Available Online at www.ijcsmc.com

International Journal of Computer Science and Mobile Computing

A Monthly Journal of Computer Science and Information Technology

ISSN 2320–088X

IJCSMC, Vol. 4, Issue. 6, June 2015, pg.592 – 600

RESEARCH ARTICLE

Confidential and Secure Query Services in the Cloud with RASP

Neeta R.Padashetty[1], Asst. Prof. Ranjana B.Nadagoudar[2]

[1]Computer Science and Engineering, VTU, India

[2]Computer Science and Engineering, VTU, India

[1] neetapadashetty@gmail.com; [2] ranjanapriya8@gmail.com

Abstract— Range query is one of the most frequently used queries for online data analytics. Providing such a query service could be expensive for the data owner. With the development of services computing and cloud computing, it has become possible to outsource large databases to database service providers and let the providers maintain the range-query service. With outsourced services, the data owner can greatly reduce the cost in maintaining computing infrastructure and data-rich applications. We propose the Random Space Encryption (RASP) approach that allows efficient range search with stronger attack resilience than existing efficiency-focused approaches. The RASP data perturbation method combines order preserving encryption, dimensionality expansion, random noise injection, and random projection, to provide strong resilience to attacks on the perturbed data and queries. It also preserves multidimensional ranges, which allows existing indexing techniques to be applied to speedup range query processing. The kNN-R algorithm is designed to work with the RASP range query algorithm to process the kNN queries. We use RASP to generate indexable auxiliary data that is resilient to prior knowledge enhanced attacks. Range queries are securely transformed to the encrypted data space and then efficiently processed with a two-stage processing algorithm.

Keywords— RASP, query services in the cloud, privacy, range query, kNN query

I. INTRODUCTION

This document is a template. An electronic copy can be downloaded from the conference website. For questions on paper guidelines, please contact the conference publications committee as indicated on the conference website. Information about final paper submission is available from the conference website. With the wide deployment of cloud infrastructures, it has become popular to host services and big data in public clouds. This new paradigm is especially attractive for data-intensive query and analysis services for its great scalability and significant cost savings. It is well known that maintaining and mining data incurs much higher cost than initial data acquisition. By moving data services to the cloud, data owners can cut costs in almost every aspect of managing and mining data. However, data privacy is still haunting data owners minds as the underlying infrastructure is out of their control. In particular, data owners may not be aware of information leakage, which can happen in all kinds of possibilities, if the cloud provider does not want to report the leakage. Cloud computing is an emerging technology. Posting data-intensive query services in the cloud is increasingly popular because of the unique advantages in scalability and cost saving.

With the cloud infrastructure data owner can scale up or down the services and pay as they use the server. Work load of query services in the public cloud is highly dynamic and it will be expensive to serve the queries with the in-house infrastructure. However service provider may lose the control over the data privacy and

confidentiality has become the major concerns. Curious service providers can possibly make a copy of the database or eavesdrop users queries, which will be difficult to detect and prevent in the cloud infrastructures. While new approaches are needed to preserve data confidentiality and query privacy, the efficiency of query services and the benefits of using the clouds should also be preserved. It will not be meaningful to provide slow query services as a result of security and privacy assurance. It is also not practical for the data owner to use a significant amount of in-house resources, because the purpose of using cloud resources is to reduce the need of maintaining scalable in-house infrastructures.

