

Efficient Processing of RDF Graph Pattern Matching on MapReduce Platforms

Padmashree Ravindra Seokyong Hong HyeongSik Kim Kemafor Anyanwu

Department of Computer Science, North Carolina State University

{pravind2, shong3, hkim22, kogam}@ncsu.edu

ABSTRACT

Broadened adoption of the Linking Open Data tenets has led to a significant surge in the amount of Semantic Web data, particularly RDF data. This has positioned the issue of scalable data processing techniques for RDF as a central issue in the Semantic Web research community. The RDF data model is a fine grained model representing relationships as binary relations. Thus, answering queries (typically *graph pattern matching* queries) over RDF data requires several join operations to reassemble related data. While MapReduce based processing is emerging as the *de facto* paradigm for processing large scale data, it is known to be inefficient for join-intensive workloads. In addition, most of the existing techniques for optimizing RDF data processing do not transfer well to the MapReduce model and often require significant lead time for pre-processing. Such a requirement may not be desirable for *on-demand* cloud database scenarios where the goal is to reduce the Time-To-Result (TTR). In this position paper, we argue that some of these challenges can be overcome by rethinking the operators for graph pattern processing, as well as adopting dynamic optimization techniques that exploit information from the previous execution steps in the current execution steps. We present some preliminary evaluation results of the proposed techniques.

Categories and Subject Descriptors

H.3.4 [Semantic Web]:

General Terms

Algorithms, Languages, Performance

Keywords

RDF graph pattern matching, SPARQL, MapReduce, Hadoop

1. INTRODUCTION

There has been a rapid increase in the amount of available Semantic Web data in the past few years as Semantic Web tenets gained broadened adoption. This includes data from domains ranging from scientific (e.g. DrugBank, Linked Clinical Trials), business (e.g. BBC, New York Times), government (e.g. data.gov, data.gov.uk) and general purpose communities (Wikipedia, Linked Open Data). To give a sense of the kind of growth, we note that the Linked Open Data cloud grew from around 26 billion RDF¹ triples in September 2010 to over 31 billion triples by the following year [27]. Therefore, one of the important issues in Semantic Web community is the development of scalable and efficient techniques for processing large amounts of Semantic Web data. The most common representation model for Semantic Web is called the Resource Description Framework (RDF). RDF database is a collection of triples (*Subject*, *Predicate*, *Object*) where *Predicates* are named binary relations between *Subject* and *Object* that represent either resources or literal values in the Web.

The RDF data model can also be viewed as a directed and labeled graph in which edges are labeled with predicate names, connecting between graph nodes labeled with URIs and values, respectively.

The key construct for processing RDF data is *graph pattern matching*. Here, users describe the structure of patterns that they are interested in as queries and systems return as answers all occurrences of the pattern found in the data. The standard query language for expressing graph pattern matching queries on RDF data is called SPARQL². Each triple pattern in a given SPARQL query is a triple in which at least one of the *Subject*, *Predicate* or *Object* is a variable, denoted by a leading question mark. Therefore, the query attempts to match those patterns to sub graphs in the database and the result of a graph pattern query is a list of all variables substituted from those graphs. For example, consider the triple {Vendor1, homepage, www.vendors.org/v1} in Figure 1 (a) which models that a resource Vendor1's homepage is www.vendors.org/v1. Now we have the following triple pattern in a SPARQL query: {Vendor1 homepage ?homepage}. It will try to match its variable ?homepage from the triples, therefore, the value of ?homepage will be www.vendors.org/v1.

Traditionally, RDF data is stored as a ternary relation and graph pattern matching queries are processed as a series of relational join operations. Due to the fine-grained nature of the RDF data model, it is common to require several join operations to answer a slightly complex query. Another commonly performed task on Semantic Web data is *reasoning*, that applies the standardized semantic inference rules to compute all inferable facts from the set of explicitly represented facts. Depending on the approach for performing reasoning, this task also relies heavily on joins. Thus, the ability to process join-intensive workloads is crucial for the Semantic Web.

When considering large scale processing and analytics in data-intensive applications using cloud environment, the recent *de facto* standard is the MapReduce [5] programming model made popular by Google and its open source implementations such as Hadoop [26]. Such platforms have been explored for graph pattern matching [8][9][21][23]. However, existing approaches naively translate join-intensive workloads such as graph pattern matching into a long chain of MapReduce jobs, leading to significant I/O, communication, and sorting overhead as discussed in the next section. Further, we also show how adopting existing optimization techniques used for single-node environments to MapReduce framework is not straightforward. Therefore, it is our position that the development of new strategies is crucial.

