International Journal For Advanced Research In Science & Technology
A peer reviewed international journal
www.ijarst.in
ISSN: 2457-0362

DiploCloud: Efficient and Scalable Management Of RDF Data in the Cloud

GEDDADA SAI TEJASRI[1], K RAJA RAJESWARI [2]

[1]MCA Student, B V Raju College, Kovvada, Andhra Pradesh, India.
[2]Assistant Professor, B V Raju College, Kovvada, Andhra Pradesh, India.

ABSTRACT:

Despite recent advances in distributed RDF data management, processing large-amounts of RDF data in the cloud is still very challenging. In spite of its seemingly simple data model, RDF actually encodes rich and complex graphs mixing both instance and schema-level data. Sharing such data using classical techniques or partitioning the graph using traditional min-cut algorithms leads to very inefficient distributed operations and to a high number of joins. In this paper, we describe DiploCloud, an efficient and scalable distributed RDF data management system for the cloud. Contrary to previous approaches, DiploCloud runs a physiological analysis of both instance and schema information prior to partitioning the data. In this paper, we describe the architecture of DiploCloud, its main data structures, as well as the new algorithms we use to partition and distribute data. We also present an extensive evaluation of DiploCloud showing that our system is often two orders of magnitude faster than state-of-the-art systems on standard workloads.

Keywords: RDF, Diplocloud, cloud, data encryption

I INTRODUCTION

In the era of big data, a huge amount of data can be generated quickly from various sources (e.g., smart phones, sensors, machines, social networks, etc.). Towards these big data, conventional computer systems are not competent to store and process these data. Due to the flexible and elastic computing resources, cloud computing is a natural fit for storing and processing big data. With cloud computing, end-users store their data into the cloud, and rely on the cloud server to share their data to other users (data consumers). In order to only share end-users' data to authorized users, it is necessary to design access control mechanisms according to the requirements of end-users. When outsourcing data into the cloud, end-users lose the physical control of their data. Moreover, cloud service providers are not fully-trusted by end-users, which make the access control more challenging. For example, if the traditional access control mechanisms

(e.g., Access Control Lists) are applied, the cloud server becomes the judge to evaluate the access policy and make access decision. Thus, end-users may worry that the cloud server may make wrong access decision intentionally or unintentionally, and disclose their data to some unauthorized users. In order to enable endusers to control the access of their own data, some attributebased access control schemes are proposed by leveraging attribute-based encryption. In attribute-based access control, end-users first define access policies for their data and encrypt the data under these access policies. Only the users whose attributes can satisfy the access policy are eligible to decrypt the data. Although the existing attribute-based access control schemes can deal with the attribute revocation problem, they all suffer from one problem: the access policy may leak privacy. This is because the access policy is associated with the encrypted data in plaintext form. From the plaintext of access policy, the adversaries may obtain some privacy information about the end-user. For example, Alice encrypts her data to enable the "Psychology Doctor" to access. So, the access policy may contain the attributes "Psychology" and "Doctor". If anyone sees this data, although he/she may not be able to decrypt the data, he/she still can guess that Alice may suffer from some psychological problems, which leaks the privacy of Alice. To prevent the privacy leakage from the access policy, a straightforward method is to hide the attributes in the access policy. However, when the attributes are hidden, not only the unauthorized users but also the authorized users cannot know which attributes are involved in the access policy, which makes the decryption a challenging problem. Due to this reason, existing methods do not hide or anonymize the attributes. Instead, they only hide the values of each attribute by using wildcards, Hidden Vector Encryption, and Inner Product Encryption. Hiding the values of attributes can somehow protect user privacy, but the attribute name may also leak private information. Moreover, most of these partially hidden policy schemes only support specific policy structures (e.g., AND-gates on multi-valued attributes). In this paper, we aim to hide the whole attribute instead of only partially hiding the attribute values. Moreover, we do not restrict our method to some specific access structures. The basic idea is to express the access policy in LSSS access structure (M;r) where M is a policy matrix and r matches each row Mi of the matrix M to an attribute

