

Storage Balancing in P2P Based Distributed RDF Data Stores

Maximiliano Osorio and Carlos Buil-Aranda

Universidad Técnica Federico Santa María, Valparaíso, Chile.
{mosorio,cbuil}@inf.utfsm.cl

Abstract. Centralized RDF repositories have been designed to support RDF data storage and retrieval. However, they suffer from the traditional limitations of centralized approaches which are scalability and fault tolerance. Peer to Peer (P2P) networks can provide the scalability, fault-tolerance and robustness, features that the current solutions to local RDF storage do not provide which are needed by the existing Semantic Web applications. A common strategy from state-of-the-art P2P-RDF data stores is to store triples at three locations so each triple can be found using a look-up by subject, predicate, or object identifier. One major issue of this strategy is the lack of load-balancing, since occurrences in triples are not uniformly distributed. Consequently, this issue leads an unbalance query processing load distribution and unfair storage load in the network. To solve this problem caused by load imbalance, we propose new scheme to split the data in the stressed nodes which is based in evenly distributing excess of data across neighboring nodes providing a Prefix Hash Table for fast accessing to such data. We provide an empirical evaluation of our novel approach and compare with other state of the art systems for storage balancing showing the feasibility of our approach.

1 Introduction

Centralized RDF repositories and look-up systems have been designed to support RDF data storage and retrieval. These systems normally suffer from the traditional limitations of centralized approaches which are scalability and fault tolerance [1].

As an alternative to these centralized systems, P2P approaches have been proposed to overcome some of these limitations by building (fully) decentralized data storage and retrieval systems [2]. P2P networks and especially distributed hash tables (DHTs) have gained much attention recently, the reasons are the scalability, fault-tolerance and robustness features the can provide to Internet applications. However, it is crucial to the system that the data could be well distributed over the P2P network, otherwise, the effects are an unfair distribution of the data causing a high cost of computation and transmission. In this paper, we propose a structure for indexing triples using Distributed Hash Table and Prefix Hash Tree. Our approach is based on evenly distributing the excess

of data across neighboring nodes based on a dynamic changing threshold. For maintaining fast access to the data stored we propose to use a distributed data structure called Prefix Hash Tree (PHT). In PHT each node in the trie has a label with a prefix that is defined recursively to allow fast access to the nodes in the network containing the related data. The rest of the paper is organized as follows: in Section 2 we describe the existing state of the art related to RDF storage in P2P networks. Next in Section 3 we present our solution to the problems identified in Section 2. We evaluate our approach in Section 4 and finally we present our conclusions in Section 5.

2 Related work

The Resource Description Framework (RDF) [3] is the recommended data model by the W3C aiming to improve the World Wide Web with machine-processable semantic data for data interchange on the Web. The notion of RDF triple is the basic building block of the RDF model. It consists of a subject (s), a predicate (p) and an object (o). More precisely, given a set of IRI [4] references I , a set of blank nodes B , and a set of literals L , a triple $(s, p, o) \in (I \cup B) \times I \times (I \cup B \cup L)$. The subject of a triple denotes the resource that the statement is about, the predicate denotes a property of the subject and the object presents the value of the property.

Centralized RDF repositories and lookup systems such as Jena¹, RDFDB² or Virtuoso³ have been designed to support RDF data storage and retrieval. Although these systems are highly optimized RDF stores, they suffer from the traditional limitations of centralized approaches such as scalability and fault-tolerance [1]. As an alternative to these centralized systems, P2P approaches [5,6,7] have been proposed to overcome some of these limitations by building decentralized data storage and retrieval systems [2].