¹Resource Description Framework (RDF) <http://www.w3.org/RDF>

² SPARQL Query Language for RDF. <http://www.w3.org/TR/rdf-sparql-query/>

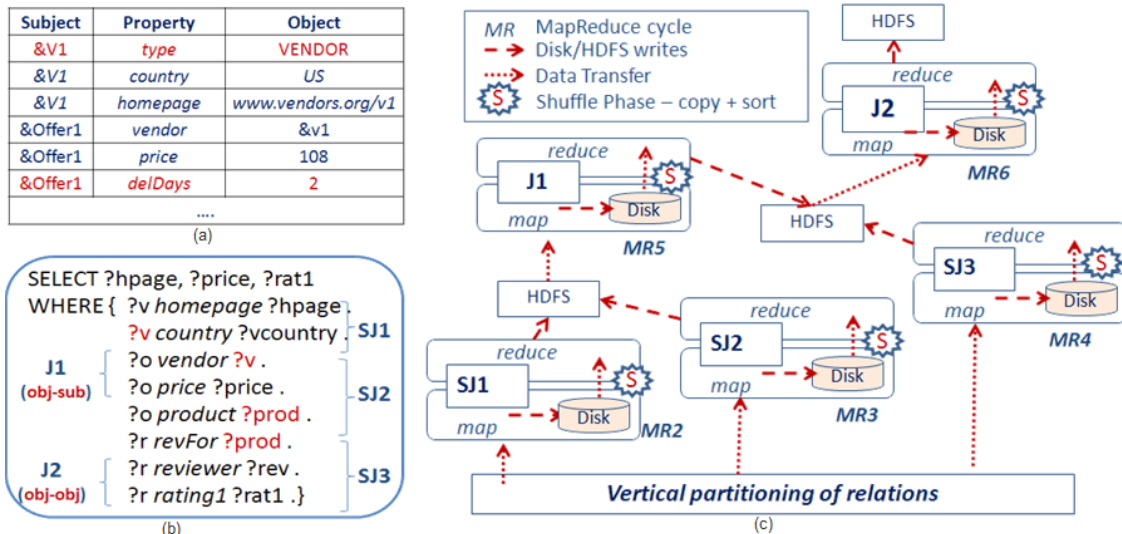


Figure 1 (a) Example RDF triples representing Vendors and their Product Offers (b) SPARQL query to retrieve details about Vendors, their Offers and the Product Reviews (c) Query processing in Pig

We overview two complimentary strategies that can be used to optimize graph pattern matching queries without the need for time consuming preprocessing, (i) *algebraic optimization based on a new algebra called the Nested TripleGroup Algebra* and (ii) *information passing*. The remainder of this paper is organized as follows: Section 2 provides some relevant background. Section 3 gives an overview of the challenges that we encountered. Section 4 describes the approach we are proposing to overcome the challenge.

2. BACKGROUND

2.1 Join Processing in MapReduce

In the MapReduce programming model, users encode their tasks in terms of two functions: the *map* whose general signature is $\text{map}(Key1, Value1) \rightarrow \text{list}(Key2, Value2)$, processes input key-value pairs and produces intermediate data as a list of key-value pairs; the *reduce* with signature $\text{reduce}(Key2, \text{list}(Value2)) \rightarrow (Key3, Value3)$ merges values into a group according to the intermediate key, and generates a final key-value pair for each group. Hadoop is an open source implementation of the MapReduce model and provides fault-tolerance and automatic parallelization of the map and reduce functions. Its architecture includes a master process (*JobTracker*) which schedules m instances of the map function on nodes (*mappers*) and r instances of the reduce function on nodes (*reducers*) over the cluster. The JobTracker splits the input data into “chunks” which are assigned to mappers. Once the mappers produce and write their output i.e. collection of key value pairs to disk, each key is mapped to a unique reduce task. This mapping forms the basis of data distribution among the reducers and thus they may get unequal amounts of data to process. The reduce phase consists of 3 sub-phases: *copy* - the map output is copied from the disk at the mapper nodes to reducer nodes, *sort* - the collected map output is sorted based on key values and *reduce* - reduce function is applied to the data. After the reduce function is completed, their outputs are saved to the *Hadoop Distributed File System - HDFS*.

To interpret the relational join operation in the MapReduce paradigm, each tuple from the participating relations is tagged based on the join columns. This will enable all the records with the same join key to be assigned and transferred to the same

reducer. In reduce-phase, each of r reducers accumulates the records into separate in-memory buffers according to their tags and performs the join. This is called the *Standard Repartitioning Join* [3]. One thing to note is that all intermediate data generated in map-phase has to be sorted and transferred into the reduce-phase. These data sorting and transferring steps are crucial cost factors that may affect the data processing time on MapReduce [7][10][13]. Moreover, in a case that the join selectivity is high, a large number of records are involved in such data sorting and transferring steps even if those records are not eventually joined. *Fragment-Replication Join* and *Map-Merge Join* are alternative join approaches that can remove the data sorting and shuffling overhead by processing the join in the map-phase [3]. However, the fragment-replication join is applicable only when one of the input tables is small enough to fit in memory. In the case of the map-merge join, an additional pre-processing phase is required to partition input data according to its join key. More importantly, this approach cannot effectively deal with scenarios where one of the inputs to the join is an intermediate result generated from earlier query operators. Another recently proposed technique is the *Semi-Join* [3], but it also suffers from similar limitations.

