

CSIS0801 - Final Year Project

SecureDB - A Secure Query Processing System with data interoperability

Supervisor: Professor Benjamin Kao

Group Member: Lin Haibin

Table of Contents

1. Introduction	3
2. Project Background and Literature Review	4
3. Project Methodology	6
3.1 Theoretical Background	6
3.1.1 Encryption Procedure	6
3.1.2 Secure Operators	7
3.2 System Architecture	10
3.2.1 Overall Architecture	10
3.2.2 System Components	11
3.3 Query Rewriting	14
3.3.1 Query Rewriting for Multiplication	14
3.3.2 Query Rewriting for Addition	15
3.3.3 Query Rewriting for Subtraction	15
3.3.4 Query Rewriting for Comparison	16
3.4 Web Interface	18
3.4.1 Web Interface for Data Upload	18
3.4.1 Web Interface for Query	19
4. Experiments and Results	21
4.1 Experiment Environment	21
4.2 Experiment Result	22
5. Conclusion and Future Works	25
5.1 Secure Operator Extension	25
5.2 Cryptographic Optimization	25
5.3 Conclusion	26
6. References	27
7. Appendix	28
7.1 Rewrite Rules for Other Secure Operators	28
7.2 Experiment Data	

1. Introduction

As Database-as-a-service(DBaaS) such as Amazon Web Service and Microsoft Azure gets more economically available, security issue in the public cloud becomes one of the major concerns of data owners(DO). Despite the advantages of computation elasticity and scalability provided by the cloud service provider(SP), data owner cannot afford the risk of having highly sensitive data compromised on the cloud, such as clients' credit card information, as the public cloud may be vulnerable to privileged database administrators and hackers who may gain accesses to disk-resident data, or learn from the query results requested by DO.

SDB[1] is a secure query processing system that supports *data interoperability* with secret sharing encryption scheme, where the output of an operator can be taken as input of another operator, which leads to a wide range of SQL queries executable by cloud database, without revealing sensitive information. Besides security, we also demonstrate that SDB is practically efficient.

In this project, we implemented prototype of SDB by query rewriting, on top of Apache Spark.

Specifically, this version of SDB supports the following kind of data interoperable query operations:

- Encryption for integer type data
- Secure operators: addition, subtraction, multiplication, comparison
- Secure aggregation function: count
- Web interface for table creation, data uploading and secure querying

2. Project Background and Literature Review

There exist many approaches tackling the challenge of data security. One common way is to encrypt data before uploading to the cloud server. For instance, Oracle database 11g provides Transparent Data Encryption to encrypt disk-resident data, as a service of reliable data storage and administration. However, as the data is not preserved after conventional encryption, query processing on encrypted data is impossible and the computational power of cloud service provider is thus lost. To execute queries, DO has to adopt a *Decrypt-Before-Query(DBQ)* approach, that is, shipping all ciphertext to DO, decrypting them before executing queries.

On the other hand, *fully homomorphic encryption(FHE)* techniques are developed to realize computation on encrypted data. Despite its theoretical significance, such encryption scheme is too computationally expensive for practical use, not to mention data-intensive queries on the cloud. For example, a recent implementation of FHE[2] lets us process 180 AES blocks in around 18 minutes using 3.7GB RAM, whose running time is far from satisfactory for industry use.

Several partially homomorphic encryption schemes are proposed to process particular types of operations on encrypted data, such as OPES[3](which supports homomorphic comparison encryption) and RSA(which supports homomorphic multiplication encryption). However, a simple selection query such as “`uni-cost * quantity < budget`” cannot be processed on encrypted data by piecing different encryption functions together.

TrustedDB[4] and Cipherbase are hardware-based secure processors, which use secure hardware components to store private keys of sensitive data during processing. However, since security is

dependent on the specific hardware, such secure processors cannot scale up as clusters easily and have to be replaced in case of hardware failure, which makes them not economically scalable.

CryptDB[5] is a well-known system for processing encrypted data, using an onion encryption approach, implemented upon MySQL. CryptDB uses different levels of encryptions to support different security strength and computational support. For example, it uses RSA to allow equality check on encrypted data, and uses OPES to achieve ordering preserving comparison. Under CryptDB, data security is gradually released to a level to achieve appropriate computation requirement. However, some operators of CryptDB are not secure against chosen plaintext attack, and it's capable of a limited range of secure operations.

SDB is a secure query processing system for relational database, using secret sharing encryption schema. With the data interoperable secure operators supported by secret sharing encryption, SDB simulates fully homomorphic encryption for sensitive data stored in database, which leads to a wide range of queries supported by cloud database server. A typical query “SELECT name, salary + 1000 FROM employee WHERE salary * 12 > 100000 AND age < 25”(with salary being sensitive), which requires additive, multiplicative and order-preserving homomorphism, is supported by SDB. Such query will be executed at cloud server first, afterwards data owner receives the encrypted query result (decrypted values of salary + 1000 and plaintext of name) and decrypts it to get the original result. This way, SecureDB turns out to be both secure and efficient.

3. Project Methodology

3.1 Theoretical background

3.1.1 Encryption Procedure

SDB encrypts sensitive data with a secret sharing method per column. The DO maintains the secret keys of columns while SP maintains the encrypted value of sensitive columns, plaintext of non-sensitive columns together with some crypto helper columns.

To encrypt data with secret sharing, the DO first generate two secret numbers, g and n . n is the product of two big random prime numbers ρ_1, ρ_2 , while g is a positive number that is co-prime with n . We define

$$n = \rho_1 \rho_2$$
$$\phi(n) = (\rho_1 - 1)(\rho_2 - 1).$$

Consider a sensitive column A to be encrypted. DO generates a pair of *column keys* $\langle m, x \rangle$ for column A randomly. Also, for each row of column A , DO generates a distinct *row-id* randomly. We require that $0 < r, m, x < n$.

We use $\llbracket v \rrbracket$ to denote a sensitive value to be encrypted, v_e to denote the ciphertext after encryption, v_k to denote the item key of $\llbracket v \rrbracket$.

The *item key* is generated by

$$v_k = \text{gen}(r, \langle m, x \rangle) = mg^{(rx \bmod \phi(n))} \bmod n.$$

The encrypted value is computed by

$$v_e = \mathcal{E}(\llbracket v \rrbracket, v_k) = \llbracket v \rrbracket v_k^{-1} \bmod n$$

To recover plaintext from ciphertext, we compute

$$\llbracket v \rrbracket = \mathcal{D}(v_e, v_k) = v_e v_k \bmod n.$$

Row-id is encrypted by additive homomorphic encryption $E()$.

The encryption procedure is illustrated in Figure 2. It shows how sensitive data is transformed into encrypted values. To reveal the plaintext, DO only needs to store the column keys, while SP maintains the bulk encrypted data.

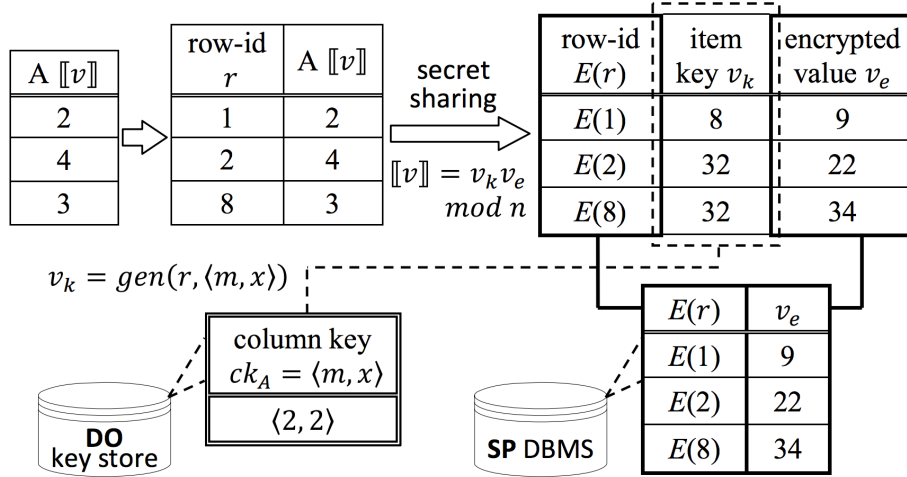


Figure 2: Encryption procedure ($g = 2, n = 35$).