A straightforward method is to encrypt datasets before exporting them to the cloud. However, searchable encryption is very challenging, showing limited successes in some specific areas such as document search [4]. Boneh et al. [2] showed that it is possible to construct a public-key system for range query, which is one of the basic database queries (another popular one is k nearest neighbor (kNN) query as we will discuss). However, it requires a significant amount of storage and computational costs, only applicable to linear scan of the entire database. Database queries such as range and kNN queries normally demand fast processing time (logarithmic or sub linear time complexity) with the support of indexing structures. However, if not impossible, there is no efficient indexing structure developed for encrypted data yet, which renders the current encryption schemes [2] unusable for search in large databases. We recently proposed the RAndom Space Perturbation (RASP) method [5] for the protection of tabular data, which is secure under the assumption of limited adversarial knowledge - only the perturbed data and the data distributions are known by adversaries. This assumption is appropriate in the context of cloud computing. The RASP perturbation is a unique combination of Order Preserving Encryption (OPE) [1], dimensionality expansion, noise injection, and random projection, which provides sufficient protection for the privacy of query services in the cloud. It has a number of unique features, such as preserving the topology of range query, non-deterministic results for duplicate records, and resilience to distributional attacks [5]. We develop the secure half-space query transformation method that casts any enclosed range in the original space to an irregularly shaped range in the perturbed space. Therefore, we are able to use a two-stage range query processing method: an existing multidimensional index, such as R*-Tree in the perturbed space is used to find out the records in the bounding box of the irregularly shaped range, which is then filtered with the transformed query condition. This processing strategy is fast and secure under the security assumption.

To allow the readers to fully appreciate the intuition and the ideas behind the RASP based perturbation and query processing, we propose this RASP Query Services (RASP-QS) demonstration system. This system consists of the following major components: (1) the user interface for perturbation parameter generation that allows users to observe the details of RASP perturbation, (2) the visualization of the two-stage range query processing procedure to understand the transformed query ranges and the query results, (3) the visualization of the progressive steps in the kNN query processing that is based on RASP range query processing, and (4) the performance comparison on index-aided processing on non-encrypted data, linear-scan query processing on encrypted data [2], and the RASP query processing.

II. RELATED WORK

A. Order Preserving Encryption

Order preserving encryption technique [1], Encryption is a well known technology for protecting sensitive data Integration of encryption method with database system cause performance reduction example if a column of contain sensitive information is encrypted, and is used in query predicate with a comparison operator an entire table scan would be needed to evaluate the query. Reason is that current encryption techniques do not preserve order and there data base indices such as B-tree cannot be used. Order preserving encryption allows comparison operation to be applied directly on encrypted data without decrypting the operands MAX, MIN, and COUNT queries can be directly processed over encrypted data. “Group by” and “order by” operations can also be applied. SUM, AVG and “group by” values need to be decrypted. A value in the column can be modified or a new value can be inserted in a column without requiring changes in the encryption of other values. OPES can easily be integrated with the existing database system as it has been designed to work with the existing indexing structure such as B-trees.

B. Crypto Index

A Privacy-Preserving Index for Range Queries[6] In the database-as-a-service (DAS) model since data is stored at the service provider, many security and privacy challenges arise. Most approaches to DAS define a notion of a security perimeter around the data owner. Environment within the perimeter is trusted (client/data owners) whereas environment outside the perimeter is not (service provider). Query processing: Data is stored in an encrypted form outside the perimeter but accessed within. Split the query Q into two components Qsec + Qinsec where Qinsec executes at the server on the encrypted representation to compute a result for Q and Qsec executes within the security perimeter to filter out the false positive. For the purpose of splitting the queries

Bucketization approach is used, each of the bucket is identified by a tag. These bucket tags are maintained as an index and are utilized by server to process the queries.

Drawback of OPE and crypto-index

OPE and crypto-index assumes the attacker knows only the cipher text. However, If the attacker has some prior knowledge, such as the attribute domains (maximum and minimum values), the attribute distributions, and even a few pairs of plaintext and cipher text, these encryption methods will be vulnerable to attacks.

C. RASP efficient multi-dimensional range queries

RASP: Efficient Multidimensional Range Query on Attack Resilient Encrypted Databases[7] Range query is the most frequently used query in online data analytics (OLAP) that requires the service provider to quickly respond to concurrent user queries. Most existing encryption based approaches require linear scan over the entire database, which is inappropriate for online data analytics on large databases. It was reported that maintaining data and supporting query-based services incur much higher cost than initial data acquisition.

RAndom Space encryption (RASP) approach for efficient range query processing on encrypted data, assume the outsourced data are multidimensional data and thus the data records can be treated as vectors (or points) in the multidimensional space. The RASP method randomly transforms the multidimensional space, while preserving the convexity of datasets. The framework assumes a secure proxy server at the client side that handles data encryption/decryption and query encryption. The data owner and authorized users submit the original data and queries to the proxy server; the proxy server then sends the encrypted data/queries to the service provider. The service provider is able to index the encrypted data and use it to efficiently process encrypted queries.