2.2 Automatic Generation of MapReduce Execution Workflows

Typical data processing tasks may involve several operations or multiple instances of the same operation. A task may include one or more filtering conditions, join, grouping, and aggregation operations. Executing such a task on MapReduce requires users to determine the appropriate map and reduce functions and implement them in a way that achieves the most efficient workflow, which is not always trivial. In order to eliminate this burden on the users, Hadoop extensions such as Apache *Hive* [24] and *Pig* [18] offer high-level declarative query and dataflow languages, *HiveQL* and *Pig Latin* respectively, that are compiled to generate efficient MapReduce workflows. This is similar in spirit to database systems where users write high-level declarative queries and the system is responsible for generating an optimized execution plan. The high-level query expressed by the user is compiled into a series of MapReduce cycles where the output generated at the end of one cycle is saved onto the HDFS and fed as input into the next MR cycle forming a MapReduce workflow. Naïve compilers in some existing systems follow rule-based

translation of each relational operator into one or more MR cycles. Such naïve translations in the case of join-intensive workloads such as RDF graph pattern matching result in lengthy MapReduce workflows with a separate MR cycle for each join operation. The example graph pattern query in Figure 1(b) would be compiled into 7 self-join operations on a triple data model, leading to 7 MR cycles. Lengthy execution workflows lead to performance inefficiency since each MR cycle compounds the overall I/O, sorting, and communication cost, in addition to the materialization cost between two contiguous cycles. This motivates the need for optimization techniques that minimize the length of MapReduce workflows and also reduce the intermediate data. A possible optimization to reduce the required #MR cycles is to group related operations into the same MR cycle. For example, n-way join operations can be used to replace joins on the same column. Using this optimization, the example graph pattern in Figure 1 (b) can be executed as 3 n-way join operations (SJ1, SJ2, and SJ3 respectively), which reduces the overall #MR cycles to 5 (3 for the n-way joins, and 2 to further join them) as shown in Figure 1 (c). Further, we can avoid the expensive self-join operations on large ternary relations and reduce the initial load costs by adopting the *vertical partitioning* (VP) [1] storage model that partitions the input into property-based relations. The VP approach can be implemented in Pig Latin by using the `SPLIT` operator. Figure 2 (top) shows Pig’s logical plan corresponding to our example query in Figure 1 (b) using the VP approach. In the next section, we discuss other optimization techniques and review the challenges of optimizing graph pattern matching queries on MapReduce-based platforms.

3. CHALLENGES

The key challenges to overcome in order to achieve efficient processing of RDF data processing are (1) minimizing the length of MapReduce execution workflows – reduces the overall costs of processing by reducing the number of iterations for disk I/Os, communication and sorting steps; (2) minimizing the size of intermediate data - reduces the amount of data written and read from disk i.e. disk I/O, the amount of data communicated between nodes and, the size of data sorted during the sorting steps. A natural inclination is to explore how existing optimization strategies for graph pattern matching designed for single-node environments can be adopted into the MapReduce model to achieve these goals. As the following discussion will show, there are several challenges that limit the possibility of doing so, suggesting the need for other kinds of strategies.

The state-of-the-art techniques for graph pattern matching [15][1][25] rely heavily on comprehensively indexing RDF data. This means that data is indexed in all possible ways that allows every execution step to be supported by some index. In particular,

the goal is always to enable the use of the merge-join algorithm for as many join operations as possible because it is the most efficient join algorithm for sorted data. Consequently, most indexing schemes are ordered index structures such as B+Trees that maintain data in a desired order. These comprehensive indexing strategies enable “index-only” plans that do not require any other files for query processing except the index files. Given this scenario, the major challenge in the context of MapReduce is how to use these indexes during the map and reduce phases of each join. These data structures will need to be “chunked” and partitioned across nodes similar to the way data is done. Such a partitioning strategy is not straightforward because such tree structures will need to be chunked in meaningful ways that maintain root-to-leaf path relationships. Further, this approach would require performing a single-join-operation-per-MR-cycle. Note that the cost based optimization used in RDF-3X just changes the order of joins but does not change the number of join steps.

