Selectivity Estimation for SPARQL Triple Patterns with Shape Expressions

Abdullah Abbas

Univ. Grenoble Alpes, CNRS, Grenoble INP, Inria,
LIG
Grenoble, France
[firstname.lastname]@inria.fr

Cécile Roisin

Univ. Grenoble Alpes, CNRS, Grenoble INP, Inria,
LIG
Grenoble, France
[firstname.lastname]@inria.fr

Pierre Genevès

Univ. Grenoble Alpes, CNRS, Grenoble INP, Inria,
LIG
Grenoble, France
[firstname.lastname]@cnrs.fr

Nabil Layaïda

Univ. Grenoble Alpes, CNRS, Grenoble INP, Inria,
LIG
Grenoble, France
[firstname.lastname]@inria.fr

ABSTRACT

ShEx (Shape Expressions) is a language for expressing constraints on RDF graphs. In this work we optimize the evaluation of conjunctive SPARQL queries, on RDF graphs, by taking advantage of ShEx constraints. Our optimization is based on computing and assigning ranks to query triple patterns, dictating their order of execution. We first define a set of well formed ShEx schemas, that possess interesting characteristics for SPARQL query optimization. We then define our optimization method by exploiting information extracted from a ShEx schema. We finally report on evaluation results performed showing the advantages of applying our optimization on the top of an existing state-of-the-art query evaluation system.

KEYWORDS

SPARQL evaluation, ShEx, Shape Expressions, RDF constraints

1 INTRODUCTION

The Shape Expressions (ShEx) language [18] allows to describe constraints on RDF graph structures [20]. These descriptions identify predicates and their associated cardinalities and datatypes. ShEx shapes can be used to validate RDF documents, generate RDF documents, or communicate expected graph patterns associated with some process or interface.

In this work we investigate how the evaluation of SPARQL queries [21] on big RDF graphs can be optimized in the presence of ShEx constraints. We propose a method for optimizing the order of evaluation of subqueries, by taking advantage of the information on the data described in ShEx. While SPARQL query optimization by static analysis is important and well-studied, the emergence of constraint languages (such as ShEx) raises new questions on how the knowledge of additional constraints can be effectively leveraged as a part of the static analysis and optimization, especially in the context of big data structures and the increased attention to its applications. In this work, we focus on the logical query structure

and in particular on subquery ordering that can be automatically inferred from a set of data constraints. More specifically, we consider SPARQL basic graph patterns (BGPs), and we focus on the order of execution of triple patterns that aim to minimize the overall execution cost of the query.

We postulate that ShEx constraints contain useful information for selecting the order of execution of triple patterns. Optimization opportunities arise from the presence of joins between query triple patterns, and common variables. In several situations, the order of execution of triple patterns can be rearranged so that the size of intermediate results for join variables are minimized. Consider the arbitrary query example of Listing 1 with 3 triple patterns and a join on the variable $?x$.

```
SELECT ?x WHERE {  
  ?x :p1000 :a .  
  ?x :p700 :b .  
  ?x :p1 :c . }
```

Listing 1: SPARQL query example

Assume that we know that the triple with predicate $:p1000$ will return 1000 values for $?x$, that of $:p700$ will return 700, and that of $:p1$ will return 1. The join between the first two triple patterns may give up to 700 values which should be reserved in memory for another join with the third triple. A wise choice in this case is to reorder the triple pattern execution, knowing that the third triple is more selective than the other two triples. Executing the third triple first will guarantee that at most one value will be reserved in the memory for the next join. Such an order will provide an optimized execution of the whole query.

Contribution. Hence, our main purpose in this work is to infer better execution orders for triple patterns in more general queries, based on the knowledge extracted from a ShEx document. We first define a set of well formed ShEx schemas, that possess interesting characteristics for deciding optimal execution orders. We then define our optimization

method by exploiting information extracted from a ShEx schema. To the best of our knowledge, this is the first work addressing SPARQL query optimization based on ShEx. We implemented our procedure on the top of SPARQLGX, which is one of the most efficient engine for distributed SPARQL evaluation and known to outperform many competitors in the field [8]. SPARQLGX already implements various query optimization techniques including reordering triple patterns [8], but without considering schema constraints. We show that our technique further improves the efficiency of query execution times. With large amounts of data, the optimization may exceed 85% for queries that are highly affected by the triple patterns execution order, and by 25% on average.

Outline. The rest of the paper is organized as follows: In Sect. 2 we review the closest related works. In Sect. III we introduce some preliminaries necessary for understanding the rest of the paper. In Sect. IV we define well-formation rules for a data-schema pair that are of interest for our optimization process. In Sect. V we define a graph representation of a ShEx schema. In Sect. VI we formally define our optimization process. In Sect. VII we report on experimental results with our optimization technique. Finally, we conclude in Sect. VIII.

2 RELATED WORKS

Query optimization for the RDF data model has been studied with various different approaches in the literature. Most existing works are based on scanning the data a-priori and either saving new pieces of information about it, or providing alternative data representations. The works in [7, 9, 13, 14, 16] are based on techniques that mainly focus on join optimizations by indexing the data. These works do not consider structured data and data typing. Another approach that also does not consider the query structure, yet provides optimized query processing, is vertical data partitioning [1] (which is already used in SPARQLGX [8]).

The works in [6, 10, 15, 17] also provide query optimization techniques, by proposing new data representations that are more compact after scanning the data. Additionally, in [11, 25, 26] the authors study the structure of the data and provide structural summaries or representative schemas. None of these works is based on a given schema, and thus they require an extensive data scan. An important advantage of our approach is that optimizations are computed by analysing only a schema (whose size is much smaller than the actual data); and our optimizations remain valid in the presence of data updates.

For works based on typed data, an approach was proposed in [4] that considers typed XML data trees. Unlike RDF, the tree data-type model of XML allows for extremely efficient subtree pruning. In [2] semantic query optimization for object-oriented databases is considered. In [19, 22] the authors, as in our work, consider query optimization for typed RDF graphs. These works are mainly oriented towards schema violations. In our work, we mainly assume non-violating queries, and

we study the effect of reordering which is not studied in the previous works.