3.1.1 Secure Operators

Data interoperability is one of the most important qualities of SDB. Table 1 shows a list of primitive secure operators implemented in current version of SDB.

Operator	expression	description
\times	$A \times B$	vector dot product of two columns of the same table
$+, -$	$A + B, A - B$	vector addition/subtraction of two columns of the same table
$=$	$A = B$	equality comparison on two columns of the same table and output a binary column of '0' and '1'
$>$	$A > B$	ordering comparison on two columns of the same table and output a binary column of '0' and '1'
π	$\pi S(R)$	project table R on attributes specified in an attribute set $S \otimes$
Count	Count(R)	count the number of rows in a relation

Table 1: list of primitive secure operators

Notice that SDB supports data operability in different modes. For instance, the \times operator, can be applied in three scenarios: operands are both encrypted(*EE Mode*); one operand is encrypted, the other is a constant(*EC Mode*); one operand is encrypted, the other is plain(*EP Mode*). Besides row-id, column R(with encrypted values of random number) and column S(with encrypted values of 1's) help achieve secure operation.

Data: Column A, B with column key $\langle m_A, x_A \rangle$ and $\langle m_B, x_B \rangle$

Result: $C = AB$ with C 's column key $\langle m_C, x_C \rangle$

Client-protocol:

$x_C = x_A + x_B \bmod \phi(n)$;

$m_C = m_A m_B \bmod n$;

Server-protocol:

for each row r do

 Let a_e, b_e be the encrypted values on A, B ;

 Set encrypted value of C $c_e = a_e b_e \bmod n$;

end

Algorithm 1: EE multiplication




Figure 4: $C = A \times B$ ($g = 2, n = 35$).

The protocol for secure multiplication in EE mode, and an illustrative figure are listed above. To calculate the secure multiplication between two sensitive columns, DO simply generates a new column key while SP performs expensive computation on bulk data. The encrypted result is then sent back to DO and decrypted by the new column key. The proof for correctness and security can be found in [1].

Key update is a helper operator which takes an column A and a target column key $\langle m_C, x_C \rangle$ and output a column C that shares the same plaintext of A with target column key, without revealing sensitive information. Other protocols of secure operators implemented in SDB are also listed below.

Data: Column A with column key $\langle m_A, x_A \rangle$ and a constant p

Result: $C = pA$ with C 's column key $\langle m_C, x_C \rangle$

Client-protocol:

$x_C = x_A$;

$m_C = pm_A \bmod n$;

Indicate that C 's encrypted column is A 's encrypted column;

Server-protocol:

Nil

Algorithm 2: EC multiplication

Data: Column A, B with column key $\langle m_A, x_A \rangle$ and $\langle m_B, x_B \rangle$

Result: $C = A + B$ with C 's column key $ck_C = \langle m_C, x_C \rangle$

Client-protocol:

Generate random m_C, x_C

$A' = \kappa(A, ck_C)$; // DO executes client-protocol

$B' = \kappa(B, ck_C)$; // of key update.

Set C 's column key as $\langle m_C, x_C \rangle$;

Server-protocol:

$A' = \kappa(A, ck_C)$; // SP executes server-protocol

$B' = \kappa(B, ck_C)$; // of key update.

for each row r do

 Let a'_e, b'_e be the encrypted values on A', B' ;

 Set encrypted value of C $c_e = a'_e + b'_e \bmod n$;

end

Algorithm 4: EE addition/subtraction

Data: (i) Column A with column key $\langle m_A, x_A \rangle$; (ii) target column key $\langle m_C, x_C \rangle$

Result: $C = A$ with C 's column key $\langle m_C, x_C \rangle$

Client-protocol:

Let $\langle m_S, x_S \rangle$ be the column key of S ;

$p = x_S^{-1}(x_C - x_A) \bmod \phi(n)$;

$q = m_A m_S^p m_C^{-1} \bmod n$;

Send p, q to SP;

Set C 's column key as $\langle m_C, x_C \rangle$;

Server-protocol:

Obtain p, q from DO;

for each row r do

 Let a_e, s_e be the encrypted values on A, S ;

 Set encrypted value of C $c_e = qa_e s_e^p \bmod n$;

end

Algorithm 3: Key update

Data: Column A, B with column key $\langle m_A, x_A \rangle$ and $\langle m_B, x_B \rangle$

Result: A column of comparison results C : = 0 if $A = B$; = 1 if $A > B$; return = -1 if $A < B$

Client-protocol:

$Z = R(A - B)$; // EE addition, EE multiplication

$Z' = \kappa(Z, \langle 1, 0 \rangle)$; // DO executes client-protocol

Server-protocol:

$Z = R(A - B)$; // Corresponding server protocol

$Z' = \kappa(Z, \langle 1, 0 \rangle)$; // Corresponding server protocol

for each row r do

 Let z'_e be the values on Z' ;

switch z'_e **do**

case = 0

 | $c_e = 0$; // c_e is the result of this row

end

case > 0

 | $c_e = 1$;

end

case < 0

 | $c_e = -1$;

end

endsw

end

Algorithm 5: Comparison

3.2 System Architecture

3.2.1 Overall Architecture

SDB is implemented as a software-layer on top of *Apache Spark*. *Apache Spark* is a fast and general engine for large-scale data processing, with in-memory primitives that enhance fast cluster computing. *Apache Hadoop* serves as the cluster manager and distributed storage system for very large data sets to support *Spark*'s cluster computation. *Apache Hive* is the distributed data warehouse to store data.

Spark SQL is a component on top of *Spark Core* that provides support for structured data and SQL support. *Spark SQL* supports *User-Defined Functions*, which provides a mechanism for extending the functionality of the cluster database server by adding cryptographic function that can be evaluated in *Spark SQL*. This way, SDB's server protocol is pushed to *Spark* execution engine and executed as UDFs.

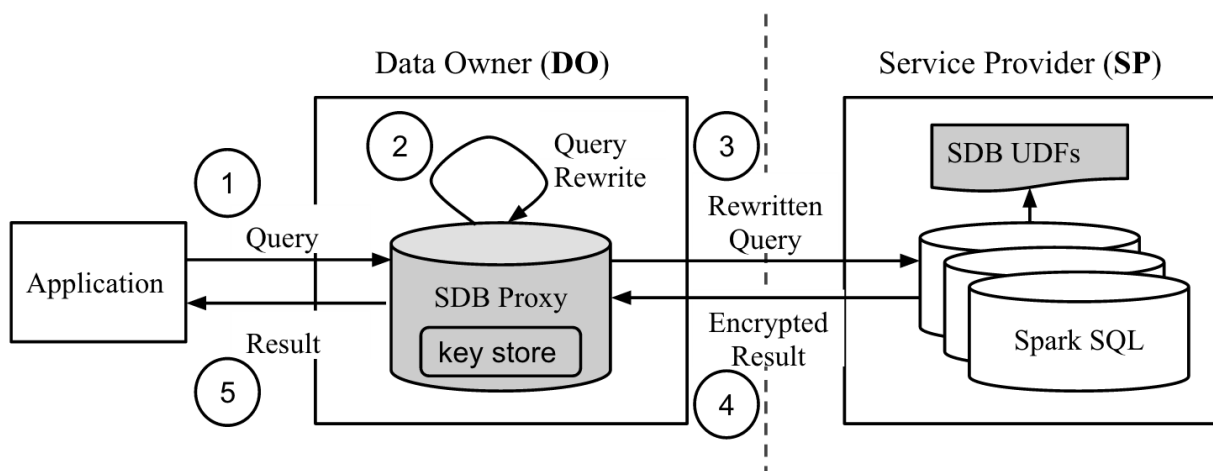


Figure 5: Architecture of SDB

The SDB Proxy resides on client side and stores the secret keys in the *key store*. A typical life cycle of a query consists of 5 stages:

1. Application submits a query to SDB proxy.
2. SDB proxy parses, analyses and rewrites the query with SDB UDFs.

3. SDB proxy submits the rewritten query to Spark SQL for server protocol execution.
4. Spark SQL sends back the encrypted query result to SDB proxy.
5. SDB proxy decrypts the query result and sends it back to application.