While the previously discussed techniques index data at the level of binary relations, other approaches try to index data at a coarser level of granularity. This is motivated by the fact that the task of reassembling related data in RDF is very common in graph pattern matching queries and hence a good fraction of the joins in queries focus on this reassembly task. The reassembly process results in star-shaped join structures that can be exploited for optimization (refer to SJ1, SJ2, and SJ3 in Figure 1(b)). In order to efficiently match the star-structured graph patterns to star subgraphs in data, we can cluster related data prior to query processing. By doing so, it might appear that we could eliminate the need for several join operations potentially reducing the number of MapReduce cycles. In fact, it may be possible to implement such join operations as a special kind of complex filter operation in this storage model. Unfortunately, realizing this is not as straightforward as it might first seem. Since RDF is semi-structured and represented using binary relations, clustering related data results in variable length tuples containing data and metadata. As example, in the strategy employed in SHARD [21], the example data will appear as the following after clustering:

(&V1, type, VENDOR, country, US, country, China, homepage, www.vendors.org)

(&Offer1, vendor, &V1, price, 108, delDays, 2)

To match the star-join pattern with the *country* and *homepage* predicates, the algorithm will first need to identify the relevant columns in each tuple, since the columns in this clustered model are not identified in a meaningful way like in the relational model.

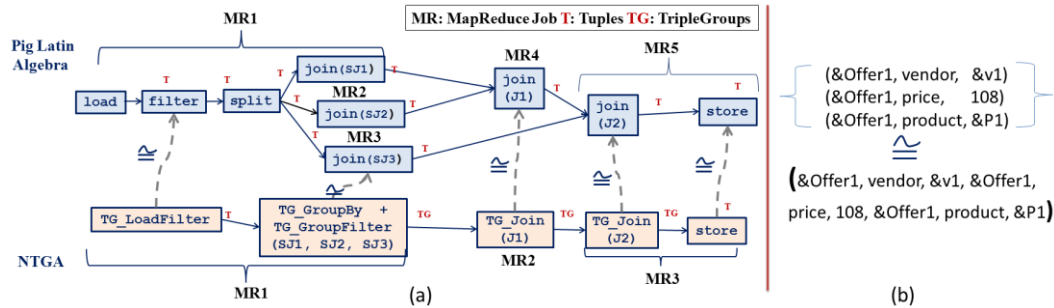


Figure 2: (a) Comparison and mapping between relational and NTGA for the example query plan (b) An example TripleGroup and its content equivalent n-tuple

Here attribute names as well as their values are part of the data and the number of occurrences of a particular attribute may differ from row to row. Therefore, extracting the relevant fields does not translate to the relational filter operation, which motivates the need for a specialized operator when using this storage model. In SHARD, the matching process is done using an iterative join strategy where during each iteration, the columns related to one particular predicate are identified using a special function and the values in those columns are joined to the current intermediate result. This approach results in as many MapReduce cycles as there as join operations.

Another limitation common to all these discussed approaches is that they all require preprocessing either for computing indexes or clustering data and so on. Often these preprocessing steps are in the order of hours which can be limiting in the case of on-demand data processing tasks where the main goal is to reduce the Time-To-Result.

In the sequel, we will present some techniques that address the two challenges earlier mentioned. To reduce the length of the workflow, we propose to *reinterpret the set of star joins in a query as a GROUP-BY operation* [19][20]. The results are equivalent to the results using join operations with respect to content but differ in their structure and representation. However, this allows all star-joins to be computed concurrently in a single MR cycle. To deal with the problem of reducing intermediate data, we propose an *information passing* approach that summarizes the content of intermediate execution steps and passes the information to future execution steps to help inform their disk retrieval tasks about what data items are no longer relevant to processing.

4. PROPOSED APPROACH

4.1 Algebraic Optimization - TripleGroup based Processing

We propose to re-interpret the star-join operations in a graph pattern as a GROUP BY operation. The rationale for the grouping approach is that typical graph pattern matching queries involve multiple star-join structures (50% of BSBM benchmark queries have 2 or more star-patterns), which would have otherwise led to a separate MR cycle for each star-join. However, by using a single GROUP BY operation, all the required star-joins can be computed concurrently in a single MR cycle. For example, the 3 star-joins

(SJ1, SJ2, and SJ3) in Figure 1 (b) can be computed using a single GROUP BY on *Subject* in a single MR cycle, thus reducing the required #MR cycles. However, this grouping strategy produces “groups of triples” or *TripleGroups* instead of n-tuples, and requires specialized operators for efficient processing of TripleGroups. Figure 2 (b) (top) shows an example TripleGroup corresponding to the group of triples with *Subject* & Offer1.