In [8] the authors provide a SPARQL query evaluator, SPARQLGX, that relies on a translation of SPARQL queries into executable Spark code that adopts evaluation strategies according to the storage method used and statistics on data. Within the system, optimized joins are considered by reordering BGP triple patterns by combining those with common variables, but this reordering does not consider the selectivity of triples based on the structure of the RDF data. Their approach scales better than the state-of-the-art systems they compare with [8]. For this reason we implemented and tested our optimization technique on the top of this system. In the experimental part, we compare the results that we obtain with and without our optimizations to the results obtained with SPARQLGX.

A novelty of our work is that it provides a new optimization technique that is based on a simple analysis of a given schema. optimizations hold as long as the data conform to the schema, i.e. possibly at no cost in the presence of updates. Our work could also be applied as a supplement to a wide range of other previously known optimizations techniques like [1, 8].

3 DEFINITIONS

3.1 SPARQL

SPARQL is an RDF query language and a W3C Recommendation, where RDF is a directed, labeled graph data format for representing information in the web [12, 20]. SPARQL contains capabilities for querying required and optional graph patterns along with their conjunctions and disjunctions [21].

A SPARQL graph pattern is defined inductively from triple patterns. Given disjoint infinite sets of IRIs - Internationalised Resource Identifiers - (I), blank nodes (B), literals (L), and variables (V), a triple pattern is defined as an instance of $(I \cup B \cup V)(I \cup V)(I \cup B \cup L \cup V)$ denoted by $IBV \times IV \times IBLV$.

In this work we focus on the conjunctive SPARQL fragment, including only BGPs (basic graph patterns), which can be defined abstractly as $q ::= t \mid q \text{ AND } q'$ where t is a triple pattern.

3.2 ShEx

ShEx (or Shape Expressions) is intended to be an RDF constraint language. Logical operators in Shape Expressions such as grouping, conjunction, disjunction and cardinality constraints, are defined to make as closely as possible to their counterparts in regular expressions and grammar languages like BNF [23]. A Shape Expression has a name and describes the constraints associated with a subject RDF node. For example, Listing 2 shows a definition of a shape in ShEx, where `ex:name` and `ex:phone` are predicates and `xsd:string` is a basic type. This definition means that for an RDF node n to satisfy the requirements of `Shape1`, there must exist two outgoing edges starting from n and labeled with predicates `ex:name` and `ex:phone`, each one being connected to exactly one object node. The objects corresponding to these predicates must be of type `xsd:string`. In general, a ShEx

document may define several shapes, built with more sophisticated logical operators, which we now define more formally. For a more complete definition of ShEx please refer to [23].

```
<Shape1> {
  ex:name xsd:string,
  ex:phone xsd:string }
```

Listing 2: Sample Shape in ShEx

3.2.1 Abstract Syntax of the Considered ShEx Fragment.

Given a finite set of edge labels Σ and a finite set of types Γ , we define a shape expression e over $\Sigma \times \Gamma$ as follows:

$$e ::= \epsilon \mid \Sigma \times \Gamma \mid e^+ \mid (e|e') \mid (e||e')$$

where “ $|$ ” denotes disjunction, “ $||$ ” denotes unordered concatenation, and “ $+$ ” denotes repetition for a positive number of times. From this definition we also further define the following operators as macros:

- $e^? := (\epsilon \mid e)$ (optional)
- $e^* := (\epsilon \mid e^+)$ (unordered Kleene star)
- $e^{[m;n]}$ (e repeated i times with i in the interval from m to n)

which are also parts of the ShEx syntax. In the sequel we write $a :: t$ as a shorthand for $(a, t) \in \Sigma \times \Gamma$.

A shape expression schema (ShEx), or simply a schema, is a tuple $S = (\Sigma, \Gamma, \delta)$, where Σ is a finite set of edge labels, Γ is a finite set of types, and δ is a type definition function that maps elements of Γ to shape expressions e over $\Sigma \times \Gamma$. If δ is not defined for some type $t \in \Gamma$, the default definition is $\delta(t) = \epsilon$.

We notice that a ShEx shape (or simply a shape) is itself a type. While a shape is considered as a user-defined type, more generally a type may also be a basic built-in type (like `xsd:string` in the concrete syntaxes of ShEx).

3.2.2 Semantics of ShEx [23]. Semantically, an RDF graph is valid against a ShEx schema if it is possible to assign types to the nodes of the graph in a manner that satisfies the type definitions of the schema. We assume a fixed graph $G = (V, E)$ which resembles an RDF graph, and a fixed schema $S = (\Sigma, \Gamma, \delta)$. A typing of G w.r.t. S is a function $\lambda : V \rightarrow 2^\Gamma$ that associates with every node of G a set of types. Next, the conditions that a typing needs to satisfy are identified. Given a typing λ and a node $n \in V$ we define the neighborhood-typing of n w.r.t. λ as bag (i.e. multiset) over $\Sigma \times 2^\Gamma$ as $neighborTyping_G^\lambda(n) = \{|a :: \lambda(m) \mid (n, a, m) \in E|\}$. We note by $L(e)$ the bag language of a shape expression e , i.e. $L(e)$ is the set of bags allowed by the language of e . Now, λ is a valid typing of S on G if and only if every node satisfies the type definitions of its associated type i.e., for every $n \in V$, $neighborTyping_G^\lambda(n) \in L(\delta(t))$, for all $t \in \lambda(n)$.

3.2.3 Preliminary Definitions. In accordance with the aforementioned abstract syntax and semantics, we define the following shape expression inclusion relation which we use in the rest of the paper.

DEFINITION 1. Given a shape expression e , a predicate p and a ShEx shape s . The inclusion relation $(p, s) \in e$ is defined inductively as follows:

- $(p, s) \in e$ if $e = (p, s)$
- $(p, s) \in e$ if $e = e_1^+$ and $(p, s) \in e_1$
- $(p, s) \in e$ if $e = e_1 \mid e_2$, and $(p, s) \in e_1$ or $(p, s) \in e_2$
- $(p, s) \in e$ if $e = e_1 || e_2$, and $(p, s) \in e_1$ or $(p, s) \in e_2$

We also define the following shape expression optional condition.

