sensitive values over the following 5 sensitive bins: $SB_0 \{s_5, s_{10}\}$, $SB_1 \{s_1, s_6\}$, $SB_2 \{s_2, s_7\}$, $SB_3 \{s_3, s_8\}$, $SB_4 \{s_4, s_9\}$; see Figure 3. Now, QB distributes non-sensitive values associated with the sensitive values over two non-sensitive bins, resulting in the bin $NSB_0 \{ns_5, ns_1, ns_2, ns_3, *\}$ and $NSB_1 \{*, ns_6, *, *, *\}$, where a $*$ shows an empty position in the bin. In the sequel, QB needs to fill the non-sensitive bins with the remaining 5 non-sensitive values; hence, ns_{11} is assigned to the last position of the bin NSB_0 , and the bin NSB_1 contains the remaining 4 non-sensitive values such as $\{ns_{12}, ns_6, ns_{13}, ns_{14}, ns_{15}\}$.

Example 3: QB example (continued) Step 2: Bin-retrieval.

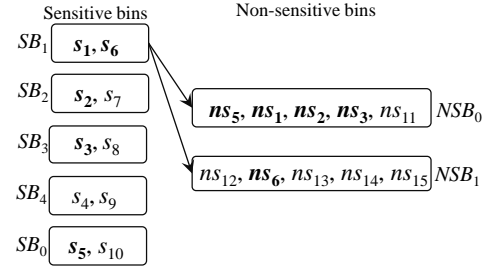


Figure 3: QB for 10 sensitive and 10 non-sensitive values.

We show how to retrieve tuples. If a query is for the sensitive value s_2 (see Figure 3), then the DB owner fetches two bins SB_2 and NSB_0 . If a query is for the non-sensitive value ns_{13} or sensitive value s_7 , then the DB owner fetches two bins SB_2 and NSB_1 . Thus, it is impossible for the adversary to find (by observing the adversarial view) which is an exact query value from the non-sensitive bin and which is the sensitive value associated with one of the non-sensitive values. This fact is also clear from Table IV, which shows that the adversarial view is not enough to leak information from the joint processing of sensitive and non-sensitive data, unlike Example 2. In Table IV, $E(s_i)$ shows the encrypted value of s_i , and we are showing the adversarial view only for queries for s_2, s_7 , and ns_{13} , due to space restriction. In this example, note that the bin SB_2 gets associated with both the non-sensitive bins NSB_0 and NSB_1 , due to following Algorithm 2.

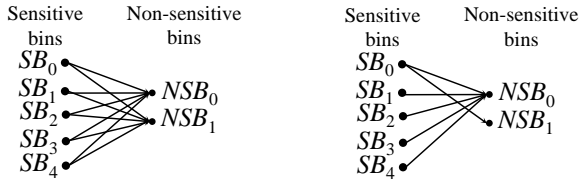
Exact query value	Returned tuples/Adversarial view	
	Sensitive bin and data	Non-sensitive bin and data
s_2 or ns_2	$SB_2: E(s_2), E(s_7)$	$NSB_0: ns_1, ns_2, ns_3, ns_5, ns_{11}$
s_7	$SB_2: E(s_2), E(s_7)$	$NSB_1: ns_6, ns_{12}, ns_{13}, ns_{14}, ns_{15}$
ns_{13}	$SB_2: E(s_2), E(s_7)$	$NSB_1: ns_6, ns_{12}, ns_{13}, ns_{14}, ns_{15}$

Table IV: Queries and returned tuples/adversarial view after retrieving tuples according to Algorithm 2.

Algorithm Correctness. We will prove that QB does not lead to information leakage through the joint processing of sensitive and non-sensitive data. To prove correctness, we first define the concept of *surviving matches*. Informally, we show that QB maintains surviving matches among all sensitive and non-sensitive values, resulting in all sensitive bins being associated with all non-sensitive bins. Thus, an initial condition: a sensitive value is assumed to have an identical value to one of the non-sensitive value is preserved.