We propose an intermediate *Nested Triplegroup algebra* (NTGA) that supports special TripleGroup based operators for efficient processing of RDF graph patterns. The NTGA comprises of operators to load RDF statements into TripleGroups (TG_Load), project a *Subject / Object* component of a triple (TG_Project), value-based filtering of TripleGroups based on filter conditions (TG_filter), structure-based filtering of TripleGroups to eliminate sub-structures with missing edges (TG_Groupfilter), and join between TripleGroups representing different structures (TG_Join). It also supports flattening operators to create an n-tuple equivalent of TripleGroup (TG_Flatten), and to flatten a nested TripleGroup (TG_Unnest). Figure 2 (a) shows a comparison of the MR execution workflow to process the example query in Figure 1 (b) using the Pig Latin and the NTGA operators. It also demonstrates that the example query can be executed in just 3 MR cycles using NTGA as opposed to 5 MR cycles using the Pig approach. Figure 2 (a) also represents the mapping between the two plans through the concept of *content equivalence* that enables lossless transformation between queries written in relational algebra and NTGA. Figure 2 (b) (bottom) shows the equivalent n-tuple resulting from the TG_Flatten operator on the example TripleGroup (top).

4.2 Dynamic Optimization - Information Passing

Existing sideways information passing techniques assume that operators in a query plan run in parallel. However, this assumption is not valid in MapReduce-based data processing systems since operators in a query plan are physically and temporally isolated from each other. Moreover, since there is no effective communication method that enables operators to communicate with each other, we cannot directly adopt such traditional information passing techniques over MapReduce frameworks.

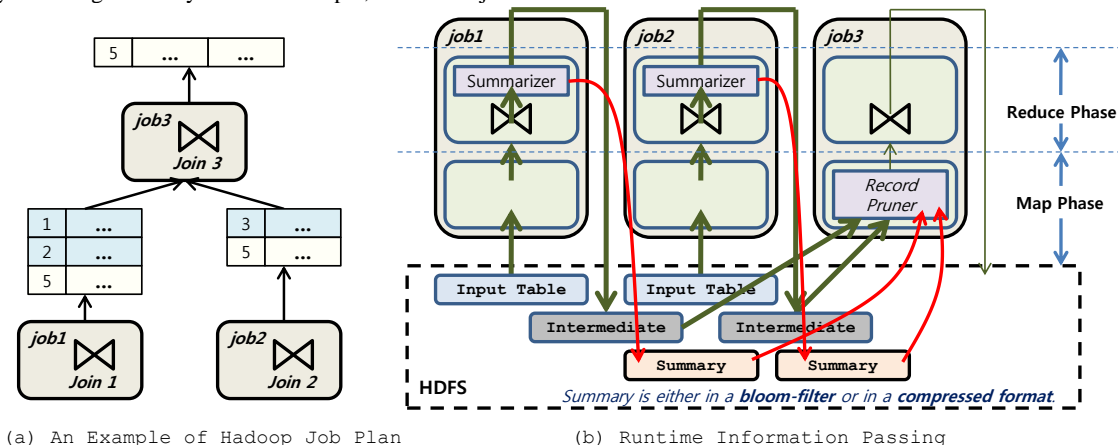


Figure 3: Enabling information passing in Hadoop query plans

Hence, we design and develop a new adaptive information passing technique called *Inter-job Information Passing* that runs At compile-time, the query optimizer collects all the possible dataflow between MapReduce jobs from the physical query plan, calculates all the equivalent classes, creates an information passing plan, and finally embeds the plan into the Hadoop job plan. Figure 3 (a) is an example of a Hadoop job plan where the Hadoop job *job3* receives its inputs from the previous jobs, *job1* and *job2*. Figure 3 (b) shows a brief representation of the run-time information passing operations. When *job1* and *job2* are executed, the summarizer receives the streams of outputs, and creates the summary information on the columns that will be joined in *job3*. Then, the summary information is represented in a more compact format to reduce the data transfer costs, and is stored in the shared information store that resides on HDFS. When *job3* runs, the previous summary information is loaded into *PSO (Parameterized Filter Operator)* at Map-phase and used to filter out the input data streams. For example, in Figure 3(b), *job1* generates a filter that contains a list of distinct join key values - 1, 2 and 5 and *job2* generates a filter containing a list of 3 and 5 respectively. When *job3* runs, it loads the two lists into the in-memory buffer and filters out all records from each input stream that do not match the values in the corresponding list.

Another important issue is that the overhead resulting from the creation and the transmission of summary information may exceed the performance gain achieved by the information passing technique in some cases. Therefore, we provide a *benefit estimation model* and integrate it with the query optimizer to dynamically enable or disable the feature of the technique. Our benefit estimation model considers the possible benefits in sorting, transferring, and merging steps which can be earned from data reduction caused by a given size of summary information. Also, the model considers the overhead caused by generating and transferring summary information via Hadoop’s DistributedCache.

4.3 General Query Plan Strategy

We address the two challenges in MapReduce-based graph pattern matching (described in section 3) by applying the two proposed techniques. First, we minimize the #MR cycles required for the star-join computation phase by using the TripleGroup-based processing using the NTGA operators. In addition, we can collect the summary information during NTGA’s star-join computation phase and integrate the information-passing technique to reduce the intermediate data in the subsequent MR cycles. Thus, in the map phase of the `TG_Join` operator, we can eliminate irrelevant TripleGroups that do not join in the subsequent phases. Both these techniques help us minimize the overall I/O, communication, sorting and materialization costs, leading to efficient processing of RDF data.