DEFINITION 2. Given a shape expression e , a predicate p and a ShEx shape s , we say that the atomic shape expression (p, s) is optional in e , written as $(p, s) \in_{opt} e$ if:

- $e = e_1 \mid e_2$ and $(p, s) \notin e_1$ and $(p, s) \in e_2$, OR
- $e = e_1 \mid e_2$ and $(p, s) \in e_1$ and $(p, s) \notin e_2$, OR
- $e = e_1 \mid e_2$ and $(p, s) \in_{opt} e_1$ or $(p, s) \in_{opt} e_2$, OR
- $e = e_1 || e_2$ and $(p, s) \in_{opt} e_1$ and $(p, s) \in_{opt} e_2$, OR
- $e = e_1^+$, and $(p, s) \in_{opt} e_1$

4 WELL FORMED DATA-SCHEMA PAIRS

Using RDF as a data format often raises a number of data modelling issues for which choices must be made. The same information might end up being represented in different ways according to the designer choices. In addition, different point of views might be considered for processing the same graph data, yielding to different sets of constraints for the same information. Thus, different ShEx schemas – all of which correctly and usefully describe different aspects of the same data graph – might be suggested. Accordingly, we introduce a notion of well-formed data-schema pairs. The set of rules that will be introduced on the data-schema pairs in this section will help us to better identify efficient SPARQL query designs, by static analysis of the schema.

The rules for well-formation guarantee that the necessary information needed for our ranking can be deduced from the ShEx schema, yet the ranking procedure is not deterministic. For some shapes, the relations attached to them are not indicative for the selectivity of those shapes. We also define in this section the schema formation rules that makes our shape ranking procedure deterministic.

DEFINITION 3 (WELL-FORMED DATA-SCHEMA PAIR). A data-schema pair (G, S) is well-formed if and only if the following rules hold.

- (1) **Cardinality rule:** Every m -to- n relation between two schema shapes in S , where $m > n$ or m is not bound, is modelled from the m -sided shape to the n -sided shape.
- (2) **Shape distinction rule:** For every 4 schema shapes $s_1, s_2, so_1, so_2 \in S$ (not necessarily distinct), and for every 2 predicates p_a and p_b (not necessarily distinct), so_1 and so_2 are distinct if the following conditions hold:
 - $(p_a, so_1) \in \delta(s_1)$
 - $(p_b, so_2) \in \delta(s_2)$
 - O_1 is the set of nodes of G which occur as objects of p_a whose subject belongs to the shape s_1

- O_2 is the set of nodes of G which occur as objects of p_b whose subject belongs to the shape s_2
- $O_1 \cap O_2 = \emptyset$

The well-formation rules impose a schema design that gives preference to some constraints among others, all of which are respected in the data-schema pair, but it does not force new constraints to be added. Well-formed data-schema pairs provide the maximal set of desired information that can be inferred from the ranking procedure described in Sect. 6, without adding new constraints.

Although the well-formation rules are sufficient for optimization of real life data-schema examples, the ranking system is not totally deterministic. A restrictive set of rules on a data-schema pair that make our ranking system deterministic are given in the following definition.

DEFINITION 4 (RANKING-DETERMINISTIC DATA-SCHEMA PAIR). A data-schema pair (G, S) is ranking-deterministic if and only if the following rules hold.

- (1) **Well-formedness:** (G, S) is well-formed.
- (2) **Cardinality rule:** There is no closure cardinality $(+, *)$ in S .
- (3) **Shape distinction rule:** For every 3 shapes $s_o, s_1, s_2 \in S$ (not necessarily distinct), if there exist 2 predicates p_a and p_b (not necessarily distinct) where $(p_a, s_o) \in \delta(s_1)$ and $(p_b, s_o) \in \delta(s_2)$, then s_1 and s_2 refers to the same shape, and p_a and p_b refers to the same predicate.
- (4) **Data nodes isolation rule:** For every data IRI instance d , every 2 shapes $s_1, s_2 \in S$, and every predicate p , if $(p, s_2) \in \delta(s_1)$ and d belongs to the shape s_2 , then there exists a data IRI instance d' such that the RDF triple $\langle d', p, d \rangle \in G$.

In the following subsections (4.1, 4.2, and 4.3), we give examples and additional descriptions of the well-formedness and ranking-deterministic rules, aiding to understand how they contribute to our ranking procedure.

4.1 Cardinality

EXAMPLE 1. Assume we want to model a schema describing the relation between students and schools. If we know that the relation in the data between schools and students will be 1-to-many, then the following two schema examples are legitimate, but only the first one is well formed w.r.t. the data.

Schema proposition 1: (Well Formed)

```
<Student> { :name xsd:string, :school @<School> }
<School> { :name xsd:string }
```

Schema proposition 2: (Not Well Formed)

```
<Student> { :name xsd:string }
<School> { :name xsd:string, :student @<Student> + }
```

As it is evident from Example 1, the *well-formation cardinality rule* tries to avoid the usage of positive and Kleene closures $(+, *)$. Formally, the semantics of the two proposed schemas are different. Schema proposition 1 is more restrictive. Schema proposition 2 misses the restriction that a student should belong to 1 and only 1 school, although it is still an

acceptable schema even if this restriction is inherent in the data.

Indeed, the *well-formation cardinality rule* helps us to determine the relative quantity of shape occurrences in the data. For example Schema proposition 1 allows us to know that the $\langle Student \rangle$ instances definitely occur in the data more than $\langle School \rangle$ instances.

4.2 Shape Distinction

A shape in a ShEx schema can be as general as allowing any node in any RDF graph to belong to it. The more the shape has restrictions, the more it describes a specific type of nodes. The *well-formation shape distinction rule* puts restrictions on shapes that seem to be too general that they surely miss expressing some constraints that are inherent in the data.

EXAMPLE 2. Assume we want to model a schema describing the relation between students and researchers to their corresponding schools and research companies. Knowing that schools are not research companies, then the following two schema examples are legitimate, but only the first one is well formed w.r.t. the data.

Schema proposition 1: (Well Formed)

```
<Student> { :name xsd:string, :school @<School> }
<Researcher> { :name xsd:string, :company @<Company> }
```

Schema proposition 2: (Not Well Formed)

```
<Student> { :name xsd:string, :school @<Establishment> }
<Researcher> { :name xsd:string, :company @<Establishment> }
```

In Example 2, Schema proposition 2 will not allow us to determine the relative quantity of $\langle Student \rangle$ instances to those of $\langle Establishment \rangle$ instances in the data, while with Schema proposition 1 we are sure that the quantity of $\langle Student \rangle$ instances are more than that of $\langle School \rangle$ instances.

4.3 Data Nodes Isolation

The *data nodes isolation rule* for deterministic ranking states that a data instance shall not be isolated from other data instances unless isolation is required by the given schema.

