

Towards a Rule Based Distributed OWL Reasoning Framework

Raghava Mutharaju, Prabhaker Mateti, and Pascal Hitzler

Wright State University, OH, USA.

{mutharaju.2, prabhaker.mateti, pascal.hitzler}@wright.edu

Abstract. The amount of data exposed in the form of RDF and OWL continues to increase exponentially. Some approaches have already been proposed for the scalable reasoning over several language profiles such as RDFS, OWL Horst, OWL 2 EL, OWL 2 RL etc. But all those approaches are limited to the particular ruleset that the reasoner supports. In this work, we propose a rule-based distributed reasoning framework that can support any given ruleset.

1 Introduction

The W3C recommendations RDF and OWL, are primarily used to represent data in the Semantic Web. Large amount of data in these formats is now available and it only continues to grow. Several billions of RDF triples are available as Linked Open Data (close to 90 billion¹). Automated generation of OWL axioms from streaming data [11] and text [12,13] can result in very large knowledge bases. Reasoning is one of the most important operations that can be performed over OWL and RDF knowledge bases. It is required to infer logical consequences and to check the consistency of the knowledge base. Reasoning is memory and compute intensive. So reasoning over large knowledge bases needs a scalable approach. Currently, all reasoners are single machine reasoners. It is not possible for a single machine to keep up with the growth rate of data. Also, for some reasoning tasks the output is several times larger than the input. Distributed memory reasoning provides a viable alternative.

There are some existing approaches for scalable reasoning over each individual Semantic Web language profile such as RDFS, OWL Horst, OWL 2 EL, OWL 2 RL (see Section 4). Reasoning over ontologies in each of these profiles is performed using a set of rules that vary with each profile (there is some overlap among the different rulesets). Generally, the existing solutions are tuned towards a particular ruleset. In this paper, we propose the idea for a unified distributed reasoning framework that can work on any given ruleset. This framework can not only handle the aforementioned language profiles but also avoids the need to develop a customized scalable approach for any new ruleset.

¹ <http://stats.lod2.eu/>

1: $s p o$ (if o is literal)	\Rightarrow $..n$ rdf:type rdfs:Literal
2: p rdfs:domain x & $s p o$	\Rightarrow s rdf:type x
3: p rdfs:range x & $s p o$	\Rightarrow o rdf:type x
4a: $s p o$	\Rightarrow s rdf:type rdfs:Resource
4b: $s p o$	\Rightarrow o rdf:type rdfs:Resource
5: p $\text{rdfs:subPropertyOf}$ q & q $\text{rdfs:subPropertyOf}$ r	\Rightarrow p $\text{rdfs:subPropertyOf}$ r
6: p rdf:type rdf:Property	\Rightarrow p $\text{rdfs:subPropertyOf}$ p
7: $s p o$ & p $\text{rdfs:subPropertyOf}$ q	\Rightarrow $s q o$
8: s rdf:type rdfs:Class	\Rightarrow s rdfs:subClassOf rdfs:Resource
9: s rdf:type x & x rdfs:subClassOf y	\Rightarrow s rdf:type y
10: s rdf:type rdfs:Class	\Rightarrow s rdfs:subClassOf s
11: x rdfs:subClassOf y & y rdfs:subClassOf z	\Rightarrow x rdfs:subClassOf z
12: p rdf:type $\text{rdfs:ContainerMembershipProperty}$	\Rightarrow p $\text{rdfs:subPropertyOf}$ rdfs:member
13: o rdf:type rdfs:Datatype	\Rightarrow o rdfs:subClassOf rdfs:Literal

Table 1. RDFS closure rules

2 Challenges

Some of the challenges in the design and implementation of a rule-based distributed reasoning framework are discussed here.

2.1 Rule dependency analysis

If the input to a rule R_i , depends on the output of another rule R_j , then the rule R_i is dependent on R_j . The more independent the rules are, better can be the rule distribution among the nodes in the cluster.

A rule dependency graph can be constructed in order to determine the inter-dependency among the rules. Each vertex represents a rule and an outgoing edge between vertex v_i and v_j represents the dependency of vertex v_i on v_j . Isolated vertices are independent of each other and can be executed in parallel. For each vertex v_i , all the vertices that are reachable from it are dependent on each other.

RDFS rules from [24] are shown in Table 1 and its dependency graph is shown in Figure 1. There are several rules that are independent and can be given to separate nodes in the cluster. Rules 5, 6 and 7 are dependent on each other and can be grouped together i.e., all the three rules can be executed by one node. Same holds for the group of rules 9, 2, 3, 10, 11.