5. CASE STUDY

In this section, we present an empirical study of the performance of the two proposed optimization techniques to enable efficient RDF graph pattern matching on MapReduce-based platforms.

5.1 Task Description

Task A - Scalability of TripleGroup-based processing: We evaluated the performance of the NTGA approach with increasing size of RDF graphs. We extended Pig to integrate the NTGA operators [11] and compared its performance with that of RDF graph processing using relational operators in naïve Pig.

Task B - Performance Evaluation of Hybrid Plans: We also performed a comparative evaluation of a hybrid approach containing a mix of Pig Latin and NTGA operators. The notion of

content equivalence allows for lossless translations between the algebra as shown in Figure 2 (a). For example, each TripleGroup in the result of the star-join computation phase using NTGA’s `TG_GroupBy` (MR1 in NTGA plan in Figure 2 (a)) is content equivalent to some n-tuple in the result of the corresponding phase using Pig’s `JOIN` operator (MR1, MR2, and MR3 in Pig Latin plan). For the hybrid approach, we considered two execution plans, (i) *NTGA-StarJoin* – using NTGA’s `TG_GroupBy` operator for star-join computations, `TG_Flatten` to flatten the TripleGroups into n-tuples, and Pig’s `JOIN` operator to compute the join between the stars; (ii) *Pig-StarJoin* – Pig’s n-way `JOIN` operator for star-join computations, NTGA’s `TG_Unflatten` to convert n-tuples to TripleGroups, and NTGA’s `TG_Join` to compute the join between the stars.

Task C – Scalability of Information-Passing technique: We evaluated the benefit of *information passing*, with increasing size of RDF graphs. We extended Hive to enable inter-job information passing (IIP-enabled Hive) and compared its performance to the default Hive implementation. However, the proposed information passing technique uses a general and abstract query processing model and, therefore, is applicable to other MapReduce-based data processing systems such as Pig.

5.2 Setup and Testbed

Environment: The experiments were conducted on a 5-node cluster on NC State’s VCL [22], an on-demand virtual computing environment. Each node in the cluster was a dual core Intel x86 machine with 2.33 GHz processor speed, 3GB RAM and running Red Hat Linux. Experiments used Pig release 0.8.0 and Hive 0.5.0 running on Hadoop 20.0 with block size 256MB.

Testbed – Data Set and Queries: Synthetic BSBM³ dataset generator was used to generate n-triple files for Task A, with data size ranging from 13GB to 44GB (approx. 170 million triples). The query used for Task A consisted of 6 triples patterns requiring 5 join operations, which comprised of two star-join structures. Two test queries with varying star cardinality were used for Task B (i) *q-small* – three star sub-patterns with 1, 3 and 1 triple pattern in each star respectively, (ii) *q-dense* – three star sub-patterns with 3 triple patterns in each star. Task B was evaluated on input data containing 50,000 products, which corresponds to approx. 17 million triples (approx. 4.3 GB data size). For Task C, three separate Hive tables were generated from BSBM in the format of SQL-dump with the #products ranging from 1000 to 500,000 (approx. 50GB data size). Task C evaluated a query requiring 3 join operations on these three tables.

5.3 Evaluation Results

Task A: Figure 5 (a) shows the execution times for the NTGA and naïve Pig implementations. For all the four data sizes, we see up to 60% improvement in the execution times with the NTGA approach. In the Pig approach, the most efficient execution plan is to compute the two star-patterns using n-way joins, which can be achieved in 3 MR cycles. NTGA’s grouping approach computes the 2 star-joins in a single MR cycle, reducing the required #MR cycles to just 2. The advantage is more obvious in complex graph pattern queries with several star-joins.

Task B: Figure 4 shows the comparative evaluation of the two hybrid plans along with the NTGA and Pig execution plans for the two query patterns *q-small* and *q-dense* respectively.

³<http://www4.wiwiw.fu-berlin.de/bizer/BerlinSPARQLBenchmark/>

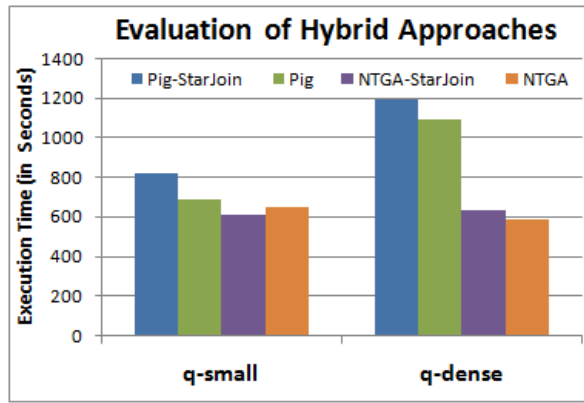


Figure 4: Performance evaluation of hybrid plans with varying density of star sub-patterns (4.3GB, 5-node cluster)