Surviving matches. We define surviving matches, which are classified as either *surviving matches of values* or *surviving matches of bins*, as follows:

Before query execution. Before retrieving any tuple, having an assumption that only the DB owner can decrypt an encrypted sensitive value, $E(s_i)$, the adversary cannot learn which non-sensitive value is associated with the value s_i . Thus, the adversary will consider that $E(s_i)$ is associated with one of the non-sensitive values. Based on this fact, the adversary can create a complete bipartite graph having $|S|$ nodes on one side and $|NS|$ nodes on the other side. The edges in the graph are called *surviving matches of the values*. For example, before executing any query, the adversary can create a bipartite graph



(a) Surviving matches after the (b) Surviving matches without fol-
 tuple retrieval following Algo- lowing Algorithm 2 for ns_{12} , ns_{13} ,
 rithm 2. ns_{14} , ns_{15} ; also see Table IV.

Figure 4: An example to show security of QB using surviving matches for 10 sensitive and 10 non-sensitive values.

for 10 sensitive and 10 non-sensitive values.

After query execution. Recall that the query execution on the datasets creates an adversarial view that guides the adversary to create a (new) bipartite graph containing SB nodes on one side and NSB nodes on the other side. The edges in the new graph (obtained after the query execution) are called *surviving matches of the bins*. *E.g.*, after executing queries according to Algorithm 2, an adversary can create a bipartite graph having 5 nodes on one side and 2 nodes on another side (Figure 4a). Note that since bins contain values, the surviving matches of the bins can lead to the surviving matches of the values. Hence, from Figure 4a, the adversary can also create a bipartite graph for 10 sensitive and 10 non-sensitive values.

We show that a technique for retrieving tuples that drops some surviving matches of the bins leading to drop of the surviving matches of the values is not secure, and hence, results in the information leakage through non-sensitive data.

Example 4: Dropping surviving matches. In Figure 3, for answering queries for associated values $s_1, s_2, s_3, s_5, s_6, ns_1, ns_2, ns_3, ns_5$, or ns_6 , the DB owner must follow Line 3 or 6 of Algorithm 2 for retrieving the two bins holding corresponding sensitive and non-sensitive data; otherwise, she cannot retrieve two bins that share a common value. However, for answering values $s_4, s_7, s_8, s_9, s_{10}, ns_6, ns_{12}, ns_{13}, ns_{14}$, or ns_{15} (recall that these values are not associated), if the DB owner does not follow Algorithm 2 and retrieves the bin containing the desired value with any randomly selected bin of the other side, then it could result in the following adversarial view; see Table V. (We show the case when NSB_1 is only associated with bin SB_1 , and bins SB_2 is only associated with bin NSB_0 , since Algorithm 2 is not followed.)

Having such an adversarial view (Table V), the adversary can learn the following fact that: (i) Encrypted sensitive tuples of the bin SB_2 have associated non-sensitive tuples only in the bin NSB_0 , not in NSB_1 (Figure 4b). (ii) Non-sensitive tuples of the bin NSB_1 have their associated sensitive tuples only in the bin SB_1 (see Figure 4b). Based on this adversarial view (Table V), the bipartite graph drops some surviving matches of the bins (see Figure 4b). However, note that in Table IV, bin SB_2 was associated with both non-sensitive bins. Hence, a random retrieval of bins is not secure to prevent information leakage through non-sensitive data accessing.

However, if the DB owner uses Line 3 or 6 of Algorithm 2 for retrieving values that are not associated, the above-

Exact query value	Returned tuples/Adversarial view	
	Sensitive bin and data	Non-sensitive bin and data
s_2 or ns_2	$SB_2; E(s_2), E(s_7)$	$NSB_0; ns_1, ns_2, ns_3, ns_5, ns_{11}$
s_6 or ns_6	$SB_1; E(s_1), E(s_6)$	$NSB_1; ns_6, ns_{12}, ns_{13}, ns_{14}, ns_{15}$
s_7	$SB_2; E(s_2), E(s_7)$	$NSB_0; ns_1, ns_2, ns_3, ns_5, ns_{11}$
ns_{12}	$SB_1; E(s_1), E(s_6)$	$NSB_1; ns_6, ns_{12}, ns_{13}, ns_{14}, ns_{15}$
ns_{13}	$SB_1; E(s_1), E(s_6)$	$NSB_1; ns_6, ns_{12}, ns_{13}, ns_{14}, ns_{15}$
ns_{14}	$SB_1; E(s_1), E(s_6)$	$NSB_1; ns_6, ns_{12}, ns_{13}, ns_{14}, ns_{15}$
ns_{15}	$SB_1; E(s_1), E(s_6)$	$NSB_1; ns_6, ns_{12}, ns_{13}, ns_{14}, ns_{15}$