Drawback

Client need to take care of most of the things like encryption, decryption, indexing etc.

D. Private information retrieval

There is a significant risk to the privacy of the user, since a curious database operator can follow the user's queries and infer what the user is after. One thing a user can do to preserve his privacy is to ask for a copy of the whole database [8].

Same database is replicated at several sites, viewing the database as binary string x=x1, x2…….xn of length n. identical copies of strings are stored by k>=2 servers. The user has some index i, and he is interested in obtaining the value of bit X_i. To achieve this goal, the user queries each of the server and gets replies from which the direct bit Xi can be computed.

Drawback

Cost increases because of creating more than one copy of database.

E. Nearest neighbour search with strong location privacy

Nearest Neighbor Search with Strong Location Privacy [9] Applications like GPS in mobile devices facilitate the location based services which is an emerging application in the wireless market. Special queries pose an additional threat to privacy because location of a query may be sufficient to reveal sensitive information about the queries.

Location dependent queries may disclose sensitive information about an individual health, financial information, political affiliation etc. for example, user wishes to find the nearest restaurant, to get the information user may choose to transmit the query through an anonymous network that hides his/her IP address. It is not sufficient to hide IP address, if the information like co-ordinates of queries and background knowledge is known then information can be easily be hacked.

The solution can be classified as

- Location obfuscation.
- Data transformation.
- Private information retrieval.

III.METHODOLOGY

A. *Query Services In The Cloud*

Query is mainly used to search. Queries are constructed by using structured query language. It is mainly used to retrieving the needed information from the database. Query services are the method for services that are exposed through an implementation of service provider. Here by using RASP, range query and kNN query in cloud provide secure, fast storing and retrieving process of encryption and decryption of a data from database.

Range query is an important type of query for many data analytic tasks from simple aggregation to more sophisticated machine learning tasks. Let T be a table and X_i, X_j, and X_k be the real valued attributes in T, and a and b be some constants. Take the counting query for example. A typical range query looks like

select count (*) from T

where $X_i \in [a_i,b_i]$ and $X_j \in (a_j,b_j)$ and $X_k=a_k$

which calculates the number of records in the range defined by conditions on X_i, X_j, and X_k. Range queries may be applied to arbitrary number of attributes and conditions on these attributes combined with conditional operators "and"/"or." We call each part of the query condition that involves only one attribute as a simple condition. A simple condition like $X_i \in [a_i,b_i]$ can be described with two half space conditions $X_i \leq b_i$ and $-X_i \leq -a_i$. Without loss of generality, we will discuss how to process half-space conditions like $X_i \leq b_i$ in this paper. A slight modification will extend the discussed algorithms to handle other conditions like $X_i < b_i$ and $X_i = b_i$.

kNN query is to find the closest k records to the query point, where the Euclidean distance is often used to measure the proximity. It is frequently used in location based services for searching the objects close to a query point, and also in machine learning algorithms such as hierarchical clustering and kNN classifier. A kNN query consists of the query point and the number of nearest neighbors, k.

B. *System Architecture*

We assume that a cloud computing infrastructure, such as Amazon EC2, is used to host the query services and large data sets. The purpose of this architecture is to extend the proprietary database servers to the public cloud, or use a hybrid private-public cloud to achieve scalability and reduce costs while maintaining confidentiality.

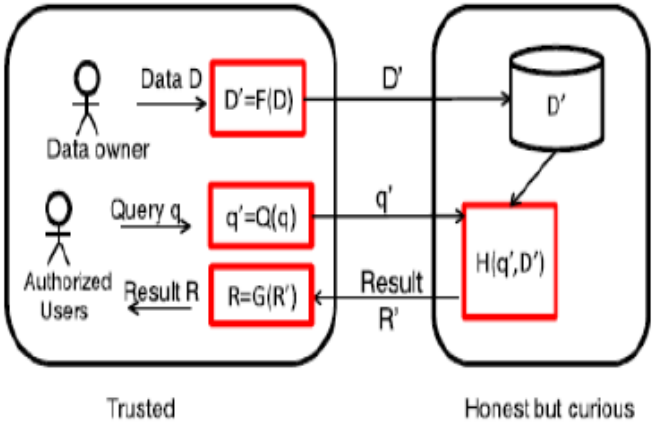

Figure: 1. The system architecture for RASP-based query services.