P2P networks and especially distributed hash tables (DHTs [8], hash tables in which the responsibility for maintaining the mapping from keys to values is distributed among the nodes in the network) have gained much attention recently [2]. The reasons are the scalability, fault-tolerance and robustness features key features for most Internet and Web applications. We say that a responsible node is the node which stores the key for the RDF triple. Some RDF systems use DHTs to store and query RDF data at Internet scale like RDFPeers [9], Atlas [10] or GridVine [11]. These solutions index triples 3 times for each component of triple, these are regularly disseminated to the nodes in the network by calculating the hash function of the subjects, predicates, and objects and sending the triples to the nodes responsible received. This indexing technique provides the possibility to find triples based on any search criteria as long there exist at least one constant in a triple pattern [12]. Even though, the frequency of the triples usually is not uniformly distributed, thus, it is possible that the

¹ <http://jena.apache.org/documentation/rdf/>

² <https://github.com/scor/rdfdb>

³ <https://virtuoso.openlinksw.com/>

responsible peer will be heavily loaded. And if a peer is heavily loaded it can affect the system's behavior, causing a high cost of computation and transmission, consequently, it is crucial to solve the problem of triples distribution. In [9] the authors approach the storage balancing problem by simply not indexing the most common RDF triples such those having an `rdf:type` predicate and the requesting node must then find an alternative way of resolving the query and ask to another target node. The result of such an approach is cost of possibly losing the complete result.

To solve this problem authors in [13] propose the use of four RDF databases in each peer of the network (local triples, received triples, generated triples, and replica triples database). The local triples database stores RDF triples that originate from the particular node, the received triples database stores all local triples, these are regularly disseminated to the nodes in the network by calculating the hash function of the subjects, predicates, and objects and sending the triples to the nodes responsible received. Also, each node hosts a database for generated triples that originate from forward chaining. Finally, the replica database has a copy data to the k whose IDs that are nearest to the target hash value determined by the hash function. The authors addressed the issue of load-balancing to build an overlay tree over a DHT, if a node detects that it is overloaded, the node performs a split operation, half of the triples remained in the local part of the remote triples database and half of them were moved to the new node. This approach allows an easier access to the data, however the overhead generated by the five databases in each node of the network and the use of a forward chaining approach for RDF(S) reasoning results in higher storage and bandwidth costs for a single peer [14].

Atlas [7] is a distributed RDF system that uses DHTs for distributing RDF data across the peers. For storing an RDF triple Atlas sends three DHT put requests using as key the subject, property and object together, and the triple itself as an item. The key is hashed to create the identifier that leads to the responsible node where the triple is stored. Atlas also suffers from load imbalances [14], due to RDF triples some triple components are more frequent, for instance `rdf:type` and `rdfs:label`, the peer responsible for such a key store more triples and the built-in load balancing is not able to balance this higher load.

In [12] the authors propose an indexing scheme "3-tuple index" where 3 combined routing indexes are created on triples subject, predicate and object components. The output combinations are subject+predicate, predicate+object and object+subject.

In summary, for storing and accessing RDF triples in a P2P network the main approach is to use DHTs. However one critical problem is not correctly balancing the amount of data stored in each node, leading to extra processing by some of the nodes, and removing some of the advantages of the P2P framework. We now proceed to describe our solution to such problem.

3 System model and data model

An obvious approach to solve the node saturation problem described before is to add a maximum amount of triples that a node can store. Once a node reaches that maximum size, the node should split the data into two child nodes, and these keys should be redistributed among the children equally. After that, each child has a half of triples and each half of the triples can be split again if a node reaches again the maximum number of triples. However the problem now is to find the triples in the tree since a triple can stay in any node of such tree. To fix such problem, we propose to use a distributed data structure called Prefix Hash Tree (PHT), in PHT each node in the trie has a label with a prefix that is defined recursively to allow fast access to the nodes in the network containing the related data. Consider the node l in which l_0 and l_1 are the left hand and right hand side nodes. For each two children we have the following properties:

1. Each node has either 0 or 2 children.
2. A key K is stored at a leaf node whose label is a prefix K .
3. Each leaf node stores at most B keys.
4. Each internal node contains at least $(B + 1)$ keys in its sub-tree.

The PHT structure is a routine binary trie where each node of the trie is a node of the DHT ring. Atlas uses a mapping dictionary, and since URIs and literals may consist of long strings, they are mapped to integer values and then, triple storage, and query evaluation is performed using these integer values. [7]. Using the previous idea, the domain for the indexing is $\{0, 1\}^D$, thus, the triples are stored according to their integer value.

Notice that this problem cannot be solved by just redistributing again the triples across the network since the hash function for distributing the data would return similar keys all the time and some nodes would be still overloaded. Besides generating more complexity at search time.