The advantages of SDB's architecture include

1. Performance wise
 - Secure operators are implemented as UDF which are executed in the same address space of the SparkSQL, which leads to less memory copy, less network transfer and no IPC.
2. Engineering wise
 - SDB leverages the SparkSQL's optimizer to optimize the server side queries.
 - SDB leverages all the normal operators provided by SparkSQL.
 - Machine failures of the working nodes, disk-based processing and parallelism are well taken care of by Spark.

3.2.2 System Components

SDB proxy is composed of 4 major components:

1. SDB *Query Parser*
2. SDB *Semantic Analyser*
3. SDB *Query Rewriter*
4. SDB *Query Executor*

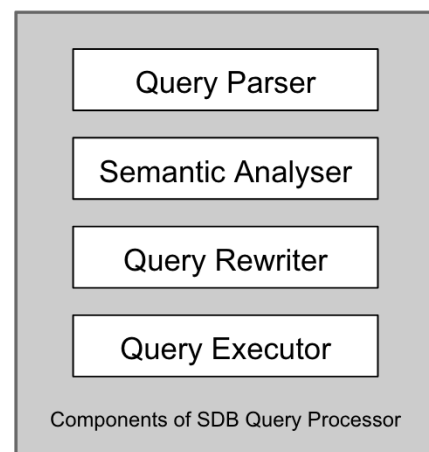


Figure 7: Components of SDB Proxy

SDB Query Parser parses a query string and transforms it into an *abstract syntax tree(AST)*. It is built based on HiveQL Parser, with additional tokens for column sensitivity declaration. For example, to declare a table T with sensitive column A and non-sensitive column B, the user has to submit query “CREATE TABLE T (A int ENC, B int)”, which is transformed into the AST in figure 8.

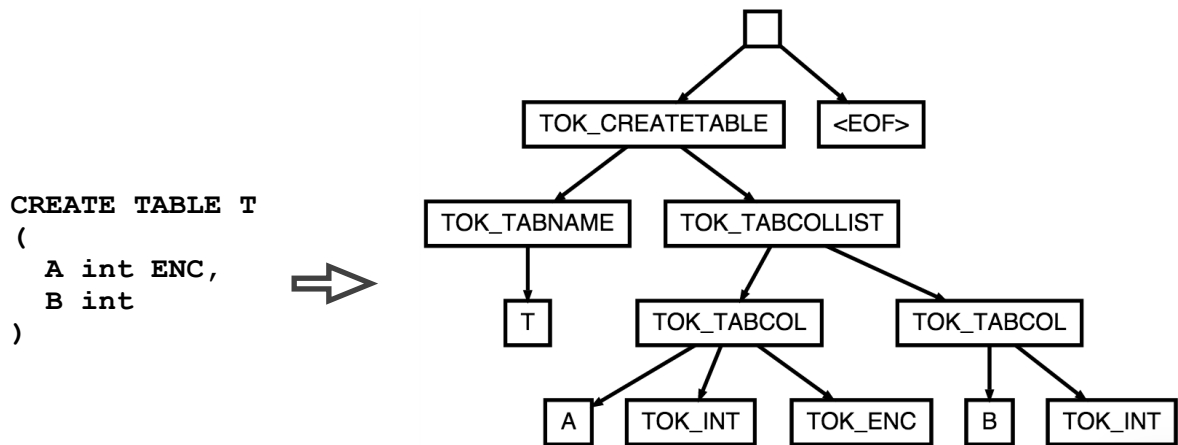


Figure 8: AST of “CREATE TABLE T (A int ENC, B int)”

SDB Semantic Analyzer transforms an AST into a *Logical Plan Tree(LPT)* as the preparation for query rewrite. The semantic analyser traverses the AST to construct a LPT, which is easier to manipulate with, and access the key store to

1. verify whether the reference of queried tables and columns are valid
2. annotate a column with its sensitivity and its column key, if any

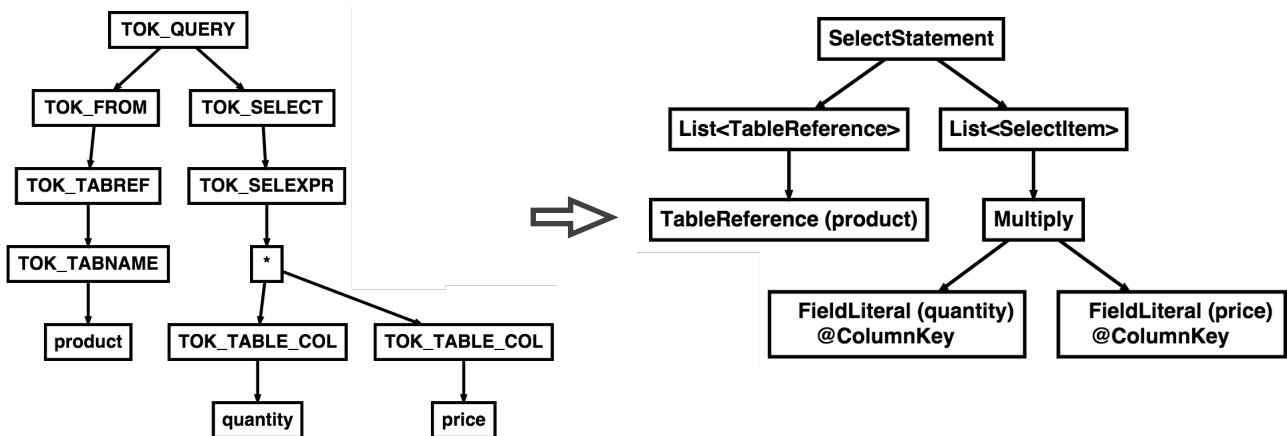


Figure 9: Transformation from AST to LPT

SDB Query Rewriter rewrites the LPT with secure operator, and constructs a *Physical Plan Tree(PPT)*. While traversing LPT, every time the rewriter identifies a normal arithmetic operator involving sensitive columns, it replaces the normal operator with the corresponding secure operator, based on SDB rewrite rule(discussed in detail in next section). Finally the rewrite construct a physical plan tree for SDB Executor to execute.

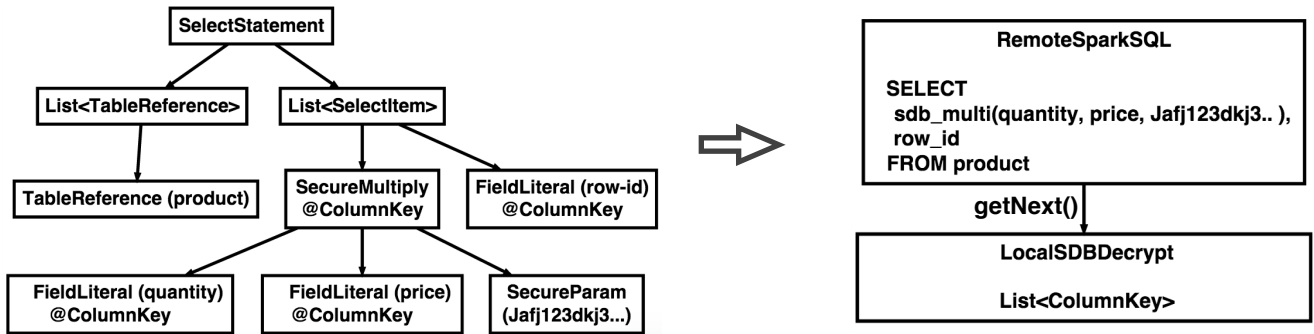


Figure 9: Transformation from LPT to PPT

SDB Executor executes the physical query plan and returns the result set. Finally, the executor executes the physical plan tree which normally splits the execution into two phases. In the first phase, it sends server-side query and waits for query result. In the second phase, it receives server-side query and decrypts query result. The decrypted result is then sent back to client application via connector.