Pig-StarJoin and Pig show low performance in both cases since the star-join computation using Pig’s JOIN leads to lengthy execution workflows. Pig-StarJoin does worse than Pig due to the overhead of converting n-tuples to TripleGroups before the join between the stars. On the other hand, NTGA-StarJoin and NTGA show a performance gain of 42% and 46% respectively over Pig

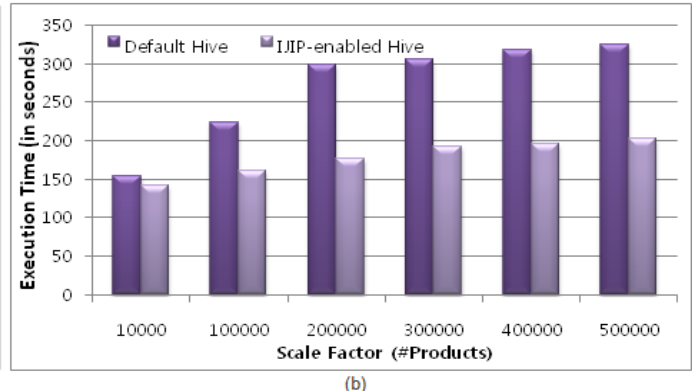
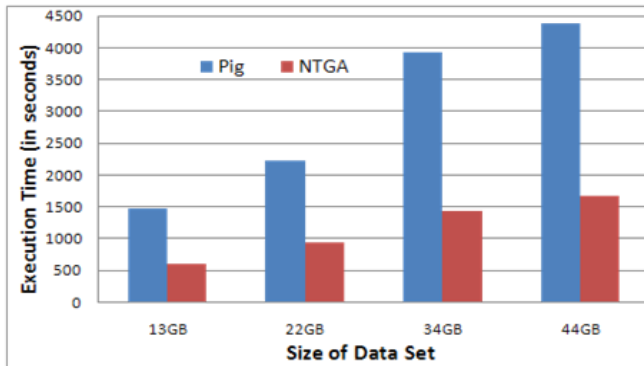


Figure 5: Scalability study with increasing size of RDF graphs for (a) TripleGroup-based processing of RDF graph patterns and (b) Inter-job Information Passing technique for join-intensive workloads (5-node cluster)

8. REFERENCES

- [1] Daniel J. Abadi, Adam Marcus, Samuel R. Madden, and Kate Hollenbach. Scalable Semantic Web Data Management using Vertical Partitioning. In Proc. International Conference on Very large data bases, 2007.
- [2] Ron Avnur and Joseph M. Hellerstein. Eddies: Continuously Adaptive Query Processing. In Proc. International conference on Management of Data, pages 261–272, 2000.
- [3] Spyros Blanas, Jignesh M. Patel, Vuk Ercegovic, Jun Rao, Eugene J. Shekita, and Yuanyuan Tian. A Comparison of Join Algorithms for Log Processing in MapReduce. In Proc. International conference on Management of Data, 2010
- [4] Michael J. Carey. BEA Liquid Data for WebLogic: XML-Based Enterprise Information Integration. In Proc. International Conference on Data Engineering, pages 800–, 2004.
- [5] Jeffrey Dean and Sanjay Ghemawat. MapReduce: Simplified Data Processing on Large Clusters. Proceedings of the Conference on Symposium on Operating Systems Design & Implementation. Vol. 6. 10-10. 2004
- [6] Jens Dittrich, Jorge-Arnulfo Quiané-Ruiz, Alekh Jindal, Yagiz Kargin, Vinay Setty, and Jörg Schad. Hadoop++:

approach for the denser query pattern. This demonstrates the advantage of shorter execution plans using the NTGA approach to compute star-joins.

Task C: Figure 5(b) shows that the IP-enabled implementation outperforms the default Hive implementation as the size of the input grows. This provides a promising approach to reduce the size of intermediate data, specifically for processing join-intensive workloads on MapReduce based platforms.

6. CONCLUSION

In this position paper, we propose two optimization techniques that enable efficient processing of join-intensive workloads such as RDF graph pattern matching. These two techniques help to shorten the MR execution workflows, and reduce the intermediate tuples, resulting in less I/O, sorting, and communication costs involved in traditional MapReduce based RDF graph processing.

7. ACKNOWLEDGMENTS

This work was partially funded by NSF grant IIS-0915865.