Each record x in the outsourced database contains two parts: the RASP-processed attributes D'=F(D,K) and the encrypted original records, Z = E(D,K'), where K and K' are keys for perturbation and encryption, respectively. The RASP-perturbed data D' are for indexing and query processing. [**Figure-1]** shows the system architecture for both RASP-based range query service and kNN service.

There are two clearly separated groups: the trusted parties and the untrusted parties. The trusted parties include the data/service owner, the in-house proxy server, and the authorized users who can only submit queries. The data owner exports the perturbed data to the cloud. Meanwhile, the authorized users can submit range queries or kNN queries to learn statistics or find some records. The untrusted parties include the curious cloud provider who hosts the query services and the protected database. The RASP-perturbed data will be used to build indices to support query processing.

There are a number of basic procedures in this framework:

1) F(D) is the RASP perturbation that transforms the original data D to the perturbed data D'; 2) Q(q) transforms the original query q to the protected from q' that can be processed on the perturbed data; and 3) H(q',D') is the query processing algorithm that returns the result R'. When the statistics such as SUM or AVG of a specific dimension are needed, RASP can work with partial homomorphic encryption such as Paillier encryption to compute these statistics on the encrypted data, which are then recovered with the procedure G'(R').

C. Threat Model

The cloud server is considered as "honest-but-curious" in our model, which is consistent with related works on cloud security. Specifically, the cloud server acts in an "honest" fashion and correctly follows the designated protocol specification. However, it is "curious" to infer and analyze data (including index) in its storage and message flows received during the protocol so as to learn additional information.

Assumptions: Our security analysis is built on the important features of the architecture. Under this setting, we believe the following assumptions are appropriate:

- Only the authorized users can query the proprietary database. Authorized users are not malicious and will not intentionally breach the confidentiality. We consider insider attacks are orthogonal to our research; thus, we can exclude the situation that the authorized users collude with the untrusted cloud providers to leak additional information.
- The client-side system and the communication channels are properly secured and no protected data records and queries can be leaked.
- Adversaries can see the perturbed database, the transformed queries, the whole query processing procedure, the access patterns, and understand the same query returns the same set of results, but nothing else.

Adversaries can possibly have the global information of the database, such as the applications of the database, the attribute domains, and possibly the attribute distributions, via other published sources (e.g., the distribution of sales, or patient diseases, in public reports).

Protected assets: Data confidentiality and query privacy should be protected in the RASP approach. While the integrity of query services is also an important issue, it is orthogonal to our study. Existing integrity checking and preventing techniques [10], [11], [12] can be integrated into our framework. Thus, the integrity problem will be excluded from the paper, and we can assume the curious cloud provider is interested in the data and queries, but it will honestly follow the protocol to provide the infrastructure service. Attacker modeling. The goal of attack is to recover (or estimate) the original data from the perturbed data, or identify the exact queries (i.e., location queries) to breach users' privacy. According to the level of prior knowledge the attacker may have, we categorize the attacks into two categories:

- Level 1: The attacker knows only the perturbed data and transformed queries, without any other prior knowledge. This corresponds to the ciphertext-only attack in the cryptographic setting.
- Level 2: The attacker also knows the original data distributions, including individual attribute distributions and the joint distribution (e.g., the covariance matrix) between attributes. In practice, for some applications, whose statistics are interesting to the public domain, the dimensional distributions might have been published via other sources.

D. RASP: RAndom Space Perturbation

RASP denotes Random Space Perturbation. RASP is one type of multiplicative perturbation, with a novel combination of OPE, dimension expansion, random noise injection, and random projection.

Random projection is mainly used to process the high dimensional data into low dimensional data representations. It contains features like good scaling potential and good performances. Random noise injection is mainly used to adding noise to the input to get proper output when we compare it to the estimated power. The RASP method and its combination provide confidentiality of data and this approach is mainly used to protect the multidimensional range of queries in secure manner and also with indexing and efficient query processing will be done. RASP has some important features.