3.3 Query Rewriting

We denote $ck_A = \langle m_A, x_A \rangle$ as the column key for column A, denote $\kappa_{AC} = \langle p, q \rangle$ as the client-side result of key update column A with target column key $\langle m_C, x_C \rangle$. Besides, we use ck_Z to denote the special column key $\langle 1, 0 \rangle$.

Function	Explanation
$\text{rand}(\rho_1, \rho_2)$	Client-side function which generates a random number co-prime with n , $\Phi(n)$
$\text{key_update_client}(ck_A, ck_C, cks, \rho_1, \rho_2)$	Client-side key update which computes κ_{AC}
$\text{decrypt}(C_e, ck_C, n, g, \text{row_id})$	Client-side decryption which computes $\llbracket c \rrbracket$
$\text{sdb_mul}(A_e, B_e, n)$	Server-side multiplication ($A \times B$) UDF
$\text{sdb_add}(A_e, B_e, S_e, \kappa_{AC}, \kappa_{BC}, n)$	Server-side addition ($A + B$ as C) UDF
$\text{sdb_compare}(A_e, n)$	Server-side comparison ($A > 0$) UDF
$\text{sdb_key_update}(A_e, S_e, \kappa_{AC}, n)$	Server-side key update UDF

Table 2: List of functions for query rewriting

3.3.1 Query Rewriting for Multiplication

According to the secret sharing scheme, query rewriting for secure multiplication is straightforward.

Input: $\text{SELECT } A \times B \text{ as } C \text{ FROM } T$

$$ck_c = \langle m_A \times m_B, x_A + x_B \rangle$$

Output: $\text{SELECT row_id, sdb_mul}(A_e, B_e, n) \text{ AS } C_e \text{ FROM } T$

Input: $\text{SELECT } A \times u \text{ as } C \text{ FROM } T$

$$ck_C = \langle m_A \times u, x_A \rangle$$

Output: $\text{SELECT row_id, } A_e \text{ AS } C_e \text{ FROM } T$

Rule 1: Rewrite Rule for Multiplication (EE & EC mode)

3.3.2 Query Rewriting for Addition

For secure addition in EE Mode, addition is achieved by applying key update to both A and B so that they share the same column key.

Input: SELECT A + B as C FROM T

$ck_C = \langle \text{rand}(\rho_1, \rho_2), \text{rand}(\rho_1, \rho_2) \rangle$
 $\kappa_{AC} = \text{key_update_client}(ck_A, ck_C, cks, \rho_1, \rho_2)$
 $\kappa_{BC} = \text{key_update_client}(ck_A, ck_B, cks, \rho_1, \rho_2)$

Output: SELECT row_id, sdb_add($A_e, B_e, S_e, \kappa_{AC}, \kappa_{BC}, n$) AS C_e FROM T

Rule 2: Rewrite Rule for Addition (EE mode)

For secure addition in EC mode, the expression $A+u$ can be considered as $A+(S \times u)$, as the plaintext of S is 1. Thus an EC addition operation can be done as an EE addition after EC Multiplication.

Input: SELECT A + u as C FROM T

$ck_{Su} = \langle u \times ms, x_S \rangle$
 $ck_C = \langle \text{rand}(\rho_1, \rho_2), \text{rand}(\rho_1, \rho_2) \rangle$
 $\kappa_{AC} = \text{key_update_client}(ck_A, ck_C, cks, \rho_1, \rho_2)$
 $\kappa_{SuC} = \text{key_update_client}(ck_{Su}, ck_C, cks, \rho_1, \rho_2)$

Output: SELECT row_id, sdb_add($A_e, S_e, S_e, \kappa_{AC}, \kappa_{SuC}, n$) as C_e FROM T

Rule 3: Rewrite Rule for Addition (EC mode)

3.3.3 Query Rewriting for Subtraction

For secure subtraction in EE Mode, the expression $A-B$ can be considered as $A+(B \times (-1))$, similar to addition in EC mode. Based on the modular equation $-1 \bmod n == (n-1) \bmod n$, we compute the column key of $-B$ as follows.

Input: SELECT A - B as C FROM T

$ck_{Bi} = \langle (n-1) \times ms, x_B \rangle$ // inverse the value of B
 $ck_C = \langle \text{rand}(\rho_1, \rho_2), \text{rand}(\rho_1, \rho_2) \rangle$
 $\kappa_{AC} = \text{key_update_client}(ck_A, ck_C, cks, \rho_1, \rho_2)$
 $\kappa_{BiC} = \text{key_update_client}(ck_{Bi}, ck_C, cks, \rho_1, \rho_2)$

Output: SELECT row_id, sdb_add($A_e, B_e, S_e, \kappa_{AC}, \kappa_{BiC}, n$) AS C_e FROM T

Rule 4: Rewrite Rule for Subtraction (EE mode)

For secure subtraction in EC mode, the expression $A-u$ can be considered as $A-S \times u$, thus we simply compute based on the rule of EE subtraction.

Input: SELECT A – u as C FROM T

$ck_{Su} = \langle u \times ms, x_s \rangle \quad // S \times u$
 $ck_{Sui} = \langle (n-1) \times ms_u, x_{Su} \rangle \quad // -(S \times u)$
 $ck_C = \langle \text{rand}(\rho_1, \rho_2), \text{rand}(\rho_1, \rho_2) \rangle$
 $\kappa_{AC} = \text{key_update_client}(ck_A, ck_C, cks, \rho_1, \rho_2)$
 $\kappa_{SuiC} = \text{key_update_client}(ck_{Sui}, ck_C, cks, \rho_1, \rho_2)$

Output: SELECT row_id, sdb_add(Ae, Se, Se, κ_{AC} , κ_{SuiC} , n) AS C_e FROM T

Rule 5: Rewrite Rule for Subtraction (EC mode)

3.3.4 Query Rewriting for Comparison

For secure comparison in EC mode, the expression $A > B$ can be considered as $(A - B) > 0$, to ensure no sensitive information is leaked, we randomize the result by $R \times (A - B) > 0$ at server side so that the result of $A - B$ is not revealed.

Input: SELECT A FROM T WHERE A > B

$ck_{Bi} = \langle (n-1) \times ms, x_B \rangle \quad // \text{inverse the value of } B$
 $ck_C = \langle \text{rand}(\rho_1, \rho_2), \text{rand}(\rho_1, \rho_2) \rangle$
 $\kappa_{AC} = \text{key_update_client}(ck_A, ck_C, cks, \rho_1, \rho_2)$
 $\kappa_{BiC} = \text{key_update_client}(ck_{Bi}, ck_C, cks, \rho_1, \rho_2) \quad // A - B$
 $ck_{RC} = \langle m_R \times m_C, x_B + x_C \rangle \quad // R \times (A - B)$
 $\kappa_{RCZ} = \text{key_update_client}(ck_{RC}, ck_Z, cks, \rho_1, \rho_2)$

Output: SELECT A, row_id FROM T

WHERE sdb_compare(sdb_key_update(sdb_mul(R_e , sdb_add(Ae, Be, Se, κ_{AC} , κ_{BiC}), n), S_e , κ_{RCZ} , n), n) > 0

Rule 6: Rewrite Rule for Comparison (EE mode)

For secure comparison in EC mode, the expression $A > u$ can be considered as $(A - S \times u) > 0$. Notice that we give the rewriting rule for $>$, to apply it to $=$ or $<$, simply swap $>$ with the target equality operator.