On the other hand, rules for \mathcal{EL}^{++} [1] are highly inter-dependent and are difficult to parallelize.

2.2 Rule implementation

Interpreting different rulesets will be very difficult for the framework. Instead, rulesets should be converted to a common domain specific language (DSL) that is supported by the framework. This DSL should be able to define the vocabulary, syntax and semantics of the language to be used.




Fig. 1. RDFS rule dependency graph

For the choice of DSL, there are some options. 1) general purpose rule languages such as RETE, Datalog and Prolog. 2) or a custom DSL for the rules supported by the framework. DSL should support the declaration of variables and constants in the rules.

3 Evaluation Plan

The rule-based distributed reasoning framework can be evaluated along the lines of adaptation and extension.

- There are several existing specialized and scalable reasoners for rulesets such as RDFS (WebPIE, Cichlid), OWL Horst (QueryPIE) and OWL 2 EL (DistEL). The framework should be able to handle these rulesets. The performance of the general purpose framework in comparison to the specialized ones remains to be seen.
- The framework should be able to take in a new ruleset and provide sound and complete inference over the given data.

4 Related Work

There are several language profiles in the Semantic Web that support rule-based reasoning. Other reasoning approaches such as tableau algorithms are not considered here. Scalable reasoning approaches such as parallel shared memory reasoning to distributed shared-nothing reasoning exist.

RDFS has around 34 inference (entailment) rules including simple, extensional and datatype entailment rules [4]. Almost all of the existing work on scalable RDFS entailment (closure) computation considers only a subset of these rules. Marvin computes the closure of RDF triples using a peer-to-peer model [18]. In [25], triples are distributed across the cluster by making a distinction between the schema and instance triples. Finding closure becomes an embarrassingly parallel computation. Cray XMT, a massively parallel shared memory supercomputer was used in [2] for RDFS inference, dictionary encoding and SPARQL query processing. Several other scalable approaches exist for RDFS closure computation [5,21,10,7].

OWL Horst (also known as pD^*) [6] extends RDFS entailment rules to include reasoning with datatypes from a given datatype map D . Rule partitioning and data partitioning strategies are explored in [22] for computing closure over OWL Horst knowledge base. QueryPIE [23] is a backward chaining distributed reasoner that supports OWL Horst reasoning over large knowledge bases. A recent Apache Spark implementation of RDFS and OWL Horst rules named Cichlid, is 10 times faster than state-of-the-art implementations [3].

OWL 2 RL [14] is further extension of pD^* and has several entailment rules. A scalable OWL 2 RL inference engine has been implemented inside a relational database system (Oracle) in [9]. QueryPIE has partial support for OWL 2 RL.

The description logic underlying OWL 2 EL is \mathcal{EL}^{++} and there are 11 completion rules ([1]) for classification (computing the subsumption hierarchy of all concepts). Three approaches to distributed \mathcal{EL}^{++} reasoning were discussed in [16]. Among them, the most efficient one named DistEL, uses a peer-to-peer model that uses rule partitioning based on the axiom types [17]. Though not distributed, parallelization of OWL 2 EL classification has been studied in [8,20].

There are some existing generic rule-based scalable reasoning approaches. RETE implementation on GPUs for RDFS and OWL Horst rulesets is shown in [19]. Another alternative is to convert different rulesets into datalog rules. A parallel implementation of datalog programs with application to RDFS rules is shown in [15]. A distributed approach for either of these two has not been developed yet.

5 Conclusion

There are different rulesets for different language profiles and there are scalable approaches for many of these rulesets. However, such a specialized scalable reasoner does not work on other rulesets. A general purpose rule-based reasoning framework is proposed here to fill that gap.

The next step is to choose the appropriate DSL by checking the advantages and disadvantages of RETE, Datalog and Prolog. After this, we can proceed with the implementation of the rule dependency graph as well as other aspects of the framework.