In RASP the use of matrix multiplication does not protect the dimensional values so no need to suffer from the distribution based attack. RASP prevents the data that are perturbed from distance based attacks; it does not protect the distances that are occurred between the records. And also it won't protect more difficult structures it may be a matrix and other components. The range queries can be send to the RASP perturbed data and this range query describes open bounds in the multidimensional space.

Algorithm 1 RASP Data Perturbation

1: **RASP Perturb**(X,RNG,RIMG,Ko)
2: Input: X: k × n data records, RNG: random real value generator that draws values from the standard normal distribution, RIMG : random invertible matrix generator, K_{ope}: key for OPE E_{ope}; Output: the matrix A
3: $A \leftarrow 0$;
4: $A_3 \leftarrow$ the last column of A;
5: $v_0 \leftarrow 4$;
6: **while** A_3 contains zero **do**

7: generate A with RIMG;
8: **end while**
9: **for** each record x in X **do**
10: $v \leftarrow v_0 - 1$;
11: **while** $v < v_0$ **do**
12: $v \leftarrow$ RNG;
13: **end while**
14: $y \leftarrow A((E_{ope}(x,K_{ope}))T , 1, v)T$;
15: submit y to the server;
16: **end for**
17: return A;

In random space perturbation, the word perturbation is used to do collapsing this process will happen according to the key value that is given by the owner. In this module the data owner have to register as owner and have to give owner name and key value. And then the user have register and get the key value and data owner name from the owner to do access in the cloud. Here user can submit their query as range query or kNN query and get their answer. We analyze and show the result with encrypted and also in decrypted format of the data for the query construct by the user.

Algorithm 2 RASP Secure Query Transformation.
1: **QuadraticQuery**(C_{ond},A)
2: Input: C_{ond}: 2d simple conditions for d-dimensional data, 2 conditions for each dimension. A:the perturbation matrix. Output: the MBR of the transformed range and the quadratic query matrices Qi, i = 1 . . . 2d.
3: $v_0 \leftarrow 4$;
4: **for** each condition C_i in C_{ond} **do**
5: $u \leftarrow$ zeros(d + 2, 1);
6: **if** C_i is like $X_j < a_j$ **then**
7: $u_j \leftarrow 1, u_{d+1} \leftarrow -a_j$;
8: **end if**
9: **if** C_i is like $X_j > a_j$ **then**
10: $u_j \leftarrow -1, u_{d+1} \leftarrow a_j$;
11: **end if**
12: $w \leftarrow$ zeros(d + 2, 1);
13: $w_{d+2} \leftarrow 1$;
14: $w_{d+1} \leftarrow v_0$;
15: $Q_i \leftarrow (A-1)Tuw\ TA-1$;
16: **end for**
17: Use the vertex transformation method to find the MBR of the transformed queries;
18: **return** MBR and $\{Q_i, i = 1 . . . 2d\}$;

RASP has several important features. First, RASP does not preserve the order of dimensional values because of the matrix multiplication component, which distinguishes itself from order preserving encryption schemes, and thus does not suffer from the distribution-based attack. Second, RASP does not preserve the distances between records, which prevents the perturbed data from distance based attacks. Because none of the transformations in the RASP: E_{ope}, G, and F preserves distances, apparently the RASP perturbation will not preserve distances. Third, the original range queries can be transformed to the RASP perturbed data space, which is the basis of our query processing strategy. A range query describes a hyper cubic area (with possibly open bounds) in the multidimensional space.

E. *kNN Query Processing with RASP*

The RASP perturbation does not preserve distances (and distance orders), kNN query cannot be directly processed with the RASP perturbed data. In this section, we design a kNN query processing algorithm based on range queries (the kNN-R algorithm). As a result, the use of index in range query processing also enables fast processing of kNN queries.