Input: SELECT A FROM T WHERE A > u

$ck_{Su} = \langle u \times ms, x_s \rangle \quad // S \times u$
 $ck_{Sui} = \langle (n-1) \times ms_u, x_{Su} \rangle \quad // -(S \times u)$
 $ck_C = \langle \text{rand}(\rho_1, \rho_2), \text{rand}(\rho_1, \rho_2) \rangle$
 $\kappa_{AC} = \text{key_update_client}(ck_A, ck_C, cks, \rho_1, \rho_2)$
 $\kappa_{SuiC} = \text{key_update_client}(ck_{Sui}, ck_C, cks, \rho_1, \rho_2) \quad // A - S \times u$
 $ck_{RC} = \langle m_R \times m_C, x_B + x_C \rangle$
 $\kappa_{RCZ} = \text{key_update_client}(ck_{RC}, ck_Z, cks, \rho_1, \rho_2) \quad // R \times (A - B)$

Output: SELECT A, row_id FROM T

WHERE sdb_compare(sdb_key_update(sdb_mul(R_e , sdb_add(Ae, Se, Se, κ_{AC} , κ_{SuiC})), S_e , κ_{RCZ} , n), n) > 0

Rule 7: Rewrite Rule for Comparison (EC mode)

Figure 10 shows the rewritten physical plan tree of EE Comparison query.

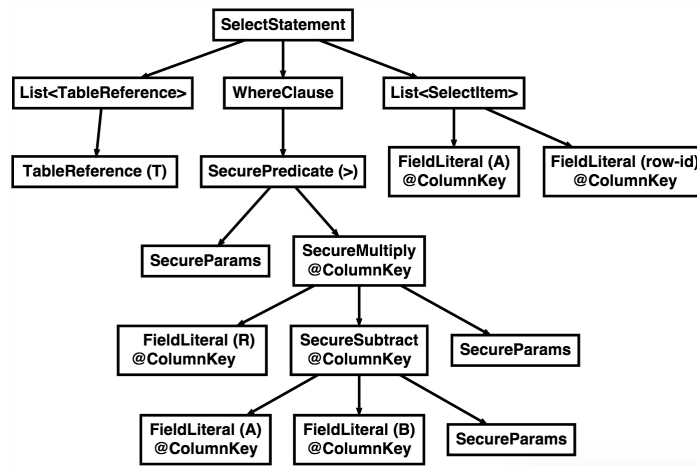


Figure 10: Rewritten LPT of SELECT A FROM T WHERE A > B

3.4 Web Interface

3.4.1 Web Interface for Uploading Data

We introduce the steps to upload a table of sensitive data as follows.

1. In the *SQL Editor* page, type in the CREATE TABLE query
2. Append keyword “ENC” to every sensitive field
3. Click “Execute SQL Query” button

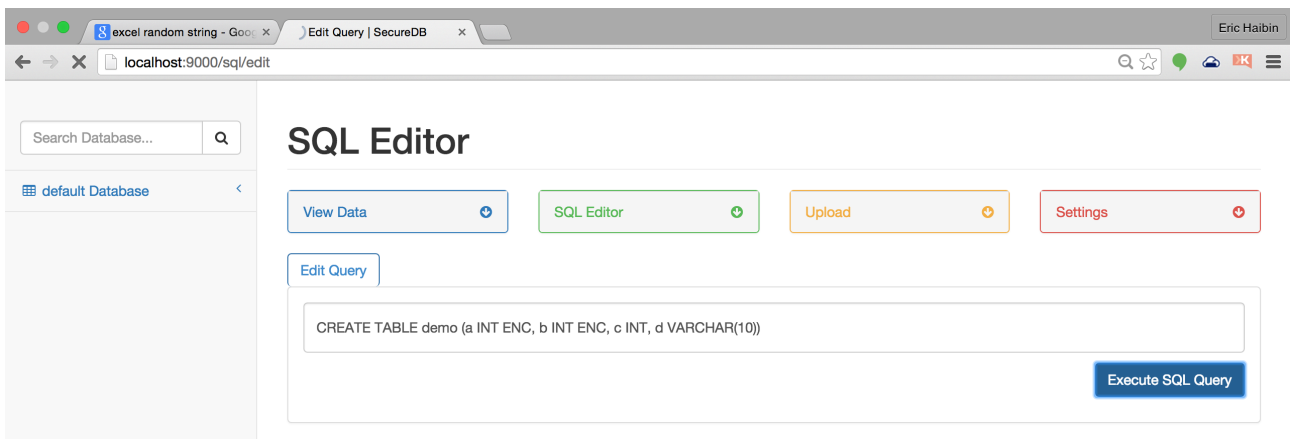


Figure 11: Web interface to create a table with two sensitive fields (A, B)

4. Receive confirmation message “Create Table Successful” after success

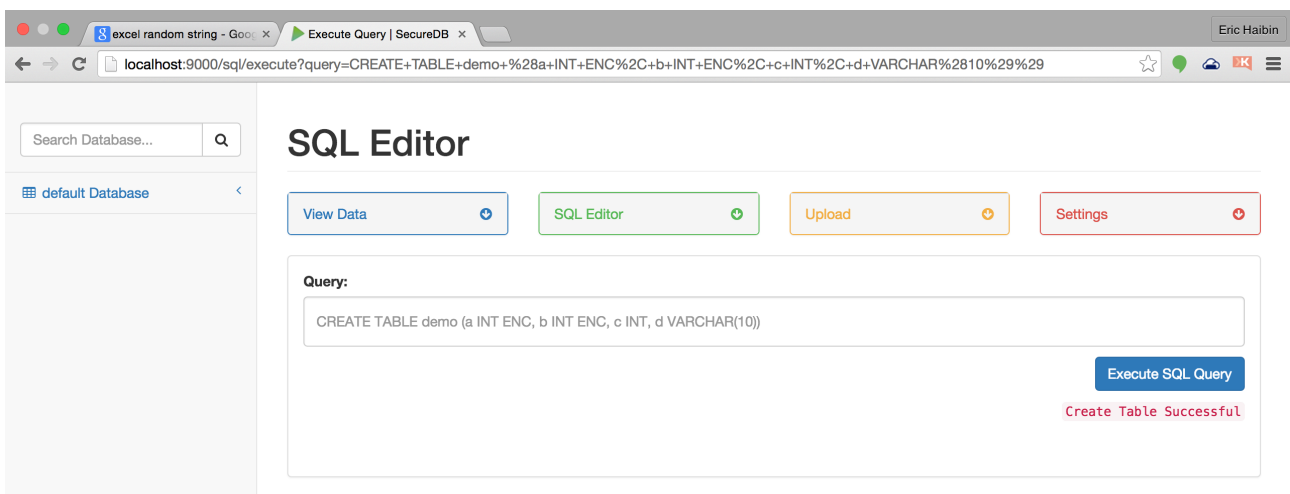


Figure 12: Confirmation message of table creation success

5. In the *Upload Data* page, select table name and data source file to upload.
6. Click “Upload Data” button to start the upload process.