Making a Yellow Elephant Run Like a Cheetah (without It Even Noticing), Proc. VLDB Endow. 3, volume 3, 515–529, 2010

- [7] Laura M. Haas, Donald Kossmann, Edward L. Wimmers, and Jun Yang. Optimizing Queries Across Diverse Data Sources. In Proc. International Conference on Very Large Data Bases, VLDB ’97, pages 276–285, 1997.
- [8] Jiewen Huang, Daniel J. Abadi, and Kun Ren. Scalable SPARQL Querying of Large RDF Graphs. In Proc. VLDB Endowment, 4(11), 2011.
- [9] Mohammad Farhan Husain, Latifur Khan, Murat Kantarcioglu, and Bhavani Thuraisingham. Data Intensive Query Processing for Large RDF Graphs Using Cloud Computing Tools. In Cloud Computing (CLOUD), IEEE International Conference on, 2010.
- [10] Zachary G. Ives and Nicholas E. Taylor. Sideways Information Passing for Push-Style Query Processing. In Proc. Conference on Data Engineering, pages 774–783, 2008.
- [11] Hyeongsik Kim, Padmashree Ravindra, and Kemafor Anyanwu. From SPARQL to MapReduce: The Journey

- Using a Nested TripleGroup Algebra. Proceedings of the VLDB Endowment, 4(12), 2011.
- [12] Grzegorz Malewicz, Matthew H. Austern, Aart J.C. Bik, James C. Dehnert, Ilan Horn, Naty Leiser, and Grzegorz Czajkowski. Pregel: A System for Large-scale Graph Processing. In Proc. International Conference on Management of Data, 2010.
- [13] Inderpal Singh Mumick and Hamid Pirahesh. Implementation of Magic-sets in a Relational Database System. In Proc. international conference on Management of Data, SIGMOD '94, pages 103–114, 1994.
- [14] Thomas Neumann and Gerhard Weikum. Scalable Join Processing on Very Large RDF Graphs. In Proc. International Conference on Management of Data, 2009.
- [15] Thomas Neumann and Gerhard Weikum. The RDF-3X engine for scalable management of RDF data. The VLDB Journal, 19:91–113, 2010.
- [16] Andrew Newman, Jane Hunter, Yuan-Fang Li, Chris Bouton, and Melissa Davis. A Scale-out RDF Molecule Store for Distributed Processing of Biomedical Data. In Semantic Web for Health Care and Life Sciences Workshop, 2008.
- [17] Andrew Newman, Yuan-Fang Li, and Jane Hunter. Scalable Semantics: The Silver Lining of Cloud Computing. In eScience. IEEE International Conference on, 2008.
- [18] Christopher Olston, Benjamin Reed, Utkarsh Srivastava, Ravi Kumar, and Andrew Tomkins. Pig Latin: A Not-So-Foreign Language for Data Processing. In Proc. International Conference on Management of Data, 2008.
- [19] Padmashree Ravindra, Vikas V. Deshpande, and Kemafor Anyanwu. Towards Scalable RDF Graph Analytics on MapReduce. In Proc. Workshop on Massive Data Analytics on the Cloud, 2010.
- [20] Padmashree Ravindra, HyeongSik Kim, and Kemafor Anyanwu. An Intermediate Algebra for Optimizing RDF Graph Pattern Matching on MapReduce. In Proc. Extended Semantic Web Conference, 2011.
- [21] Kurt Rohloff and Richard E. Schantz. High-performance, Massively Scalable Distributed Systems using the MapReduce Software Framework: the SHARD Triple-store. In Programming Support Innovations for Emerging Distributed Applications, PSI EtA '10, pages 4:1–4:5, 2010.
- [22] H.E. Schaffer, S.F. Averitt, M.I. Hoit, A. Peeler, E.D. Sills, and M.A. Vouk. NCSU's Virtual Computing Lab: A Cloud Computing Solution. Computer, 42:94–97, 2009.
- [23] Yusuke Tanimura, Akiyoshi Matono, Steven Lynden, and Isao Kojima. Extensions to the Pig data processing platform for scalable RDF data processing using Hadoop. In Data Engineering Workshops, IEEE International Conference on, 2010.
- [24] Ashish Thusoo, Joydeep Sen Sarma, Namit Jain, Zheng Shao, Prasad Chakka, Suresh Anthony, Hao Liu, Pete Wyckoff, and Raghatham Murthy. Hive: a warehousing solution over a map-reduce framework. Proc. VLDB Endow., 2:1626–1629, 2009.
- [25] Cathrin Weiss, Panagiotis Karras, and Abraham Bernstein. Hexastore: Sextuple indexing for semantic web data management. Proc. VLDB Endow., 1:1008–1019, August 2008.
- [26] A. Bialecki, M. Cafarella, D. Cutting, and O. O Malley. Hadoop: A Framework for Running Applications on Large Clusters Built of Commodity Hardware. <http://hadoop.apache.org/>
- [27] R.Cyganiak and A. Jentzsch The Linking Open Data Cloud Diagram <http://www4.wiwiw.fu-berlin.de/locloud/state/>