Table V: Queries and returned tuples/adversarial view without following Algorithm 2.

mentioned facts no longer hold. Figure 4a shows each sensitive bin is associated with each non-sensitive bin, if Algorithm 2 is followed. Thus, all the surviving matches of the bins and values are preserved after answering queries. Hence, for the example of 10 sensitive and 10 non-sensitive values, QB (Algorithms 1 and 2) is secure, so the adversary cannot find an exact association between sensitive and non-sensitive values.

Informal security proof sketch (see [34] for a detailed proof). Let v_1, v_2, v_3 , and v_4 be values containing only one sensitive and one non-sensitive tuple. Let E_1, E_2, E_3 , and E_4 be encrypted representations of these values in an arbitrary order, *i.e.*, it is not mandatory that E_1 be the encrypted representation of v_1 . In this example, the cloud stores an encrypted relation, say R_s , containing four encrypted tuples with encrypted representations E_1, E_2, E_3, E_4 and a clear-text relation, say R_{ns} , containing four clear-text tuples with values v_1, v_2, v_3, v_4 . The objective of the adversary is to deduce a clear-text value corresponding to an encrypted value. Note that before executing a query, the probability of an encrypted value, say E_i , to have the clear-text value, say v_i , $1 \leq i \leq 4$ is $1/4$, which QB maintains at the end of a query. Assume that the user wishes to retrieve the tuple containing v_1 . By following QB, the user asks a query, say $q(E_1, E_3)(R_s)$, on the encrypted relation R_s for E_1, E_3 , and a query, say $q(v_1, v_2)(R_{ns})$, on the clear-text relation R_{ns} for v_1, v_2 . Here, we need that the probability of finding the clear-text value of an encrypted representation, say E_i , $1 \leq i \leq 4$, remains identical before and after a query to satisfy the first condition of the partitioned data security. In short, the retrieval of the four tuples containing one of the following: $\langle E_1, E_3, v_1, v_2 \rangle$, results in 16 possible allocations of the values v_1, v_2, v_3 , and v_4 to E_1, E_2, E_3 , and E_4 , of which only four possible allocations have v_1 as the clear-text representation of E_1 . This results in the probability of finding $E_1 = v_1$ is $1/4$. A similar argument also holds for other encrypted values. Hence, an initial probability of associating a sensitive value with a non-sensitive value remains identical to after executing a query.

A Simple Extension of the Base Case. Algorithm 1 creates bins when the number of non-sensitive data values is not a prime number, by finding the two approximately square factors. However, Algorithm 1 may exhibit a relatively higher *cost* (*i.e.*, the number of the retrieved tuple) when the sum of the approximately square factors is high. For example, if there are 41 sensitive data values and 82 non-sensitive data values, then Algorithm 1 creates 2 non-sensitive bins having 41 values in each and 41 sensitive bins having exactly one value in each. We handle the case when the number of non-sensitive

values is close to a square number. We find the cost using Algorithm 1; in addition, we find a square number closest to the non-sensitive values (here 81 is the closest square number to 82) and the cost. Then, we use Algorithm 1 that creates bins using a method that results in fewer retrieved tuples. In this example, we create 9 non-sensitive and 9 sensitive bins.

B. General Case: Multiple Values with Multiple Tuples

This section generalizes Algorithm 1 to consider a case when different data values have different numbers of associated tuples. First, we will show that sensitive values with different numbers of tuples may provide enough information to the adversary leading to the size, frequency-count attacks, and may disclose some information about the sensitive data. Hence, in the case of multiple values with multiple tuples, Algorithm 1 cannot be directly implemented. We, thus, develop a strategy to overcome such a situation.