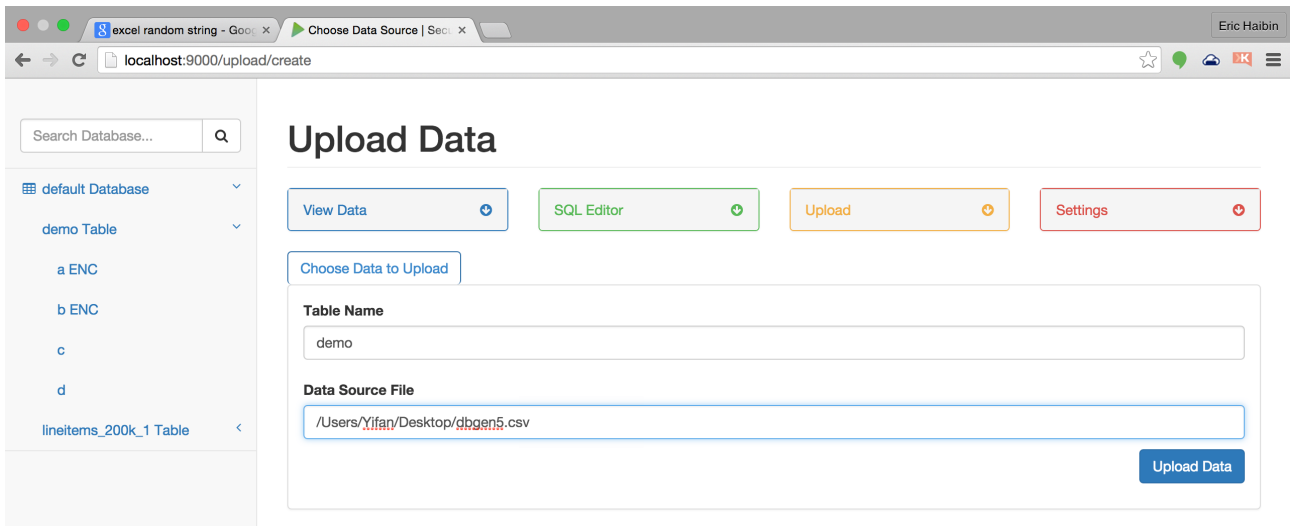


Figure 13: Web interface to upload data to table

3.4.2 Web Interface for Query

We introduce the interface for query as follows.

1. Type in the query in the SQL Editor page
2. Click “Execute SQL Query” button
3. Query result will be displayed as in Figure 14

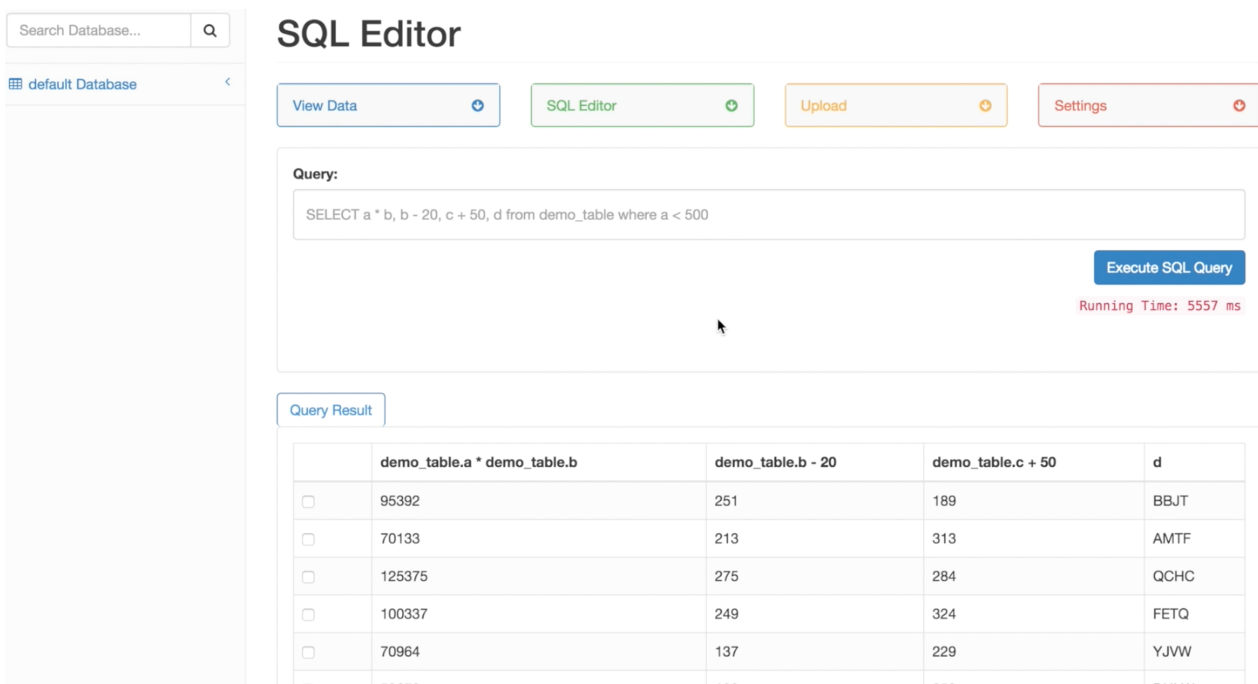


Figure 14: Web interface to execute query

- At “Execution Analysis” section, the query executed at cluster server is displayed
- Client and server side query execution cost are visualized in charts



Figure 15: Query Performance Analysis Section

Cloud service provider’s view of encrypted sensitive data is illustrated in Figure 16.



Figure 16: Cloud Service Provider’s View of Data

4. Experiments and Results

4.1 Experiment Environment

The prototype of SDB is deployed on a cluster of 10 machines. Every server machine is equipped with eight Intel(R) Core(TM) i7-3770 CPUs, four 4GB Kingston DDR3 1333MHz RAM, running on Ubuntu Release 12.04. We use Apache Hadoop 2.4.1, Spark 1.1.0 and Hive 0.12.0, running on Java 1.6.0. The Hadoop File System is configured to have 3 replication of files and 64 MB data blocks. Spark executors are configured to have memory of 256 MB. The SDB proxy resides on the client machine with a single Intel 2.40 GHz Core i5 CPU.

In our experiment, we answer the question of *Q1*: “How important it is for the secure database operators to be data interoperable in a cloud database environment?” Second, we address the questions of *Q2*: “How much performance overhead is introduced by SDB to protect sensitivity? Is it practical in query processing?”

We prepare a synthetic table *T* with two sensitive columns *A* and *B*. The values of each column are generated with a uniform distribution. We execute 3 kinds of queries:

- [*Range Query EE Mode*] `SELECT A, B from T WHERE A < q`
- [*Range Query EC Mode*] `SELECT A, B from T WHERE A < B`
- [*Count Query*] `SELECT count(A) from T WHERE A < q`

The parameter *q* controls the selectivity of a range query. A smaller *q* leads to a smaller result size.

We split the total query time into two components. (1) *Server Cost*: the time taken for the cluster server to execute rewritten queries. (2) *Client Cost*: the time taken for SDB proxy to parse, analyse, rewrite queries and decrypt query results.

4.2 Experiment Result

4.2.1 SDB vs. DBQ

To answer QI , we compare the performance of SDB with *Decrypt-Before-Query Model(DBQ)*.

DBQ is configured with the same cluster machines at server-side and runs Spark standalone mode at client side. Since no computation is done on sensitive data at server side, sensitive columns are encrypted using RSA encryption. For a sample query “SELECT A, B from T WHERE A < q”, DBQ’s server cost is simply the time taken to ship all the data back to client. DBQ client then decrypts all sensitive data to execute the range query.

As shown in Figure 17, the query processing time of DBQ is dominated by client cost, while SDB is dominated by server cost. We conclude that the client cost of SDB is significantly less than that of DBQ.

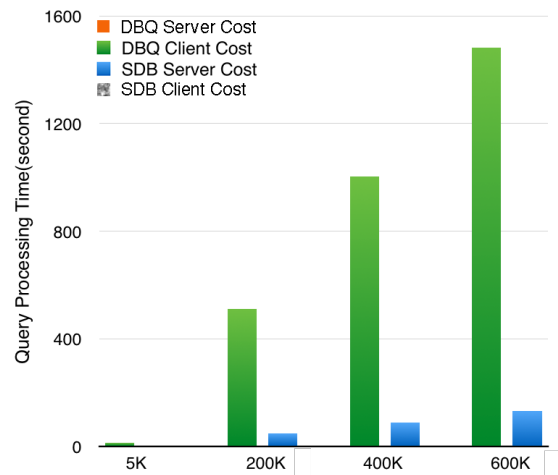


Figure 17: DBQ vs. SDB

4.2.2 SDB Cost Breakdown

We further break down the components of SDB cost as follows. Figure 18 shows the client/server cost of a range EC query(on 200K dataset). With increasing selectivity, server cost stays stable, while client cost increases linearly. Figure 19 shows the further detail of client cost of a range EC query(on 5K dataset). This shows that decryption cost increases as the number of results to decrypt increases. Decryption cost is the major contributor to client cost. We observe that the decryption cost is significant, which may become the bottleneck of query execution when result size is large. Therefore, improving the performance of decryption is one of the key optimizations to be done in the future.