5 SHAPE RELATION GRAPH

In this section we define a shape graph representation that we use to assign ranks to shapes in Sect. 6. A shape relation graph is a graphical representation focusing on the relations existing between shapes in a ShEx document. It is an intermediate structure that will be used later on for selectivity estimation analyses.

DEFINITION 5 (SHAPE RELATION GRAPH). Given a ShEx document S , we define a shape relation graph $G = SRG(S)$ as a tuple (N, E) of set of nodes N , each corresponding to a ShEx shape, and an labelled directed relation E between nodes such that:

- $E(n_1, x, n_2)$ defines an edge from n_1 to n_2 labeled with x .
- Given any two nodes $n_1, n_2 \in N$, and any predicate p , then $E(n_1, p, n_2)$ if and only if $(p, n_2) \in \delta(n_1)$ and $(p, n_2) \notin_{opt} \delta(n_1)$

- Given any two nodes $n_1, n_2 \in N$, and any predicate p , then $E(n_1, p^\epsilon, n_2)$ if and only if $(p, n_2) \in_{opt} \delta(n_1)$

Figure 6 shows the *shape relation graph* of a real life schema used in our experimentation (Sect. 7). User-defined types are shown as ovals while built-in types (like `xsd:string`) and IRIs are shown as rectangles.

We define the set of root nodes of a *shape relation graph*.

DEFINITION 6. Given a shape relation graph $G(N, E)$, we define $\mathcal{R}(G)$ as the set of all root nodes of G . A node $s \in N$ is considered a root node if and only if it has no incoming edges in G .

We also define the set of cycles of a *shape relation graph*.

DEFINITION 7. Given a shape relation graph $G(N, E)$, we define $\mathcal{C}(G)$ as the set of all cycles in G . A cycle is a subgraph of shape relation graph. A subgraph $C(N_C, E_C)$ of G is a cycle if and only if the set of edges E_C defines a directed path that starts and ends with the same node $n \in N$, and N_C is the set of all nodes that can be visited by the set of edges E_C .

6 RANKING

In order to decide the order of execution of query triple patterns, we assign them ranks inferred from the analysis of the ShEx document. These ranks are based on two main concepts: 1) The hierarchical relations between ShEx shapes, and 2) The predicate distributions among ShEx shapes.

The first concept gives rankings to shapes, and the second concept gives ranking to predicates. The ranking of query triple patterns is based on the product of both rankings together.

6.1 Hierarchical Relations between ShEx Shapes

In ShEx, the definition of a shape may be based on other shapes defined in the same schema. This notion, called shape inclusion, is explicitly represented by the edges of the *shape relation graph* defined in Sect. 5. Such edge relations between shapes allow us to infer information about the relative frequency of data corresponding to these shapes.

Consider Schema proposition 1 in of Example 1. Representing it as a *shape relation graph*, $\langle School \rangle$ is a child of $\langle Student \rangle$. Each student in the data should have exactly one registered school, and multiple students may be registered in the same school according to the schema. Such a relation between shapes allows us to know that a student instance occurs more in the database than a school instance. Actually the number of schools is at most equivalent to the number of students, where this is a worst case assumption - each student has a unique school. It is evident that this is an extreme case that should not be considered as an average distribution. Thus, it is important to study the hierarchical relations between ShEx shapes. In the example we give the $\langle School \rangle$ shape a priority ranking, since we know that they occur less than the $\langle Student \rangle$ shape, and thus rendering variables corresponding to it more selective.

Concerning cardinality, we notice that a higher cardinality is independent on the actual number of data instances of a shape. For example, if we have `:registeredIn @<School> {1,3}` instead of `:registeredIn @<School>`, that does not necessarily means an increase in the number of schools; the same set of schools may apply in both cases. For the ranking system it is sufficient to consider the relation structure rather than the structure and cardinalities together, and that is why we ignore explicit cardinalities of edges in the *shape relation graph* defined in Sect. 5.

The ranking procedure we propose starts from the root shapes (root nodes as defined in Definition 6). A root shape will have a ranking of 1. Going down through the descendant shapes from the root shape the ranking increases. If there are two (or more) incoming edges to a shape, the lower ranking is transferred. A problem in such a procedure is when there is a cycle between shapes in the graph representation of the schema, that means that the ranking will propagate forever. In such case there is no preference for any of the shapes in the cycles, and all of them must have the same ranking. In some cases, a cycle has an optional relation(s) within it, given by the cardinalities “?”, “*”, or “{0,n}”. In such case, we know that a cut in the cycle can only occur at these points. For asserting the strength of normal relations against such optional relations, the preference for ranking is to actually cut the cycles at these points and apply the ranking system by avoiding such kind of cycles.

Now we formally define all the procedures described.

6.1.1 Schema Graph Adjustment. First, given the *shape relation graph* G of a ShEx schema, we modify it by detecting optional relations and cycles.

- (1) For each cycle $C_i(N_i, E_i) \in \mathcal{C}(G)$:
 - For all predicate p , if there exist nodes $n_1, n_2 \in N_i$ such that there exists an edge $E(n_1, p^\epsilon, n_2)$, then remove this edge. Let the new resulting graph be G_{nor} .
- (2) For each cycle $C_i(N_i, E_i) \in \mathcal{C}(G_{nor})$:
 - Merge all the nodes $x \in N_i$ into a single node c_i .

6.1.2 Schema Shapes Ranking. Now let the output of the *Schema Graph Adjustment* be $G_{adj}(N, E)$. We define the ranking function $\delta_S(x)$ for all $x \in N$ as follows:

- (1) For each node $r \in \mathcal{R}(G_{adj})$, $\delta_S(r) = 1$
- (2) For each node $r \in \mathcal{R}(G_{adj})$, apply the procedure $P(r)$ defined next.

Given a node $x \in N$, the procedure $P(x)$ is defined as follows:

- For each $s \in N$, and for each predicate p where there exists an edge $E(x, p, s)$
 - (1) If $\delta_S(s)$ is not initialised: $\delta_S(s) = \delta_S(x) + 1$
 - (2) If $\delta_S(x) + 1 < \delta_S(s)$, then $\delta_S(s) = \delta_S(x) + 1$
 - (3) Apply $P(s)$

Finally, we transmit the cycle rankings to the original nodes.