References

1. Baader, F., Brandt, S., Lutz, C.: Pushing the EL Envelope. In: Kaelbling, L.P., Saffiotti, A. (eds.) IJCAI-05, Proceedings of the Nineteenth International Joint Conference on Artificial Intelligence, Edinburgh, Scotland, UK, July 30-August 5, 2005. pp. 364–369. AAAI (2005)
2. Goodman, E.L., Jimenez, E., Mizell, D., al Saffar, S., Adolf, B., Haglin, D.: High-performance Computing Applied to Semantic Databases. In: Proceedings of the 8th Extended Semantic Web Conference on The Semantic Web: Research and Applications - Volume Part II. pp. 31–45. ESWC'11, Springer-Verlag, Berlin, Heidelberg (2011)

3. Gu, R., Wang, S., Wang, F., Yuan, C., Huang, Y.: Cichlid: Efficient Large Scale RDFS/OWL Reasoning with Spark. In: 2015 IEEE International Parallel and Distributed Processing Symposium, IPDPS 2015, Hyderabad, India, May 25-29, 2015. pp. 700–709. IEEE Computer Society (2015)
4. Hayes, P., Patel-Schneider, P.F.: RDF Semantics. <http://www.w3.org/TR/rdf11-mt/> (2014)
5. Heino, N., Pan, J.Z.: RDFS Reasoning on Massively Parallel Hardware. In: Proceedings of the 11th International Conference on The Semantic Web - Volume Part I. pp. 133–148. ISWC'12, Springer-Verlag, Berlin, Heidelberg (2012)
6. ter Horst, H.J.: Combining RDF and Part of OWL with Rules: Semantics, Decidability, Complexity. In: Gil, Y., Motta, E., Benjamins, V.R., Musen, M.A. (eds.) The Semantic Web - ISWC 2005, 4th International Semantic Web Conference, ISWC 2005, Galway, Ireland, November 6-10, 2005, Proceedings. Lecture Notes in Computer Science, vol. 3729, pp. 668–684. Springer (2005)
7. Kaoudi, Z., Miliaraki, I., Koubarakis, M.: RDFS Reasoning and Query Answering on Top of DHTs. In: Sheth, A.P., et al. (eds.) Proceedings of the 7th International Semantic Web Conference, ISWC 2008, Karlsruhe, Germany, October 26-30, 2008. Lecture Notes in Computer Science, vol. 5318, pp. 499–516. Springer (2008)
8. Kazakov, Y., Krötzsch, M., Simancik, F.: Concurrent Classification of EL Ontologies. In: 10th International Semantic Web Conference, Bonn, Germany, October 23-27. Lecture Notes in Computer Science, vol. 7031, pp. 305–320. Springer (2011)
9. Kolovski, V., Wu, Z., Eadon, G.: Optimizing Enterprise-Scale OWL 2 RL Reasoning in a Relational Database System. In: Patel-Schneider, P.F., Pan, Y., Hitzler, P., Mika, P., Zhang, L., Pan, J.Z., Horrocks, I., Glimm, B. (eds.) International Semantic Web Conference (1). Lecture Notes in Computer Science, vol. 6496, pp. 436–452. Springer (2010)
10. Kotoulas, S., Oren, E., van Harmelen, F.: Mind the Data Skew: Distributed Inferencing by Speeddating in Elastic Regions. In: Proceedings of the 19th International Conference on World Wide Web. pp. 531–540. WWW '10, ACM, New York, NY, USA (2010)
11. Lécué, F., Tucker, R., Bicer, V., Tommasi, P., Tallevi-Diotallevi, S., Sbodio, M.L.: Predicting Severity of Road Traffic Congestion using Semantic Web Technologies. In: Presutti, V., d'Amato, C., Gandon, F., d'Aquin, M., Staab, S., Tordai, A. (eds.) Proceedings of the 11th Extended Semantic Web Conference (ESWC2014), Anisaras, Crete, Greece, May 25–May 29, 2014. Lecture Notes in Computer Science, vol. 8465. Springer (2014)
12. Lehmann, J.: DL-Learner: Learning concepts in description logics. *Journal of Machine Learning Research (JMLR)* 10, 2639–2642 (2009)
13. Ma, Y., Syamsiyah, A.: A hybrid approach to learn description logic ontology from texts. In: Horridge, M., Rospocher, M., van Ossenbruggen, J. (eds.) Posters & Demonstrations Track of the 13th International Semantic Web Conference, ISWC 2014, Riva del Garda, Italy, October 21, 2014. CEUR Workshop Proceedings, vol. 1272, pp. 421–424 (2014)
14. Motik, B., Grau, B.C., Horrocks, I., Wu, Z., Fokoue, A., Lutz, C. (eds.): OWL 2 Web Ontology Language Profiles. W3C Recommendation (2012), available at <http://www.w3.org/TR/owl2-profiles/>
15. Motik, B., Nenov, Y., Piro, R., Horrocks, I.: Parallel Materialisation of Datalog Programs in Main-Memory RDF Databases. In: Brodley, C.E., Stone, P. (eds.) Proceedings of the Twenty-Eighth AAAI Conference on Artificial Intelligence, July 27-31, 2014, Québec City, Québec, Canada. AAAI Press (2014)