Size attack scenario in the base QB. Consider an assignment of 10 sensitive and 10 non-sensitive values to bins using Algorithm 1; see Figure 3. Assume that a sensitive value, say s_1 , has 1000 sensitive tuples and an associated non-sensitive value, say ns_1 , has 2000 tuples, while all the other values have only one tuple each. Further, assume that *each data value represents the salary of employees*. In this example, consider a query execution for a value, say ns_1 . The DB owner retrieves tuples from two bins: SB_1 (containing encrypted tuples of values s_1 and s_6) and NSB_0 (containing tuples of values $ns_1, ns_2, ns_3, ns_5, ns_{11}$); see Figure 3. The number of retrieved tuples satisfying the values of the bins SB_1 and NSB_0 will be highest (*i.e.*, 3005) than the number of tuples retrieved based on any two other bins. Thus, the retrieval of the two bins SB_1 and NSB_0 provides enough information to the adversary to determine which one is the sensitive bin associated with the bin holding the value ns_1 . Moreover, after observing many queries and having background knowledge, the adversary may estimate that 1000 people in the sensitive relation earn a salary equal to the value ns_1 .

Thus, in the case of different sensitive values having different numbers of tuples, Algorithm 1 cannot satisfy the *second condition of partitioned data security* (*i.e.*, the adversary is able to distinguish two sensitive values based on the number of retrieved tuples, which was not possible before the query execution, and concludes that a sensitive value (s_1 in the above example) has more tuples than any other sensitive value) though preserving all surviving matches.

In order for the second condition of partitioned data security to hold (and for the scheme to be resilient to the size and frequency-count attacks, as illustrated above), sensitive bins need to hold identical numbers of tuples. A trivial way of doing this is to outsource some encrypted fake tuples such that the number of tuples in each sensitive bin will be identical. However, we need to be careful; otherwise, adding fake tuples in each sensitive bin may increase the *cost*, if all the heavy-hitter sensitive values are allocated to a single bin. This fact will be clear in the following example.

Example 5: (Illustrating ways to assign sensitive values




Figure 5: An assignment of 9 sensitive values to 3 bins.

to bins to minimize the addition of fake tuples). Consider 9 sensitive values, say s_1, s_2, \dots, s_9 , having 10, 20, 30, 40, 50, 60, 70, 80, and 90 tuples, respectively (we assume that there are 9 non-sensitive values, and computed that we need 3 sensitive and 3 non-sensitive bins). There are multiple ways of assigning these values to three bins so that we need to add a minimum number of fake tuples to each bin. Figure 5 shows two different ways to assign these values to bins. Figure 5b shows the best way – to minimize the addition of fake encrypted tuples; hence minimizing the cost. However, bins in Figure 5a require us to add 180 and 90 fake encrypted tuples to the bins SB_0 and SB_1 , respectively.

There is no need to add any fake tuple if the non-sensitive values have identical numbers of tuples, because an adversary cannot deduce which sensitive bin contains sensitive tuples associated with a non-sensitive value. However, it is obvious that fake non-sensitive tuples cannot be added in clear-text.

Adding fake encrypted tuples. As an assumption, we know the number of sensitive bins, say SB , using Algorithm 1. Our objective is to assign sensitive values to bins such that each bin holds identical numbers of tuples while minimizing the number of fake tuples in each bin. To do this, the strategy is as follows: (i) Sort all the values in a decreasing order of the number of tuples. (ii) Select SB largest values and allocate one in each bin. (iii) Select the next value and find a bin that is containing the fewest number of tuples. If the bin is holding less than y values, then add the value to the bin; otherwise, select another bin with the fewest number of tuples. Repeat this step, for allocating all the values to sensitive bins. (iv) Add fake tuples’ values to the bins so that each bin contains identical numbers of tuples. (v) Allocate non-sensitive values as per Algorithm 1 (Lines 6 and 7).

V. PERFORMANCE EVALUATION OF QB

This section explores how effective is QB in scaling expensive cryptographic techniques by eliminating the necessity of encrypted data processing over non-sensitive data. Note that while QB prevents expensive cryptographic operations, it, nonetheless, comes with an overhead of additional search as well as communication costs. Thus, a natural question is when the tradeoff offered by QB improves performance. We note that if a cryptographic mechanism is extremely inexpensive, *e.g.*, deterministic, order-preserving encryptions, or index-based mechanisms [9], [10], the overhead of QB would not justify the reduced encrypted data processing. But QB has not been designed for such techniques. Instead, our goal for QB is to couple it with techniques such as homomorphic encryptions, DPF [6], or secret-sharing [5] that offer strong security but do

not scale. Below we develop an analytical model to compare performance of cryptographic mechanisms with/without QB. We then conduct an experimental validation of the model and study QB under different choices of parameters.