International Journal For Advanced Research In Science & Technology
A peer reviewed international journal
www.ijarst.in
ISSN: 2457-0362

[6], and hide the attributes by simply removing the attribute matching function r. Without the attribute matching function r, it is necessary to design an attribute localization algorithm to evaluate whether an attribute is in the access policy and if so find the correct position in the access policy. To this end, we further build a novel Attribute Bloom Filter to locate the attributes to the anonymous access policy, which can save a lot of storage overhead and computation cost especially for large attribute universe. Our contributions are summarized as follows.

1) We propose an efficient and fine-gained big data access control scheme with privacy-preserving policy, where the whole attributes are hidden in the access policy rather than only the values of the attributes.

2) We also design a novel Attribute Bloom Filter to evaluate whether an attribute is in the access policy and locate the exact position in the access policy if it is in the access policy.

3) We further give the security proof and performance evaluation of our proposed scheme, which demonstrate that our scheme can preserve the privacy from any LSSS access policy without employing much overhead.

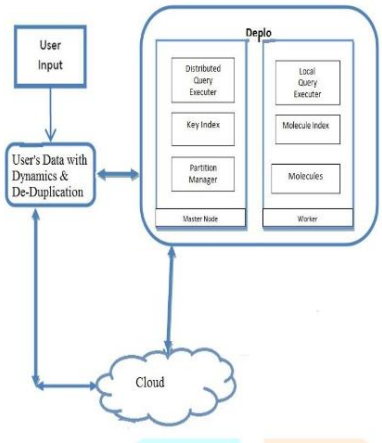

II SURVEY OF RESEARCH

1) Tracking RDF graph provenance using RDF molecules This paper investigates lossless decomposition of RDF graph and tracking the provenance of RDF graph using RDF molecule, which is the finest and lossless component of an RDF graph. A sub-graph is {em lossless} if it can be used to restore the original graph without introducing new triples. A sub-graph is {em finest} if it cannot be further decomposed into lossless sub-graphs.

2) DOGMA: A Disk-Oriented Graph Matching Algorithm for RDF Databases In this paper, we first propose the DOGMA index for fast subgraph matching on disk and then develop a basic algorithm to answer queries over this index. This algorithm is then significantly sped up via an optimized algorithm that uses efficient (but correct) pruning strategies when combined with two different extensions of the index.

International Journal For Advanced Research In Science & Technology
A peer reviewed international journal www.ijarst.in
ISSN: 2457-0362

3) The design and implementation of a clustered RDF store This paper describes the design and performance characteristics of 4store, as well as discussing some of the trade-offs and design decisions. These arose both from immediate business requirements and a desire to engineer a scalable system capable of reuse in a range of experimental contexts where we were looking to explore new business opportunities.

4) WARP: Workload-aware replication and partitioning for RDF This paper proposes a distributed SPARQL engine that combines a graph partitioning technique with workload-aware replication of triples across partitions, enabling efficient query execution even for complex queries from the workload.

5) Scaling queries over big RDF graphs with semantic hash partitioning In this paper we present a novel semantic hash partitioning approach and implement a Semantic HAsh PartitioningEnabled distributed RDF data management system, called Shape.

III WORKING METHODOLOGY

Diplo Cloud: We say that DiploCloud is a hybrid system. DiploCloud is a native, RDF database system. It was designed to run on clusters of commodity machines in order to scale out gracefully when handling bigger RDF file. Our system design follows the architecture of many modern cloudbased distributed systems. Where one (Master) node is responsible for interacting with the clients and orchestrating the operations performed by the other (Worker) nodes.

Master: The Master node is composed of three main subcomponents: a key index in charge of encoding URIs and literals into compact system identifiers and of translating them back, a partition manager responsible for the partitioning the RDF data and a distributed query executor, responsible for parsing the incoming query, rewriting the query plans into the Workers.