16. Mutharaju, R., Hitzler, P., Mateti, P.: Distributed OWL EL Reasoning: The Story So Far. In: Liebig, T., Fokoue, A. (eds.) Proceedings of the 10th International Workshop on Scalable Semantic Web Knowledge Base Systems, Riva Del Garda, Italy. CEUR Workshop Proceedings, vol. 1261, pp. 61–76. CEUR-WS.org (2014)
17. Mutharaju, R., Hitzler, P., Mateti, P., Lécué, F.: Distributed and Scalable OWL EL Reasoning. In: Gandon, F., Sabou, M., Sack, H., d’Amato, C., Cudré-Mauroux, P., Zimmermann, A. (eds.) The Semantic Web. Latest Advances and New Domains - 12th European Semantic Web Conference, ESWC 2015, Portoroz, Slovenia, May 31 - June 4, 2015. Proceedings. Lecture Notes in Computer Science, vol. 9088, pp. 88–103. Springer (2015)
18. Oren, E., Kotoulas, S., Anadiotis, G., Siebes, R., ten Teije, A., van Harmelen, F.: Marvin: Distributed reasoning over large-scale Semantic Web data. *Web Semantics: Science, Services and Agents on the World Wide Web* 7(4), 305–316 (2009)
19. Peters, M., Sachweh, S., Zündorf, A.: Large Scale Rule-Based Reasoning Using a Laptop. In: Gandon, F., Sabou, M., Sack, H., d’Amato, C., Cudré-Mauroux, P., Zimmermann, A. (eds.) The Semantic Web. Latest Advances and New Domains - 12th European Semantic Web Conference, ESWC 2015, Portoroz, Slovenia, May 31 - June 4, 2015. Proceedings. Lecture Notes in Computer Science, vol. 9088, pp. 104–118. Springer (2015)
20. Ren, Y., Pan, J.Z., Lee, K.: Parallel ABox reasoning of EL ontologies. In: Pan, J.Z., Chen, H., Kim, H., Li, J., Wu, Z., Horrocks, I., Mizoguchi, R., Wu, Z. (eds.) Proceedings of the 2011 Joint International Conference on the Semantic Web. JIST’11, vol. 7185, pp. 17–32. Springer (2012)
21. Salvadores, M., Correndo, G., Harris, S., Gibbins, N., Shadbolt, N.: The Design and Implementation of RDFS Backward Reasoning in 4Store. In: Proceedings of the 8th Extended Semantic Web Conference on The Semantic Web: Research and Applications - Volume Part II. pp. 139–153. ESWC’11, Springer-Verlag, Berlin, Heidelberg (2011)
22. Soma, R., Prasanna, V.K.: Parallel Inferencing for OWL Knowledge Bases. In: 2008 International Conference on Parallel Processing, ICPP 2008, September 8-12, 2008, Portland, Oregon, USA. pp. 75–82. IEEE Computer Society (2008)
23. Urbani, J., van Harmelen, F., Schlobach, S., Bal, H.E.: QueryPIE: Backward Reasoning for OWL Horst over Very Large Knowledge Bases. In: Aroyo, L., Welty, C., Alani, H., Taylor, J., Bernstein, A., Kagal, L., Noy, N.F., Blomqvist, E. (eds.) 10th International Semantic Web Conference, Bonn, Germany, October 23-27, 2011. Lecture Notes in Computer Science, vol. 7031, pp. 730–745. Springer (2011)
24. Urbani, J., Kotoulas, S., Oren, E., van Harmelen, F.: Scalable Distributed Reasoning Using MapReduce. In: Bernstein, A., Karger, D.R., Heath, T., Feigenbaum, L., Maynard, D., Motta, E., Thirunarayan, K. (eds.) The Semantic Web - ISWC 2009, 8th International Semantic Web Conference, ISWC 2009, Chantilly, VA, USA, October 25-29, 2009. Proceedings. Lecture Notes in Computer Science, vol. 5823, pp. 634–649. Springer (2009)
25. Weaver, J., Hendler, J.A.: Parallel Materialization of the Finite RDFS Closure for Hundreds of Millions of Triples. In: Bernstein, A., Karger, D.R., Heath, T., Feigenbaum, L., Maynard, D., Motta, E., Thirunarayan, K. (eds.) 8th International Semantic Web Conference, ISWC 2009, Chantilly, VA, USA, October 25-29, 2009. Lecture Notes in Computer Science, vol. 5823, pp. 682–697. Springer (2009)