- For each cycle $C_i(N_i, E_i)$ in G_{nor} :
 - For each node $x \in N_i$, $\delta_S(x) = \delta_S(c_i)$

6.2 Predicate Distributions Among ShEx Shapes

In the previous section we ranked shapes according to their relative frequency of occurrences based on relations between them. Such a ranking is not sufficient for deciding rankings of triple patterns in a query since such ranking is also affected by the uniqueness versus globality of predicates within shapes.

Given a predicate p used in the shapes $s_1, s_2, \dots, s_n$. The ranking of p within a shape s , denoted as $\delta_P(p, s)$ is defined as follows.

$$\begin{aligned}\delta_P(p, s) &= \frac{\delta_S(s)}{\delta_S(s_1) \times \delta_S(s_2) \times \dots \times \delta_S(s_n)}, \text{ if } s \in \{s_1, s_2, \dots, s_n\} \\ \delta_P(p, s) &= \frac{1}{\delta_S(s_1) \times \delta_S(s_2) \times \dots \times \delta_S(s_n)}, \text{ otherwise}\end{aligned}$$

The previous formula works by reducing the ranking of a predicate when it is more global, i.e. when it is used with more shapes. With more shapes the factors in the denominator will increase and thus reducing the overall ranking. Such predicates are frequent, they are used every where, and this means there will be a large set of nodes in the database associated with this predicate. The ranking system tends to leave such predicates to be executed lastly, and that is why the modelled function reduces its ranking. On the other hand, if a predicate is unique for a certain shape, its ranking tends to be bigger by reducing the number of denominators to only one, which is the shape it corresponds to.

We notice that if a predicate p corresponds to only one shape s_m , then the ranking corresponding to it will be always 1, where this value represents the highest ranking possible.

$$\delta_P(p, s_m) = \frac{\delta_S(s_m)}{\delta_S(s_m)} = 1$$

On the other hand, the lowest possible ranking is when the predicate p is used globally in all the shapes defined in the ShEx document, and particularly when the shape considered for the current ranking is a root node, which have the lowest possible shape ranking of 1, and the denominator is the largest possible which is the product of all the shape rankings.

$$\delta_P(p, s_{root}) = \frac{1}{\prod_{\forall i, \text{ where } s_i \text{ is a shape}} \delta_S(s_i)}$$

6.3 SPARQL Query Triple Rankings

Now our purpose is to rank the triple patterns given a BGP query. Triple patterns with higher ranking will be executed first. Before ranking triples, we need to validate the BGP against the ShEx document, and for each subject in the triple patterns the ShEx validator will decide to which shapes this subject may belong. A subject may belong to multiple shapes at the same time. Thus, for each subject s occurring in the triple patterns we have a set $C(s)$ of candidate shapes for s . For convenience, given a triple pattern t , we define $C(t) = C(s)$ if s is a subject of t . We also define $p(t)$ as the predicate of the triple t .

To define the triple ranking function, we use the two ranking functions δ_S and δ_P defined previously.

Given a BGP B and a triple pattern $t \in B$, we define the ranking of the triple t , denoted by $\delta_T(t)$, as follows:

$$\delta_T(t) = Avg \left[\delta_S(S_i) \times \delta_P(p(t), S_i) \right]_{\forall i, S_i \in C(t)}$$

For a given triple t , the previous function is the average of the product $\delta_S \times \delta_P$ by considering all the possible candidate shapes for the subject of t .

7 EVALUATION

We prepared an experiment that proves the soundness of the optimization procedure described in this paper. This experiment uses a real ShEx schema example. We generate data according to the Social Network Benchmark (SNB) schema of the Linked Data Benchmark Council (LDBC) [5]. The data and the workload tests are generated by gMark [3]. gMark is a graph and query workload generator based on an input schema. Our experiment demonstrates 4 different datasets of different sizes, thus we show 4 different charts corresponding to the datasets, and therefore allowing to further comment on the effect of the data size.

7.1 Generated data

Using the gMark tool, we generated 4 datasets, all according to the LDBC SNB schema (check the schema in Fig. 5 and the corresponding well-formed shape relation graph in Fig. 6). The tool allows users to define the dataset size by indicating the number of nodes to be generated, in our case 5M, 30M, 50M, and 100M nodes scenarios are used, corresponding to about 11M, 67M, 113M, and 227M RDF triples respectively.

7.2 Workload

Using gMark, we also generated a workload consisting of 12 conjunctive SPARQL queries based on the LDBC SNB schema. We setup the sizes of the queries such that in each query there are between 6 and 10 triple patterns, and there are between 4 and 6 distinguished variables. The choice of the query size is to allow for structures to form within the schema hierarchy, and not to limit it to simple variable relations. Going beyond the size where such hierarchies form is pointless for our evaluation, yet we give a small range to provide a variety of formation choices. The generated queries are given in the appendix (Section C).

Our purpose is to compare, for each query, the evaluation of the optimized triple patterns order resulting from our method with that of counter part queries, which are just equivalent to the optimized ones but with different order of their triple patterns. In our experiment we generate 50 random permutations for each query of the workload.

7.3 SPARQLGX

An advantage of our optimization technique is that it can be applied on the top of query systems like SPARQLGX [8]. This system in turns is based on SPARK coding with Hadoop underlying infrastructure for evaluating SPARQL queries [24]. SPARQLGX is known to outperform many competitors in the

field concerning conjunctive queries [8]. In our experiment we show how the application of our technique further decreases the average run time for SPARQLGX in the presence of a ShEx document.

SPARQLGX, in the current state, has a basic triple pattern ordering strategy that is based on grouping triple patterns with common join variables together, but the order of the groups themselves is not considered, and thus the ordering is not deterministic for a set of triple patterns; it also depends on their initial ordering given by the input SPARQL query. In our experiments we show that this ordering itself is important for obtaining improved results, yet we show that using our ranking strategy based on the ShEx information further improves the results.

We define 3 systems that are included in our experiment as following:

- **S1:** Is SPARQLGX with its ordering strategy turned off (the system itself provides a configuration that stops reordering triple patterns and keeps the original ordering of the query triple patterns).
- **S2:** Is SPARQLGX with its ordering strategy turned on.
- **Optimised:** Is an extension of SPARQLGX. It extends it with the application of our ordering strategy based on the ShEx information.