A. Performance Modeling of QB

For our model, we will need the following notations: (i) C_{com} : Communication cost of moving one tuple over the network. (ii) C_p (or C_e): Processing cost of a single selection query on plaintext (or encrypted data). In addition, we define three parameters: (i) α : is the ratio between the sizes of the sensitive data (denoted by S) and the entire dataset (denoted by $S + NS$, where NS is non-sensitive data). (ii) β : is the ratio between the predicate search time on encrypted data using a cryptographic technique and on clear-text data. The parameter β captures the overhead of a cryptographic technique. Note that $\beta = C_e/C_p$. (iii) γ : is the ratio between the processing time of a single selection query on encrypted data and the time to transmit the single tuple over the network from the cloud to the DB owner. Note that $\gamma = C_e/C_{com}$.

Based on the above parameters, we can compute the cost of cryptographic and non-cryptographic selection operations as follows: (i) $Cost_{plain}(x, D)$: is the sum the processing cost of x selection queries on plaintext data and the communication cost of moving all the tuples having x predicates from the cloud to the DB owner, *i.e.*, $x(\log(D)P_p + \rho DC_{com})$. (ii) $Cost_{crypt}(x, D)$: is the sum the processing cost of x selection queries on encrypted data and the communication cost of moving all the tuples having x predicates from the cloud to the DB owner, *i.e.*, $P_e D + \rho x DC_{com}$, where ρ is the selectivity of the query. Note that cost of evaluating x queries over encrypted data using techniques such as [2], [5], [6], is amortized and can be performed using a single scan of data. Hence, x is not the factor in the cost corresponding to encrypted data processing.

Given the above, we define a parameter η that is the ratio between the computation and communication cost of searching using QB and the computation and communication cost of searching when the entire data (*viz.* sensitive and non-sensitive data) is fully encrypted using the cryptographic mechanism.

$$\eta = \frac{Cost_{crypt}(|SB|, S)}{Cost_{crypt}(1, D)} + \frac{Cost_{plain}(|NSB|, NS)}{Cost_{crypt}(1, D)}$$

Filling out the values from above, the ratio is:

$$\eta = \frac{C_e S + |SB| \rho DC_{com}}{C_e D + \rho DC_{com}} + \frac{|NSB| \log(D) C_p + |NSB| \rho DC_{com}}{C_e D + \rho DC_{com}}$$

Separating out the communication and processing costs,

$$\eta = \frac{S}{D} \frac{C_e}{C_e + \rho C_{com}} + \frac{|NSB| \log(D) C_p}{C_e D + \rho DC_{com}} + \frac{\rho DC_{com} (|NSB| + |SB|)}{C_e D + \rho DC_{com}}$$

Substituting for various terms and cancelling common terms:

$$\eta = \alpha \frac{1}{(1 + \frac{\rho}{\gamma})} + \frac{\log(D)}{D} \frac{|NSB|}{\beta(1 + \frac{\rho}{\gamma})} + \frac{\rho}{\gamma} \frac{|NSB| + |SB|}{(1 + \frac{\rho}{\gamma})}$$

Note that ρ/γ is very small, thus the term $(1 + \rho/\gamma)$ can be substituted by 1. Given the above, the equation becomes: $\eta = \alpha + \log(D)|NSB|/D\beta + \rho(|NSB| + |SB|)/\gamma$. Note that the term $\log(D)|NSB|/D\beta$ is very small since $|NSB|$ is the





Figure 6: Experiments.

number of distinct values (approx. equal to $\sqrt{|NS|}$) in a non-sensitive bin, while D , which is the size of a database, is a large number, and β value is also very large. Thus, the equation becomes: $\eta = \alpha + \rho(|SB| + |NSB|)/\gamma$.