3.1 Operations

Lookup Given a part of the triple K , PHT lookup returns a unique leaf node: $leaf(K)$. The algorithm uses a binary search: if the current prefix is an internal node, the search tries a shorter prefix, and if the current prefix is not an internal node, the search tries a longer prefix.

Query Forwarding queries within the tree does not consume expensive DHT routing but can be done via direct communication. A query reaches its destination in $O(\log N + d)$ steps, where N denotes the number of nodes in the DHT network and d denotes the depth of the tree. The

Insert/Delete Insertion is a common operation in the system. When a new triple has to be inserted, the node calculates the hash function of the subject, predicates, and objects and sending the triples to the nodes responsible. These two

cases when the triple is received: the new amount of triples in the node does not reach the maximum amount of triples B or not. In the first case, the node stores the triple, otherwise, the node performs `lookup` to find the new responsible node and saved the triple in it if the triples do not reach the Maximum value amount of triples B . In both cases, the integer value of the mapping dictionary is saved in an ATLAS local database. Similarly, deletion can cause that a subtree to collapse into a single leaf node.

4 Performance Analysis

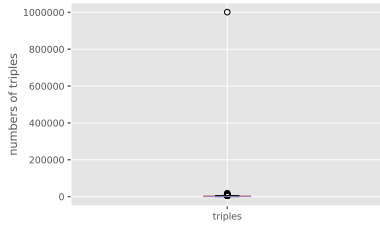
We compared our algorithm for storage balancing with the Atlas algorithm and with the approach in [12]. To do that we simulated a P2P with 1,000 nodes and distributed among these nodes 1,000,000 triples from DBpedia. We evaluated how well the data is distributed across these 1,000 nodes, we mark as future work to evaluate the query execution of a complete SPARQL benchmark. The source code for our evaluation and the data used can be found in http://inf.utfsm.cl/~mosorio/p2p_rdf_balance.

Storage balancing We analyze the effect of data indexing algorithms from Atlas and [12]. The goal is evaluate the distribution's quality of these solutions. In the figure 1 we present the results for each algorithms indexing 1 million triples from DBpedia, Figures 1a,1b show the results for Atlas's algorithm. In Figure 1a there is an outlier point which corresponds to the predicate `rdf:type`. The figure 1b has the same results but we remove the node responsible for `rdf:type` but we can see multiple outliers.


On the other hand, [12] proposes a mixed approach, 1c shows the results, in comparison with Atlas's results we see an improvement in the distribution but we see multiples outliers. On the other hand, the figure 1d shows the results for the proposed algorithms using the maximum number of triples by node $B = 1000$, in comparison with the last two algorithms, we can see that the triples are well distributed and there are not outliers in the results.

The proposed algorithm has an important parameter B which determines when the data redistribution should happen (which can be adjusted automatically), so it is crucial to evaluate the algorithm with different configurations, the figures 1d,1e,1f show the results using different amount of triples and a different values for B .


Another important aspect is the depth of tree using different values for B . Table 1 shows the results; we see the value are low. Moreover, the depth of tree decreases when the B increases. If the number of triples into the node increases over time, the parameter B can be adjusted dynamically becoming an adaptive process.




(a) Algorithm used by Atlas: we see a node which is storing more than 1M triples containing the `rdf:type` predicate.




(b) If we remove the previous triple with the `rdf:type` predicate we see a more balanced distribution of data, however there are several nodes which store large amounts of triples (10,000+ triples).




(c) Using the 3-tuple index algorithm [12] we observe a data distribution similar to 1b



(d) Using our algorithm and using a distribution with a maximum of $B=1000$ nodes we do not see any node storing large amounts of triples.



(e) Using our algorithm and using a distribution with a maximum of $B=5000$ nodes we see a few node storing medium amounts of triples.



(f) Using our algorithm and using a distribution with a maximum of $B=10000$ nodes we observe a data distribution similar to 1e.