7.4 Results

The results of our workload query evaluations with the 4 datasets are presented in Figs. 1, 2, 3, and 4 which are explained as follows:

- **The blue area** is the runtime range of system **S2** concerning the different permutations of each test query.
- **avg(S2)** marks the average runtime of all the permutations of each test query with system **S2**.
- **avg(S1)** marks the average runtime of all the permutations of each test query with system **S1**.
- Finally, **the green bars** shows the runtime of each input query with our **optimised system** that has a single deterministic triple patterns ordering for each test query.

For each query run we set an evaluation timeout, which is a maximal value result ($t_{timeout} = 15, 110, 110, \text{ and } 130$ seconds respectively) near the top of each presented results chart. Some queries timeout, and the query run time is considered equivalent to the timeout value for calculating the average. In our given charts if the average is shown to be at the top of the chart, then this means it is $\geq t_{timeout}$ (timeout value).

We notice that we don't show the runtime range of the queries with system **S1** (as done with system **S2**) since there are always input permutations that times out for the considered test queries, and thus we only show the average for this system.

The results show a faster execution of some queries (Q1, Q9, Q10, Q11, and Q12), while it preserves or slightly improves the execution time of the other queries when run using

our methodology. Some queries do not show significant improvement due to the structure of the query and its selectivity. For example if a query is asking for the results concerning two pair of variables signifying the relation between countries and languages. The results of such query is small and constant, since the number of countries and languages is constant; they do not vary even when the dataset size is exponentially increased, and thus such results are expected for some of the generated queries. Actually these kinds of queries are intentionally generated by gMark for benchmarking purposes (check [3]).

Concerning the dataset sizes, it is clear from the charts that our optimization is less evident when the 5M nodes dataset is compared to the bigger datasets. Compared to system **S2** in Fig. 1, the optimized orderings of queries Q1, Q10, and Q12 showed a slight improvement. In Fig. 2 the improvement is more significant with the latter queries, in addition to the new improvements in queries Q9 and Q11. By further increasing the size of the datasets, the improvement almost stay the same (or precisely it barely increases), which shows a threshold where the gain, although significant, is stabilised.

The average of the improvement of our system (**Optimised**) compared to **avg(S2)** is given as follows:

- **Dataset 5M nodes:** Improvement of queries ranged between 1.2% and 20.5%. The mean improvement of all test queries is 3.8%.
- **Dataset 30M nodes:** Improvement of queries ranged between 1.6% and 87%. The mean improvement of all test queries is 23.5%.
- **Dataset 50M nodes:** Improvement of queries ranged between 1.4% and 84.6%. The mean improvement of all test queries is 25.2%.
- **Dataset 100M nodes:** Improvement of queries ranged between 1.6% and 85.1%. The mean improvement of all test queries is 25.7%.

Figure 1: Comparing ranking-optimised query evaluation with other systems (SNB data 5M nodes)

8 CONCLUSION

We studied a method for SPARQL query optimization based on ranking triple patterns in order to select their execution

Figure 2: Comparing ranking-optimised query evaluation with other systems (SNB data 30M nodes)

Figure 3: Comparing ranking-optimised query evaluation with other systems (SNB data 50M nodes)

Figure 4: Comparing ranking-optimised query evaluation with other systems (SNB data 100M nodes)

order. The originality of our approach is that rankings generated by our system are based on information inferred from a schema expressed in ShEx, which is an emerging schema language for RDF data. To the best of our knowledge, this is the first attempt of leveraging ShEx constraints for SPARQL query optimization.

We first defined a well-formation notion for data-schema pairs that is useful for inferring quantitative information

about data instances. We then defined a procedure for determining rankings. We implemented a prototype of our system on top of the SPARQLGX query evaluation engine, which is known to outperform many competitors in the field. We compared the rankings found by our system, owing to the analysis of ShEx constraints, to the original reordering method of SPARQLGX in terms of query evaluation times, and with datasets of various sizes. Preliminary experimental results indicate that most rankings found by our system lead to improvements in query execution times. This illustrates the interest of considering ShEx constraints for SPARQL query optimization.

A EVALUATION: SCHEMA

Figure 5: Schema (LDBC SNB)

Source: https://github.com/ldbc/ldbc_snb_docs

B EVALUATION: SHAPE RELATION GRAPH

Figure 6: Shape relation graph (LDBC SNB)

C EVALUATION: QUERIES

Query (Q1):

```
SELECT ?x3, ?x2, ?x4, ?x1, ?x0 WHERE {
?x1 ex:pcontainerOf ?x0 .
?x1 ex:phasMember ?x2 .
?x2 ex:pworksAt ?x3 .
?x0 ex:pisSubclassOf ?x4 .
?x5 ex:plikes ?x4 .
?x6 ex:phasMember ?x5 .
?x6 ex:phasMember ?x3 . }
```

Query (Q2):

```
SELECT ?x4, ?x2, ?x0, ?x3, ?x1 WHERE {
?x0 ex:pisLocatedIn ?x8 .
?x8 ex:pknows ?x1 .
?x9 ex:pisLocatedIn ?x1 .
?x9 ex:pgender ?x2 .
?x3 ex:pemail ?x2 .
?x3 ex:phasInterest ?x4 .
?x0 ex:pbrowserUsed ?x5 .
?x6 ex:pname ?x5 .
?x6 ex:phasType ?x7 .
?x4 ex:pisSubclassOf ?x7 . }
```

Query (Q3):

```
SELECT ?x0, ?x3, ?x2, ?x1 WHERE {
?x1 ex:pname ?x0 .
?x2 ex:phasType ?x1 .
?x3 ex:phasInterest ?x2 .
?x3 ex:pbirthday ?x4 .
?x5 ex:pcreationDate ?x4 .
?x5 ex:pisLocatedIn ?x6 .
?x6 ex:pname ?x7 .
?x8 ex:pname ?x7 .
?x9 ex:pisPartOf ?x8 . }
```

Query (Q4):

```
SELECT ?x1, ?x0, ?x2, ?x3 WHERE {
?x1 ex:pemail ?x0 .
?x1 ex:plikes ?x2 .
?x2 ex:pisLocatedIn ?x3 .
?x3 ex:pisPartOf ?x4 .
?x4 ex:pname ?x5 .
?x6 ex:pname ?x5 .
?x7 ex:pstudyAt ?x6 .
?x7 ex:pisLocatedIn ?x8 . }
```