QB is better than a cryptographic approach when $\eta < 1$, *i.e.*, $\alpha + \rho(|SB| + |NSB|)/\gamma < 1$. Thus, $\alpha < 1 - \frac{\rho(|SB| + |NSB|)}{\gamma}$. Note that the values of $|SB|$ and $|NSB|$ are $\approx \sqrt{|NS|}$, we can simplify the above equation to: $\alpha < 1 - 2\rho\sqrt{|NS|}/\gamma$. If we estimate ρ to be roughly $1/|NS|$ (*i.e.*, we assume uniform distribution), the above equation becomes: $\alpha < 1 - 2/\gamma\sqrt{|NS|}$.

The equation above demonstrates that QB trades increased communication costs to reduce the amount of data that needs to be searched in encrypted form. Note that the reduction in encryption cost is proportional to α times the size of the database, while the increase in communication costs is proportional to $\sqrt{|D|}$, where $|D|$ is the number of distinct attribute values. This, coupled with the fact that γ is much higher than 1 for encryption mechanisms that offer strong security guarantees, ensures that QB almost always outperforms the full encryption approaches. For instance, the cryptographic cost for search using secret-sharing is $\approx 10ms$ [5], while the cost of transmitting a single row (≈ 200 bytes for TPC Customer table) is $\approx 4 \mu s$ making the value of $\gamma \approx 25000$. Thus, QB, based on the model, should outperform the fully encrypted solution for almost any value of α , under ideal situations where our assumption of uniformity holds. Figure 6a plots a graph of η as a function of γ , for varying sensitivity and $\rho = 10\%$.

B. Experimental Validation

We determined η values for two commercial databases that support non-deterministic encryption. We refer to them as systems A and B, respectively, to hide their identities.

Search Techniques. To support encrypted search on both systems, since they do not provide a searching facility on the non-deterministically encrypted data, we implemented the

Technique	1%	5%	20%	40%	60%
SGX-based Opaque [16]	11	15	26	42	59
MPC-based Jana [37]	22	80	270	505	749

Table VI: Time (in seconds) when mixing QB with Opaque and Jana at different levels of sensitivity.

following technique: retrieves the searching attribute of a sensitive relation at the DB owner side, decrypts the attributes, and searches for records that match $|SB|$. It then retrieves full tuples corresponding to $|SB|$ predicates’ addresses. For comparing against cryptographic searches at the cloud-side, we used SGX-based Opaque [16] and the multi-party computations (MPC) based Jana [37] for evaluating QB’s effectiveness.

Experimental setup. We used a virtual machine of 2.6 GHz, 4 core processor, 16 GB RAM, 1TB disk, and average 30Mbps download speed. We used TPC-H benchmark to generate the dataset. The DB owner stores sensitive and non-sensitive bins, whose size was propositional to the domain size of the searchable attributes and independent of the database size. For TPC-H LINEITEM table, metadata for attributes L_PARTKEY and L_SUPPKEY were 13.6MB and 0.65MB, respectively.

Exp 1: Robustness of QB. To explore the effectiveness of QB under different DB sizes, we tested QB for 3 DB sizes: 150K, 1.5M, and 4.5M tuples using No-Ind(A) and No-Ind(B) as underlying cryptographic mechanisms. Figure 6b plots η values for the three sizes for No-Ind(A) while varying α . The figure shows that $\eta < 1$, irrespective of the DB sizes, confirming that QB scales to larger DB sizes (results over No-Ind(B) are similar). Table VI shows the time taken when using QB with Opaque and Jana at different levels of sensitivity. Without using QB for answering a simple selection query, Opaque [16] took 89 seconds on a dataset of size 700MB (6M tuples) and Jana [37] took 1051 seconds on a dataset of size 116MB (1M tuples). Note that the time to execute the same query on cleartext data of size 700MB took only 0.0002 seconds. QB improves not only the performance of Opaque and Jana, but also makes them to work securely on partitioned data and resilient to output-size attack. The performance of QB will be even higher when one uses more secure cryptographic techniques that are resilient to output-size attacks, since these techniques will consume significant time for answering a query.

Exp 2: Effect of bin size. Figure 6c plots an average time for a selection query using QB with a different bin size, which is in turn governed by the values of $|SB|$ and $|NSB|$, respectively. We plot the effect of $||SB| - |NSB||$ on retrieval time and find that the minimum time is achieved when $|SB| = |NSB|$. Thus, the optimal choice is $|SB| = |NSB| = \sqrt{|NS|}$.