Worker: The Worker nodes hold the partitioned data and its corresponding local indices, and are responsible for running sub-queries and sending results back to the Master node. Conceptually, the Workers are much simpler than the Master node and are built on three main data structures: i) a type index, clustering all keys based on their types ii) a series of RDF molecules, storing RDF data as very compact subgraphs, and iii)

International Journal For Advanced Research In Science & Technology
A peer reviewed international journal
www.ijarst.in
ISSN: 2457-0362

a molecule index, storing for each key the list of molecules where the key can be found.

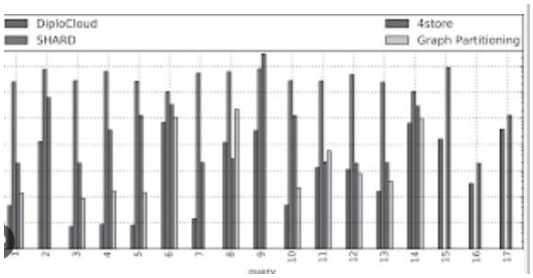

CONCLUSION

We implementing a novel search model that effectively handles the parallel data processing system and also we implements a novel architecture that will handles the distributed partitions. A novel data handle mechanism will collect similar relevance from cloud server. New data loading technique and query process analysis would be advantage of present scenario's data partitions. Our current efficient evaluation provide to our current scenario is offers two orders of magnitude faster than stat-ofthe-art systems. We further extended the system by adding data dynamics (update, delete) and data deduplication to the system.

REFERANCES

[1] K. Aberer, P. Cudre-Mauroux, M. Hauswirth, and T. van Pelt, "GridVine: Building Internet-scale semantic overlay networks," in Proc. Int. Semantic Web Conf., 2004, pp. 107–121.

[2] P. Cudre-Mauroux, S. Agarwal, and K. Aberer, "GridVine: An infrastructure for peer information management," IEEE Internet Comput., vol. 11, no. 5, pp. 36–44, Sep./Oct. 2007.

[3] M. Wylot, J. Pont, M. Wisniewski, and P. CudreMauroux. (2011). dipLODocus[RDF]: Short and long-tail RDF analytics for massive webs of data. Proc. 10th Int. Conf. Semantic Web Vol. Part I,pp. 778–793 [Online]. Available: http://dl.acm.org/citation.cfm?id=2063016.2063066

[4] M. Wylot, P. Cudre-Mauroux, and P. Groth, "TripleProv: Efficient processing of lineage queries in a native RDF store," in Proc. 23rd Int. Conf. World Wide Web, 2014, pp. 455–466.

[5] M. Wylot, P. Cudre-Mauroux, and P. Groth, "Executing provenance- enabled queries over web data," in Proc. 24th Int. Conf.World Wide Web, 2015, pp. 1275–1285.

[6] B. Haslhofer, E. M. Roochi, B. Schandl, and S. Zander. (2011). Europeana RDF store report. Univ. Vienna, Wien, Austria, Tech.Rep. [Online]. Available: http://eprints.cs.univie.ac.at/2833/1/europeana_ts _report.p df

[7] Y. Guo, Z. Pan, and J. Heflin, "An evaluation of knowledge base systems for large OWL datasets," in Proc. Int. Semantic Web Conf.,2004, pp. 274–288.

[8] Faye, O. Cure, and Blin, “A survey of RDF storage approaches,”ARIMA J., vol. 15, pp. 11–35, 2012.

[9] B. Liu and B. Hu, “An Evaluation of RDF Storage Systems for Large Data Applications,” in Proc. 1st Int. Conf. Semantics, Knowl.Grid, Nov. 2005, p. 59.

[10] Z. Kaoudi and I. Manolescu, “RDF in the clouds: A survey,” VLDB J. Int. J. Very Large Data Bases, vol. 24, no. 1, pp. 67–91, 2015.