Fig. 1: Indexing 1 million triples from DBpedia

B	Number of triples		
	100.000 triples	1 million triples	10 million triples
1000	2	6	10
5000	2	2	6
10000	1	2	5

Table 1: Tree’s depth using different numbers of triples and B

5 Conclusions

In this paper, we discussed indexing techniques for RDF in P2P networks, the proposed algorithm solved load imbalance of DHT for RDF stores. And we showed by using our new index scheme we can achieve a better storage distribution than the techniques from literature. Also, we showed by using the proposed algorithm with different configurations we can achieve a fair stability. Finally, we showed in the results that the tree’s depth is fair and considering that the worst case message complexity of this solution is $O(\log(N) + D)$ we can conclude that new scheme doesn’t increase strongly the time complexity. Summarizing, the proposed algorithm limits the maximum number of triples by each node, if the node reaches the maximum number the node performs a split operation. Also, using PHT, a node can query and find the triple in the tree. The time overhead by using PHT and the maximum number is low: $O(\log(N) + D)$

6 References

1. C. Buil-Aranda, A. Hogan, J. Umbrich, and P.-Y. Vandenbussche, “SPARQL Web-Querying Infrastructure: Ready for Action?,” in *ISWC2013*, pp. 277–293, 2013.
2. I. Filali, F. Bongiovanni, F. Huet, and F. Baude, “A survey of structured p2p systems for rdf data storage and retrieval,” in *Transactions on large-scale data- and knowledge-centered systems iii*, pp. 20–55, Springer, 2011.
3. “Rdf - semantic web standards.” <https://www.w3.org/RDF/>. (Accessed on 06/20/2017).
4. M. Dürst and M. Suignard, “Internationalized resource identifiers (iris),” tech. rep., 2004.
5. W. Nejdl, B. Wolf, C. Qu, S. Decker, M. Sintek, A. Naeve, M. Nilsson, M. Palmér, and T. Risch, “Edutella: a p2p networking infrastructure based on rdf,” in *Proceedings of the 11th international conference on World Wide Web*, pp. 604–615, ACM, 2002.
6. E. Della Valle, A. Turati, and A. Ghioni, “Page: A distributed infrastructure for fostering rdf-based interoperability,” in *DAIS*, vol. 6, pp. 347–353, Springer, 2006.
7. Z. Kaoudi, I. Miliaraki, and M. Koubarakis, “RDFS reasoning and query answering on top of DHTs,” *The Semantic Web-ISWC 2008*, pp. 499–516, 2008.
8. H. Balakrishnan, M. F. Kaashoek, D. Karger, R. Morris, and I. Stoica, “Looking up data in p2p systems,” *Communications of the ACM*, vol. 46, no. 2, pp. 43–48, 2003.
9. M. Cai and M. Frank, “Rdfpeers: a scalable distributed rdf repository based on a structured peer-to-peer network,” in *Proceedings of the 13th international conference on World Wide Web*, pp. 650–657, ACM, 2004.

10. Z. Kaoudi, M. Koubarakis, K. Kyzirakos, I. Miliaraki, M. Magiridou, and A. Papadakis-Pesaresi, "Atlas: Storing, updating and querying rdf (s) data on top of dhts," *Web Semantics: Science, Services and Agents on the World Wide Web*, vol. 8, no. 4, pp. 271–277, 2010.
11. K. Aberer, P. Cudré-Mauroux, M. Hauswirth, and T. Van Pelt, *GridVine: Building Internet-Scale Semantic Overlay Networks*, pp. 107–121. Berlin, Heidelberg: Springer Berlin Heidelberg, 2004.
12. L. Ali, T. Janson, and C. Schindelhauer, "Towards load balancing and parallelizing of rdf query processing in p2p based distributed rdf data stores," in *Parallel, Distributed and Network-Based Processing (PDP), 2014 22nd Euromicro International Conference on*, pp. 307–311, IEEE, 2014.
13. D. Battré, F. Heine, A. Hoing, and O. Kao, "Load-balancing in p2p based rdf stores," in *2nd Workshop on Scalable Semantic Web Knowledge Base System*, 2006.
14. Z. Kaoudi, "Distributed rdf query processing and reasoning in peer-to-peer networks," *Department of Informatics and Telecommunications, National and Kapodistrian University of Athens, Athens*, pp. 59–88, 2011.