VI. DESIDERATA

Below we focus on an aspect of QB, which is to a degree surprising. While QB is designed for scaling cryptographic techniques, it has a side-effect of improving security properties of an underlying cryptographic technique. In particular, a cryptographic technique that is prone to output-size, frequency-count, and workload-skew attacks becomes secure against these attacks when mixed with QB. Thus, QB offers a higher

level of security, in addition to saving the cryptographic search on non-sensitive data.

Enhancing security-levels of indexable techniques. We show how QB can be integrated with an indexable cryptographic technique, namely Arx [9] that uses a non-deterministic encryption mechanism. In Arx, the DB owner stores each domain value v and the frequency of v in the database. The technique encrypts the i^{th} occurrence of v as a concatenated string $\langle v, i \rangle$ thereby ensuring that no two occurrences of v result in an identical ciphertext. Such a ciphertext representation can then be indexed on the cloud-side. During retrieval, the user keeps track of the histogram of occurrences for each value and generates appropriate ciphertexts that can be used to query the index on the cloud. It is not difficult to see that Arx, by itself, is susceptible to the size, frequency-count, workload-skew, and access-pattern attacks. The query processing using Arx as efficient as cleartext version due to using an index, *e.g.*, β values for Arx are 1.4 on system A and 2.5 on system B.

The use of QB with Arx makes it secure against output-size, frequency-count, and workload-skew attacks. Of course, QB takes more time as compared to Arx, since the time of $|SB|$ searches cannot be absorbed in a single index scan unless all $|SB|$ values lie in a single node of the index. In the worst case, we traverse the index at most $|SB|$ times, unlike Arx, which traverses the index only once for a single selection query. It, however, significantly enhances the security of Arx by preventing output size, frequency count, and workload-skew attacks. However, QB does not protect access-patterns being revealed which could be prevented using ORAM. Determining whether coupling ORAM with Arx mixed with QB or using a more secure cryptographic solution, *e.g.*, secret-sharing, which uses a linear scan to prevent access-patterns, with QB, more efficient (while QB with both the solutions strengthen the underlying cryptographic technique) is an open question.

VII. CONCLUSION

We propose query binning (QB) technique that serves as a meta approach on top of existing cryptographic techniques to support secure selection queries when a relation is partitioned into cryptographically secure sensitive and clear-text non-sensitive sub-relations. Further, we develop a new notion of partitioned data security that restricts exposing sensitive information due to the joint processing of the sensitive and non-sensitive relations. Besides improving efficiency, while supporting partitioned security, interestingly, QB enhances the security of the underlying cryptographic technique by preventing size, frequency-count, and workload-skew attacks. Thus, combining QB with efficient but non-secure cloud-side indexable cryptographic approaches result in an efficient and significantly more secure search. Furthermore, existing access-pattern-hiding cryptographic techniques also benefit from the added security that QB offers.

REFERENCES

- [1] C. Gentry, *A fully homomorphic encryption scheme*. PhD thesis, 2009.
- [2] D. X. Song *et al.*, “Practical techniques for searches on encrypted data,” in *IEEE SP*, pp. 44–55, 2000.