Figure 18: Client/Server Cost v.s. Selectivity

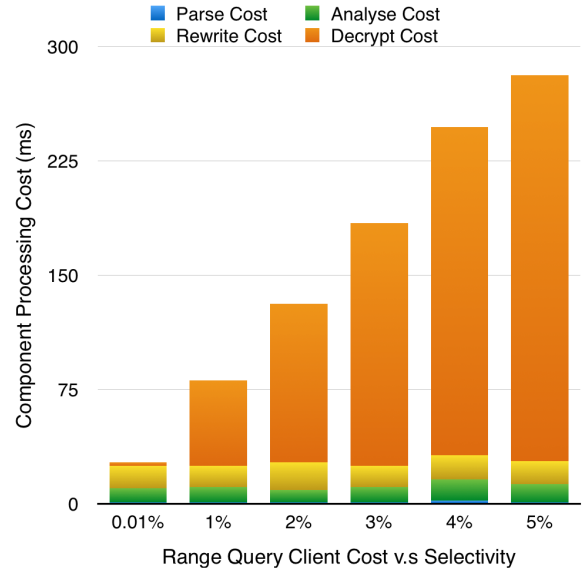


Figure 19: Client Costs v.s. Selectivity

4.2.3 SDB vs. SparkSQL

To answer Q_2 , we evaluate the performance of SDB as follows. We first execute the three queries on SparkSQL directly with all columns stored as plaintext values, bypassing all secure operators.

The time to execute under this scenario is denoted as T_{SparkSQL} . Then we execute the same queries on encrypted data on SDB, and use T_{SDB} to denote the execution time. The ratio $T_{\text{SDB}}/T_{\text{SparkSQL}}$ captures the degree to which SDB’s encryption and secure protocol slows down the entire query processing time.

We observe that SDB introduces a 170 times slowdown of the query processing time. This is mainly the result of expensive modular exponent computation of big integers. In particular, the secure comparison involves exponent computation at least three times. This forms another aspect of future optimization we will focus on.

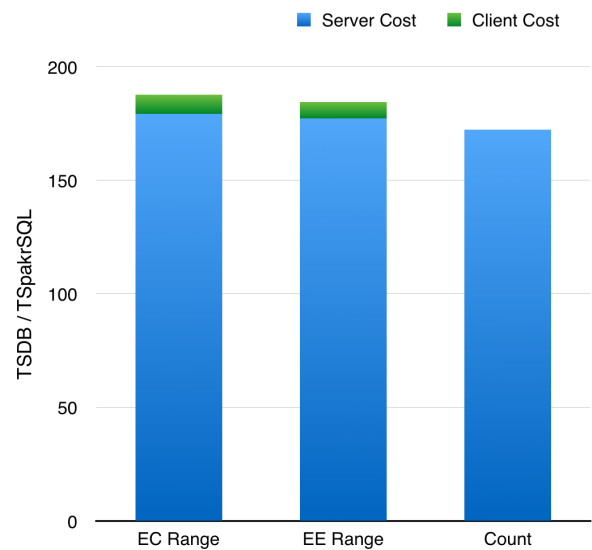


Figure 20: TSDB / TSpark

Finally, Figure 18 shows the client/server cost for the three queries with different data size.




Figure 21 Client/Server Cost vs. Data Size

5. Conclusion and Future Works

The future work of SDB is mainly divided into two parts: secure operator extension and cryptographic optimization.

5.1 Secure Operator Extension

SDB is designed to support a wide range of query with several secure operators. In Appendix 7.1, we included the query rewriting rules for the following secure operators which are not fully implemented yet:

- Secure multiplication, addition, subtraction, comparison in EP mode
- Secure cartesian product
- Secure aggregation function: `sum()`, `avg()`

Since the query executor of SDB proxy is implemented with the *iterator* pattern, SDB proxy can also be extended to adopt the *split client/server execution* approach proposed by MONOMI[6]. This way, in the case of un-supported queries, SDB proxy shall identify the parts of query supported at cluster server and push them to server as much as possible. Then SDB proxy performs the rest of the queries.

5.2 Cryptographic Optimization

As we mentioned in the experiment, the decryption cost dominates other costs at client side.

Therefore, it's crucial to improve the performance of decryption. During the computation of the item key of the same column but different rows, it can be computed as $mb^r \bmod n$ with

$b = g^x$. Thus it can be viewed as exponential computation with the same base, which can be

optimized by *Exponentiation by squaring*[7]. To optimize such computation, we first pre-compute a

table of $b^1, b^2, b^4, b^8, b^{16}, \dots$. To compute b^{13} , we only need to retrieve b^1, b^2, b^8 from table and multiply them together.

At server side, currently User Defined Functions at server side are all written in Java, whose modular exponential computation is slow. Instead, we can achieve 4x better performance[8] by calling native math methods in GMP library, which is written in C++.

Besides, there are studies on secure indexing techniques[9], which groups the value of sensitive columns into coarse ranges. Such index filtering can be applied before we process a query using our secure operators.

5.3 Conclusion

In conclusion, we gave a comprehensive description of SDB's query rewriting and analysis of performance. By deploying such secret sharing scheme that supports multiple secure operators with data interoperability, SDB is a query processing system that is both secure and practically efficient.

6. References

- [1] Wong, W. K., Kao, B., Cheung, D. W. L., Li, R., & Yiu, S. M. (2014, June). Secure query processing with data interoperability in a cloud database environment. *In Proceedings of the 2014 ACM SIGMOD international conference on Management of data* (pp. 1395-1406). ACM.
- [2] Gentry, Craig, Shai Halevi, and Nigel P. Smart. "Homomorphic evaluation of the AES circuit." *Advances in Cryptology—CRYPTO 2012*. Springer Berlin Heidelberg, 2012. 850-867.
- [3] Agrawal, R., Kiernan, J., Srikant, R., & Xu, Y. (2004, June). Order preserving encryption for numeric data. *In Proceedings of the 2004 ACM SIGMOD international conference on Management of data* (pp. 563-574). ACM.
- [4] Bajaj, Sumeet, and Radu Sion. "TrustedDB: A Trusted Hardware-Based Database with Privacy and Data Confidentiality." *Knowledge and Data Engineering, IEEE Transactions on* 26.3 (2014): 752-765.
- [5] Popa, Raluca Ada, et al. "Cryptdb: protecting confidentiality with encrypted query processing." *Proceedings of the Twenty-Third ACM Symposium on Operating Systems Principles*. ACM, 2011.
- [6] Tu, S., Kaashoek, M. F., Madden, S., & Zeldovich, N. (2013, March). *Processing analytical queries over encrypted data*. *In Proceedings of the VLDB Endowment* (Vol. 6, No. 5, pp. 289-300). VLDB Endowment.
- [7] Menezes, A. J., Van Oorschot, P. C., & Vanstone, S. A. (1996). *Handbook of applied cryptography*. CRC press.
- [8] Scott (2014, February 14). Faster RSA in Java with GMP. Retrieved 19 April, 2015, from <https://corner.squareup.com/2014/02/faster-rsa-jnagmp.html>
- [9] Hore, B., Mehrotra, S., & Tsudik, G. (2004, August). A privacy-preserving index for range queries. *In Proceedings of the Thirtieth international conference on Very large data bases-Volume 30* (pp. 720-731). VLDB Endowment.