The original distance-based kNN query processing finds the nearest k points in the spherical range that is centered at the query point. The basic idea of our algorithm is to use square ranges, instead of spherical ranges, to find the approximate kNN results, so that the RASP range query service can be used. There are a number of key problems to make this work securely and efficiently. 1) How to efficiently find the minimum square range

that surely contains the k results, without many interactions between the cloud and the client? 2)Will this solution preserve data confidentiality and query privacy? 3) Will the proxy server's workload increase? to what extent ?

The algorithm is based on square ranges to approximately find the kNN candidates for a query point, which are defined as follows.

Definition 1 : "A square range is a hypercube that is centered at the query point and with equal-length edges."

Fig. 2 illustrates the range-query-based kNN processing with 2D data. The Inner Range is the square range that contains at least k points, and the Outer Range encloses the spherical range that encloses the inner range. The outer range surely contains the kNN results (see Proposition 2) but it may also contain irrelevant points that need to be filtered out.

Proposition 1 : "The kNN-R algorithm returns results with 100 percent recall."

Proof:

The sphere in [**Figure-2]** between the outer range and the inner range covers all points with distances less than the radius r. Because the inner range contains at least k points, there are at least k nearest neighbors to the query points with distances less than the radius r. Therefore, the k nearest neighbors must be in the outer range.

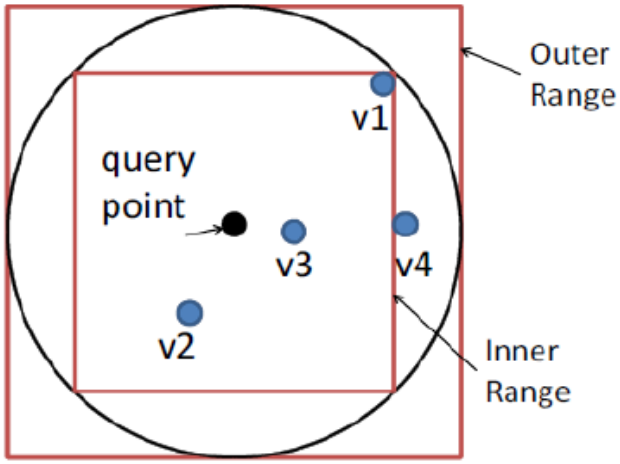

Figure: 2. Illustration for kNN-R Algorithm when k = 3.

The kNN-R algorithm consists of two rounds of interactions between the client and the server. Fig. 3 demonstrates the procedure. 1) The client will send the initial upper bound range, which contains more than k points, and the initial lower bound range, which contains less than k points, to the server. The server finds the inner range and returns to the client. 2) The client calculates the outer range based on the inner range and sends it back to the server. The server finds the records in the outer range and sends them to the client. 3) The client decrypts the records and find the top k candidates as the final result.

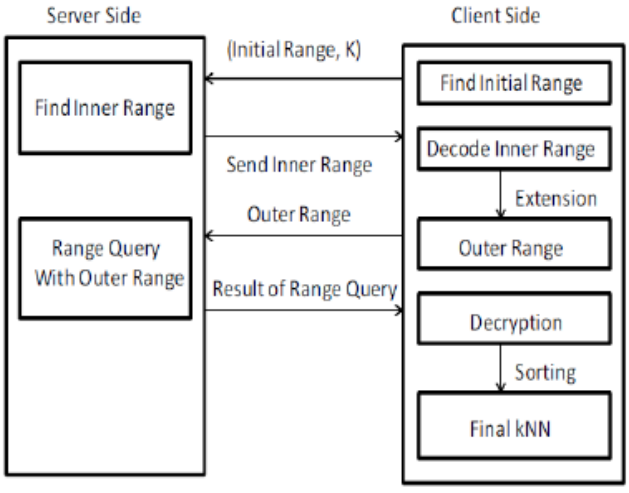

Figure: 3. Procedure of the kNN-R algorithm.

If the points are approximately uniformly distributed, we can estimate the precision of the returned result. With the uniform assumption, the number of points in an area is proportional to the size of the area. If the inner range contains m points, m > = k, the outer range contains q points, and the dimensionality is d, we can derive q = 2d=2m.