-
- [3] O. Goldreich, "Towards a theory of software protection and simulation by oblivious rams," in *STOC*, pp. 182–194, 1987.
 - [4] A. Shamir, "How to share a secret," *Commun. ACM*, vol. 22, no. 11, pp. 612–613, 1979.
 - [5] F. Emekçi *et al.*, "Dividing secrets to secure data outsourcing," *Inf. Sci.*, vol. 263, pp. 198–210, 2014.
 - [6] N. Gilboa *et al.*, "Distributed point functions and their applications," in *EUROCRYPT*, pp. 640–658, 2014.
 - [7] E. Boyle *et al.*, "Function secret sharing," in *EUROCRYPT*, 2015.
 - [8] R. A. Popa *et al.*, "CryptDB: processing queries on an encrypted database," *Commun. ACM*, vol. 55, no. 9, pp. 103–111, 2012.
 - [9] R. Poddar *et al.*, "Arx: A strongly encrypted database system," *IACR Cryptology ePrint Archive*, 2016.
 - [10] E. Shmueli *et al.*, "Designing secure indexes for encrypted databases," in *DBSec*, pp. 54–68, 2005.
 - [11] M. Naveed *et al.*, "Inference attacks on property-preserving encrypted databases," in *SIGSAC*, pp. 644–655, 2015.
 - [12] G. Kellaris *et al.*, "Generic attacks on secure outsourced databases," in *CCS*, pp. 1329–1340, 2016.
 - [13] V. Costan and S. Devadas, "Intel SGX explained," *IACR Cryptology ePrint Archive*, vol. 2016, p. 86, 2016.
 - [14] T. T. A. Dinh *et al.*, "M2R: enabling stronger privacy in mapreduce computation," in *USENIX*, pp. 447–462, 2015.
 - [15] F. Schuster *et al.*, "VC3: trustworthy data analytics in the cloud using SGX," in *SP*, pp. 38–54, 2015.
 - [16] W. Zheng *et al.*, "Opaque: An oblivious and encrypted distributed analytics platform," in *NSDI*, pp. 283–298, 2017.
 - [17] C. Priebe *et al.*, "EnclaveDB: A secure database using SGX," in *SP*, pp. 264–278, 2018.
 - [18] W. Wang *et al.*, "Leaky cauldron on the dark land: Understanding memory side-channel hazards in SGX," in *CCS*, pp. 2421–2434, 2017.
 - [19] J. Götzfried *et al.*, "Cache attacks on Intel SGX," in *EUROSEC*, pp. 2:1–2:6, 2017.
 - [20] M. Shih *et al.*, "T-SGX: eradicating controlled-channel attacks against enclave programs," in *NDSS*, 2017.
 - [21] V. Costan *et al.*, "Sanctum: Minimal hardware extensions for strong software isolation," in *USENIX Security*, pp. 857–874, 2016.
 - [22] K. Zhang *et al.*, "Sedic: privacy-aware data intensive computing on hybrid clouds," in *CCS*, pp. 515–526, 2011.
 - [23] K. Y. Oktay *et al.*, "SEMROD: secure and efficient MapReduce over hybrid clouds," in *SIGMOD*, pp. 153–166, 2015.
 - [24] K. Y. Oktay *et al.*, "Secure and efficient query processing over hybrid clouds," in *ICDE*, pp. 733–744, 2017.
 - [25] <http://www.computerworld.com/article/2834193/cloud-computing/5-tips-for-building-a-successful-hybrid-cloud.html>.
 - [26] <https://www.getfilecloud.com/blog/2015/07/5-tips-on-optimizing-your-hybrid-cloud/>.
 - [27] C. Farkas and S. Jajodia, "The inference problem: A survey," *SIGKDD Explorations*, vol. 4, no. 2, pp. 6–11, 2002.
 - [28] T. H. Hinke, "Inference aggregation detection in database management systems," in *SP*, pp. 96–106, 1988.
 - [29] G. W. Smith, "Modeling security-relevant data semantics," in *SP*, pp. 384–391, 1990.
 - [30] G. Lee *et al.*, "Hiding sensitive patterns in association rules mining," in *COMPSAC*, pp. 424–429, 2004.
 - [31] C. Li *et al.*, "Protecting individual information against inference attacks in data publishing," in *DASFAA*, pp. 422–433, 2007.
 - [32] H. Pang and X. Ding, "Privacy-preserving ad-hoc equi-join on outsourced data," *TODS*, vol. 39, no. 3, pp. 23:1–23:40, 2014.
 - [33] P. Martins *et al.*, "A survey on fully homomorphic encryption: An engineering perspective," *CSUR*, vol. 50, no. 6, pp. 83:1–83:33, 2017.
 - [34] S. Mehrotra *et al.*, "Partitioned data security on outsourced sensitive and non-sensitive data," tech. rep. <http://isg.ics.uci.edu/publications.html>.
 - [35] Stealth SDB, available at: <http://www.stealthsoftwareinc.com/>.
 - [36] R. Canetti *et al.*, "Adaptively secure multi-party computation," in *STOC*, pp. 639–648, 1996.
 - [37] D. W. Archer *et al.*, "From keys to databases—real-world applications of secure multi-party computation," eprint.iacr.org/2018/450.pdf.