7. Appendix

7.1 Rewrite Rules for Other Secure Operators

Function	Explanation
SIES(row_id, Ck _{ROW_ID} , ρ1, ρ2)	Client-side function which decrypts row_id with SIES
sdb_int_add(row_id1, row_id2, n)	Server-side row-id addition UDF
sdb_cartprod(A _e , S _{e2} , p _A , n)	Server-side cartesian product (A ⊗ B as C) UDF

Secure operation in EP mode is similar to EC mode that plaintext is treated as a column with column key <1, 0>.

```
SELECT A x P as C FROM T
```

```
ckPN = < rand(ρ1, ρ2), rand(ρ1, ρ2) >
κZPN = key_update_client(ckZ, ckPN, ckS, ρ1, ρ2)
ckC = < mA × mPN, xA + xPN >
```

```
SELECT row_id, sdb_mul(Ae, sdb_key_update(P, Se, κZPN, n), n) AS Ce FROM T
```

Rule 8: Rewrite Rule for Multiplication (EP Mode)

```
SELECT (A+P) as C FROM T
```

```
ckPN = < rand(ρ1, ρ2), rand(ρ1, ρ2) >
κZPN = key_update_client(ckZ, ckPN, ckS, ρ1, ρ2)
ckPN = < rand(ρ1, ρ2), rand(ρ1, ρ2) >
κAC = key_update_client(ckA, ckC, ckS, ρ1, ρ2)
κPNC = key_update_client(ckPN, ckC, ckS, ρ1, ρ2)
```

```
SELECT row_id, sdb_add(Ae, sdb_key_update(P, Se, κZPN, n), Se, κAC, κPNC, n) AS Ce FROM T
```

Rule 9: Rewrite Rule for Addition (EP Mode)

```
SELECT (A-P) as C FROM T
```

```
ckPN = < rand(ρ1, ρ2), rand(ρ1, ρ2) >
κZPN = key_update_client(ckZ, ckPN, ckS, ρ1, ρ2)
ckPNi = < (n-1) × mPN, xPN > // inverse the value of P
ckC = < rand(ρ1, ρ2), rand(ρ1, ρ2) >
κAC = key_update_client(ckA, ckC, ckS, ρ1, ρ2)
κPNiC = key_update_client(ckPNi, ckC, ckS, ρ1, ρ2)
```

```
SELECT row_id, sdb_add(Ae, sdb_key_update(P, Se, κZPN, n) Se, κAC, κPNiC, n) AS Ce FROM T
```

Rule 10: Rewrite Rule for Subtraction (EP Mode)

SELECT * FROM T WHERE A > P

```

ckPN = < rand(ρ1, ρ2), rand(ρ1, ρ2) >
κZPN = key_update_client(ckZ, ckPN, cks, ρ1, ρ2)
ckPNi = < (n-1) × mPN, xPN >
ckC = < rand(ρ1, ρ2), rand(ρ1, ρ2) >
κAC = key_update_client(ckA, ckC, cks, ρ1, ρ2)
κPNiC = key_update_client(ckPNi, ckC, cks, ρ1, ρ2)
ckRC = < mR × mC, xR + xC >
κRCZ = key_update_client(ckRC, ckZ, cks, ρ1, ρ2)

```

SELECT * FROM T

WHERE sdb_key_update(sdb_mul(R_e, sdb_add(A_e, sdb_key_update(P, S_e, ck_{PN}, n), S_e, κ_{AC}, κ_{PNiC})), S_e, κ_{RCZ}, n) => 0

Rule 11: Rewrite Rule for Comparison (EP Mode)

For secure Cartesian product between two tables, we make use of the additive homomorphic

property of row_id's encryption (SIES) so that T1 JOIN T2 can be decrypted with the new row_id.

SELECT A, B FROM T1 JOIN T2

```

pA = (xs2)-1 × xA mod Φ(n)
ckA' = < mA × (mS2)pA, xA >
pB = (xs1)-1 × xB mod Φ(n)
ckB' = < mB × (mS2)pB, xB >

```

SELECT sdb_int_add(T1.row_id, T2.row_id, n) as row_id, sdb_cartprod(A_e, S_{e2}, p_A, n), sdb_cartprod(B_e, S_{e1}, p_B, n) FROM T1 JOIN T2

Rule 12: Rewrite Rule for Cartesian Product

To perform secure aggregation function SUM(), we perform a key update operation on target

column A with a special column key < rand, 0 >. Particularly, we insert with a different row_id in

case the result of SUM() is referred in other nested queries. Finally, The decryption of SUM

ciphertext is handled differently by equation $[[C]] = m_z \times C_e$.

SELECT sum(A) as C FROM T

```

CkSUM = < rand(ρ1, ρ2), 0 >
κASUM = key_update_client(ckA, ckSUM, cks, ρ1, ρ2)
ckPNi = < (n-1) × mPN, xPN >
row_id = rand(ρ1, ρ2)
enc_row_id = SIES(row_id, CkROW_ID, ρ1, ρ2)

```

SELECT enc_row_id, sdb_sum(sdb_key_update(A_e, S_e, κ_{ASUM}, n)) FROM T

Rule 13: Rewrite Rule for SUM()

With computed values of SUM() and COUNT(), AVG() can be easily computed by dividing SUM

with COUNT. $[[C]] = m_z \times C_e / \text{count}$.

SELECT avg(A) as C FROM T

$Ck_{SUM} = \langle \text{rand}(\rho_1, \rho_2), 0 \rangle$
 $\kappa_{ASUM} = \text{key_update_client}(ck_A, ck_{SUM}, ck_s, \rho_1, \rho_2)$
 $ck_{PN_i} = \langle (n-1) \times mpn, x_{PN} \rangle$
 $row_id = \text{rand}(\rho_1, \rho_2)$
 $enc_row_id = \text{SIES}(row_id, Ck_{ROW_ID}, \rho_1, \rho_2)$

SELECT enc_row_id, sdb_sum(sdb_key_update(A_e, S_e, κ_{ASUM}, n)) as sum, count(A_e) as count FROM T

Rule 14: Rewrite Rule for AVG()

7.2 Experiment Data

	5K	200K	400K	600K
SDB Server	4.246	46.934	89.619	132.614
SDB Client	0.082	2.107	4.372	6.366
DBQ Server	0.230	8.93	17.860	26.8
DBQ Client	14.33	510.92	1003.931	1483.658

Figure 17: SELECT A, B from T WHERE A < p, 1% selectivity (seconds)

	0.01%	1%	2%	3%	4%	5%
Server Cost	44.239	44.327	45.035	45.502	45.627	45.714
Client Cost	0.071	2.107	4.048	5.791	7.789	9.835

Figure 18: SELECT A, B from T WHERE A < p, on 200K records (seconds)

	0.01%	1%	2%	3%	4%	5%
Parse Cost	1	1	1	1	2	1
Analyse Cost	9	10	8	10	14	12
Rewrite Cost	15	14	18	14	16	15
Decrypt Cost	2	56	104	159	215	253

Figure 19: SELECT A, B from T WHERE A < p, on 5K records (milliseconds)

	EC Range	EE Range	Count
Server Cost	44827	44867	43728
Client Cost	2107	1786	35
SparkSQL Cost	250	253	254

Figure 20: TSDB vs TSparkSQL on 200K records (milliseconds)

	5K	200K	400K	600K
Server Cost	4.633	44.867	84.053	134.601
Client Cost	0.14	1.786	3.414	5.358

Figure 20A: SELECT A, B from T WHERE A < B, 1% selectivity (seconds)

	5K	200K	400K	600K
Server Cost	4.164	44.827	85.247	126.248
Client Cost	0.082	2.107	4.372	6.366

Figure 20B: SELECT A, B from T WHERE A < p, 1% selectivity (seconds)

	5K	200K	400K	600K
Server Cost	4.45	43.728	82.991	122.435
Client Cost	0.017	0.035	0.021	0.021

Figure 20B: SELECT COUNT(A) from T WHERE A < p, 1% selectivity (seconds)