IV. CHALLENGES AND EXPECTED OUTCOME

Challenges and the expected outcome in creating the secured cloud environment is as follows:

A. Challenges

- Preserving data confidentiality: the main challenge is to preserve the privacy and confidentiality for the data.
- Should not provide slow query service as a result of security & privacy assurance: slow response should not be provided as a result of security as it increases the uses bandwidth.
- Bandwidth should be decreased.

B. Expected outcome

A secured query service: the expected outcome is to provide an environment that makes the data owner to feel that his/her data is secured in the cloud and bandwidth is reduced.

V. CONCLUSION

We propose the RASP perturbation approach to hosting query services in the cloud, which satisfies the CPEL criteria: data Confidentiality, query Privacy, Efficient query processing, and Low in-house workload. The requirement on low in-house workload is a critical feature to fully realize the benefits of cloud computing, and efficient query processing is a key measure of the quality of query services.

RASP perturbation is a unique composition of OPE, dimensionality expansion, random noise injection, and random projection, which provides unique security features. It aims to preserve the topology of the queried range in the perturbed space, and allows using indices for efficient range query processing. With the topology-preserving features, we are able to develop efficient range query services to achieve sub linear time complexity of processing queries. We then develop the kNN query service based on the range query service. The security of both the perturbed data and the protected queries is carefully analyzed under a precisely defined threat model. We also conduct several sets of experiments to show the efficiency of query processing and the low cost of in-house processing.

We will continue our studies on two aspects: (1) further improve the performance of query processing for both range queries and kNN queries; (2) formally analyze the leaked query and access patterns and the possible effect on both data and query confidentiality.

ACKNOWLEDGEMENT

I would like to express my sincere thanks to my guide Asst Prof. Ranjana.B. N for her motivation and useful suggestions which truly helped me in improving the quality of this paper. I take this opportunity to express my thanks to my teacher, family and friends for their encouragement and support.

REFERENCES

[1] R. Agrawal, J. Kiernan, R. Srikant, and Y. Xu, "Order preserving encryption for numeric data," in *Proceedings of ACM SIGMOD Conference*, 2004.

[2] Boneh, D., and Waters, B. "Conjunctive, subset, and range queries on encrypted data," In the Theory of Cryptography Conference (TCC (2007), Springer, pp. 535–554.

[3] Chen, K., and Liu, L. "VISTA: Validating and refining clusters via Visualization," Information Visualization 3, 4 (2004), 257–270.

[4] Curtmola, R., Garay, J., Kamara, S., and Ostrovsky, R. "Searchable symmetric encryption: improved definitions and efficient constructions," In ACM CCS (2006), pp. 79–88.

[5] Xu, H., Guo, S., and Chen, K. "Building confidential and efficient query services in the cloud with rasp data perturbation," IEEE Transactions on Knowledge and Data Engineering 26, 2 (2014).

[6] B. Hore, S. Mehrotra, and G. Tsudik, "A Privacy Preserving Index for Range Queries," Proc. Very Large Databases Conf. (VLDB), 2004.

[7] K. Chen, R. Kavuluru, and S. Guo, "RASP: Efficient Multidimensional Range Query on Attack-Resilient Encrypted Databases," Proc. ACM Conf. Data and Application Security and Privacy,pp. 249-260, 2011.

[8] B. Chor, E. Kushilevitz, O. Goldreich, and M. Sudan, "Private Information Retrieval," ACM Computer Survey, vol. 45, no. 6,pp. 965-981, 1998.
[9] S. Papadopoulos, S. Bakiras, and D. Papadias, "Nearest Neighbor Search with Strong Location Privacy," Proc. Very Large Databases Conf. (VLDB), 2010.
[10] M. Xie, H. Wang, J. Yin, and X. Meng, "Integrity Auditing of Outsourced Data," Proc. Very Large Databases Conf. (VLDB), pp. 782-793, 2007.
[11] R. Sion, "Query Execution Assurance for Outsourced Databases," Proc. Very Large Databases Conf. (VLDB), 2005.
[12] F. Li, M. Hadjieleftheriou, G. Kollios, and L. Reyzin, "Dynamic Authenticated Index Structures for Outsourced Databases," Proc. ACM SIGMOD Int'l Conf. Management of Data (SIGMOD), 2006.