Query (Q5):

```
SELECT ?x2, ?x1, ?x4, ?x0, ?x3 WHERE {
?x0 ex:phasMember ?x9 .
?x9 ex:pname ?x1 .
?x10 ex:phasModerator ?x1 .
?x2 ex:pspeaks ?x10 .
?x2 ex:phasModerator ?x11 .
?x11 ex:pknows ?x3 .
?x3 ex:pknows ?x4 .
?x5 ex:pisLocatedIn ?x0 .
?x5 ex:pgender ?x6 .
?x7 ex:pspeaks ?x6 .
?x8 ex:phasMember ?x7 .
?x8 ex:phasModerator ?x4 . }
```

Query (Q6):

```
SELECT ?x4, ?x2, ?x3, ?x5, ?x0, ?x1 WHERE {
?x1 ex:pname ?x0 .
?x1 ex:pname ?x2 .
?x3 ex:pname ?x2 .
?x4 ex:pisPartOf ?x3 .
?x4 ex:pisPartOf ?x5 .
?x5 ex:pname ?x6 .
?x7 ex:pgender ?x6 .
?x7 ex:pgender ?x8 .
?x9 ex:pname ?x8 .
?x9 ex:pname ?x10 . }
```

Query (Q7):

```
SELECT ?x3, ?x4, ?x5, ?x2, ?x0, ?x1 WHERE {
?x1 ex:pworksAt ?x0 .
?x1 ex:pstudyAt ?x2 .
?x2 ex:pname ?x3 .
?x4 ex:pname ?x3 .
?x4 ex:pname ?x5 .
?x6 ex:pname ?x5 .
?x6 ex:plocationIP ?x7 .
?x8 ex:pbrowserUsed ?x7 . }
```

Query (Q8):

```
SELECT ?x2, ?x1, ?x3, ?x0 WHERE {
?x0 ex:pisLocatedIn ?x1 .
?x2 ex:pisPartOf ?x1 .
?x3 ex:pisLocatedIn ?x2 .
?x3 ex:pgender ?x4 .
?x5 ex:pname ?x4 .
?x5 ex:pname ?x6 . }
```

Query (Q9):

```
SELECT ?x0, ?x3, ?x2, ?x1, ?x4 WHERE {
?x1 ex:pname ?x0 .
?x1 ex:pisLocatedIn ?x2 .
?x2 ex:pname ?x3 .
?x4 ex:pname ?x3 .
?x4 ex:pisLocatedIn ?x5 .
?x5 ex:pisPartOf ?x6 .
?x7 ex:pisPartOf ?x6 . }
```

Query (Q10):

```
SELECT ?x0, ?x2, ?x1, ?x3, ?x5, ?x4 WHERE {
?x1 ex:pisSubclassOf ?x0 .
?x2 ex:phasType ?x1 .
?x3 ex:phasInterest ?x2 .
?x3 ex:pknows ?x4 .
?x5 ex:phasType ?x0 .
?x6 ex:phasInterest ?x5 .
?x6 ex:pknows ?x7 .
?x7 ex:pknows ?x8 .
?x4 ex:phasMember ?x8 . }
```

Query (Q11):

```
SELECT ?x4, ?x5, ?x0, ?x1, ?x3, ?x2 WHERE {
?x0 ex:pname ?x1 .
?x2 ex:pname ?x1 .
?x3 ex:pisLocatedIn ?x2 .
?x4 ex:pisPartOf ?x0 .
?x5 ex:pisPartOf ?x0 .
?x6 ex:pisPartOf ?x0 .
?x7 ex:pisPartOf ?x3 .
?x8 ex:pisPartOf ?x3 .
?x9 ex:pisPartOf ?x3 . }
```

Query (Q12):

```
SELECT ?x1, ?x0, ?x2, ?x4, ?x3 WHERE {
?x1 ex:pisLocatedIn ?x0 .
?x1 ex:pisLocatedIn ?x2 .
?x3 ex:pisLocatedIn ?x2 .
?x3 ex:pisLocatedIn ?x4 .
?x5 ex:pisLocatedIn ?x0 .
?x6 ex:phasCreator ?x5 .
?x6 ex:phasTag ?x7 .
?x4 ex:phasTag ?x7 . }
```

REFERENCES

- [1] Daniel J. Abadi, Adam Marcus, Samuel R. Madden, and Kate Hollenbach. 2007. Scalable Semantic Web Data Management Using Vertical Partitioning. In *Proceedings of the 33rd International Conference on Very Large Data Bases (VLDB '07)*. VLDB Endowment, 411–422. <http://dl.acm.org/citation.cfm?id=1325851.1325900>
- [2] Karl Aberer and Gisela Fischer. 1995. Semantic Query Optimization for Methods in Object-Oriented Database Systems. In *Proceedings of the Eleventh International Conference on Data Engineering (ICDE '95)*. IEEE Computer Society, Washington, DC, USA, 70–79. <http://dl.acm.org/citation.cfm?id=645480.655431>
- [3] G. Bagan, A. Bonifati, R. Ciucanu, G. H. L. Fletcher, A. Lemay, and N. Advokaat. 2017. gMark: Schema-Driven Generation of Graphs and Queries. *IEEE Transactions on Knowledge and Data Engineering* 29, 4 (2017), 856–869.
- [4] Véronique Benzaken, Giuseppe Castagna, Dario Colazzo, and Kim Nguyen. 2013. Optimizing XML Querying Using Type-based Document Projection. *ACM Trans. Database Syst.* 38, 1, Article 4 (April 2013), 45 pages. <https://doi.org/10.1145/2445583.2445587>
- [5] Orri Erling, Alex Averbuch, Josep Larriba-Pey, Hassan Chafi, Andrey Gubichev, Arnau Prat, Minh-Duc Pham, and Peter Boncz. 2015. The LDBC Social Network Benchmark: Interactive Workload. In *Proceedings of the 2015 ACM SIGMOD International Conference on Management of Data (SIGMOD '15)*. ACM, New York, NY, USA, 619–630. <https://doi.org/10.1145/2723372.2742786>
- [6] Javier D. Fernández, Miguel A. Martínez-Prieto, Claudio Gutiérrez, Axel Polleres, and Mario Arias. 2013. Binary RDF Representation for Publication and Exchange (HDT). *Web Semant.* 19 (March 2013), 22–41. <https://doi.org/10.1016/j.websem.2013.01.002>
- [7] François Goasdoué, Zoi Kaoudi, Ioana Manolescu, Jorge Quiané-Ruiz, Stamatis Zampetakis, et al. 2013. CliqueSquare: efficient Hadoop-based RDF query processing. In *BDA'13-Journées de Bases de Données Avancées*.
- [8] Damien Graux, Louis Jachiet, Pierre Genevès, and Nabil Layaïda. 2016. *SPARQLGX: Efficient Distributed Evaluation of SPARQL with Apache Spark*. Springer International Publishing, Cham, 80–87. https://doi.org/10.1007/978-3-319-46547-0_9
- [9] Jiewen Huang, Daniel J. Abadi, and Kun Ren. 2011. Scalable SPARQL Querying of Large RDF Graphs. *PVLDB*, 4(21). (August 2011), 1123–1134 pages.
- [10] Amit Krishna Joshi, Pascal Hitzler, and Guozhu Dong. 2013. *Logical Linked Data Compression*. Springer Berlin Heidelberg, Berlin, Heidelberg, 170–184. https://doi.org/10.1007/978-3-642-38288-8_12
- [11] HyeonSik Kim, Padmashree Ravindra, and Kemafor Anyanwu. 2017. Type-based Semantic Optimization for Scalable RDF Graph Pattern Matching. In *Proceedings of the 26th International Conference on World Wide Web (WWW '17)*. International World Wide Web Conferences Steering Committee, Republic and Canton of Geneva, Switzerland, 785–793. <https://doi.org/10.1145/3038912.3052655>
- [12] Markus Lanthaler, Richard Cyganiak, and David Wood. 2014. *RDF 1.1 Concepts and Abstract Syntax*. W3C Recommendation. W3C. <http://www.w3.org/TR/2014/REC-rdf11-concepts-20140225/>.
- [13] Kisung Lee and Ling Liu. 2013. Scaling Queries over Big RDF Graphs with Semantic Hash Partitioning. *Proc. VLDB Endow.* 6, 14 (Sept. 2013), 1894–1905. <https://doi.org/10.14778/2556549.2556571>
- [14] Thomas Neumann and Gerhard Weikum. 2008. RDF-3X: A RISC-style Engine for RDF. *Proc. VLDB Endow.* 1, 1 (Aug. 2008), 647–659. <https://doi.org/10.14778/1453856.1453927>
- [15] Jeff Z. Pan, José Manuel Gómez Pérez, Yuan Ren, Honghan Wu, Haoen Wang, and Man Zhu. 2015. *Graph Pattern Based RDF Data Compression*. Springer International Publishing, Cham, 239–256. https://doi.org/10.1007/978-3-319-15615-6_18
- [16] N. Papailiou, I. Konstantinou, D. Tsoumakos, P. Karras, and N. Koziris. 2013. H2RDF+: High-performance distributed joins over large-scale RDF graphs. In *2013 IEEE International Conference on Big Data*. 255–263. <https://doi.org/10.1109/BigData.2013.6691582>
- [17] Minh-Duc Pham, Linnea Passing, Orri Erling, and Peter Boncz. 2015. Deriving an Emergent Relational Schema from RDF Data. In *Proceedings of the 24th International Conference on World Wide Web (WWW '15)*. International World Wide Web Conferences Steering Committee, Republic and Canton of Geneva, Switzerland, 864–874. <https://doi.org/10.1145/2736277.2741121>
- [18] Eric Prud'hommeaux, Jose Emilio Labra Gayo, and Harold Solbrig. 2014. Shape Expressions: An RDF Validation and Transformation Language. In *Proceedings of the 10th International Conference on Semantic Systems (SEM '14)*. ACM, New York, NY, USA, 32–40. <https://doi.org/10.1145/2660517.2660523>
- [19] Michael Schmidt, Michael Meier, and Georg Lausen. 2010. Foundations of SPARQL Query Optimization. In *Proceedings of the 13th International Conference on Database Theory (ICDT '10)*. ACM, New York, NY, USA, 4–33. <https://doi.org/10.1145/1804669.1804675>
- [20] Guus Schreiber and Yves Raimond. 2014. *RDF 1.1 Primer*. W3C Note. W3C. <http://www.w3.org/TR/2014/NOTE-rdf11-primer-20140624/>.
- [21] Andy Seaborne and Steven Harris. 2013. *SPARQL 1.1 Query Language*. W3C Recommendation. W3C. <http://www.w3.org/TR/2013/REC-sparql11-query-20130321/>.
- [22] Giorgos Serfiotis, Ioanna Koffina, Vassilis Christophides, and Val Tannen. 2005. Containment and Minimization of RDF/S Query Patterns. In *Proceedings of the 4th International Conference on The Semantic Web (ISWC'05)*. Springer-Verlag, Berlin, Heidelberg, 607–623. https://doi.org/10.1007/11574620_44
- [23] Slawek Staworko, Iovka Boneva, Jose E. Labra Gayo, Samuel Hym, Eric G. Prud'hommeaux, and Harold Solbrig. 2015. Complexity and Expressiveness of ShEx for RDF. In *18th International Conference on Database Theory (ICDT 2015) (Leibniz International Proceedings in Informatics (LIPIcs))*, Marcelo Arenas and Martín Ugarte (Eds.), Vol. 31. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, Dagstuhl, Germany, 195–211. <https://doi.org/10.4230/LIPIcs.ICDT.2015.195>
- [24] Matei Zaharia, Mosharaf Chowdhury, Tathagata Das, Ankur Dave, Justin Ma, Murphy McCauly, Michael J. Franklin, Scott Shenker, and Ion Stoica. 2012. Resilient Distributed Datasets: A Fault-Tolerant Abstraction for In-Memory Cluster Computing. In *Presented as part of the 9th USENIX Symposium on Networked Systems Design and Implementation (NSDI 12)*. USENIX, San Jose, CA, 15–28. <https://www.usenix.org/conference/nsdi12/technical-sessions/presentation/zaharia>
- [25] X. Zhang, L. Chen, Y. Tong, and M. Wang. 2013. EAGRE: Towards scalable I/O efficient SPARQL query evaluation on the cloud. In *2013 IEEE 29th International Conference on Data Engineering (ICDE)*. 565–576. <https://doi.org/10.1109/ICDE.2013.6544856>
- [26] Lei Zou, Jinghui Mo, Lei Chen, M. Tamer Özsu, and Dongyan Zhao. 2011. gStore: Answering SPARQL Queries via Subgraph Matching. *Proc. VLDB Endow.* 4, 8 (May 2011), 482–493. <https://doi.org/10.14778/2002974.2002976>