



HAL
open science

Optimising SPARQL Query Evaluation in the Presence of ShEx Constraints

Abdullah Abbas, Pierre Genevès, Cécile Roisin, Nabil Layaïda

► **To cite this version:**

Abdullah Abbas, Pierre Genevès, Cécile Roisin, Nabil Layaïda. Optimising SPARQL Query Evaluation in the Presence of ShEx Constraints. BDA 2017 - 33ème conférence sur la “ Gestion de Données - Principes, Technologies et Applications ”, Nov 2017, Nancy, France. pp.1-12. hal-01673067

HAL Id: hal-01673067

<https://inria.hal.science/hal-01673067v1>

Submitted on 28 Dec 2017

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Optimising SPARQL Query Evaluation in the Presence of ShEx Constraints

Optimisation de l'évaluation de requêtes SPARQL en présence de contraintes ShEx

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 optimise the evaluation of conjunctive SPARQL queries, on RDF graphs, by taking advantage of ShEx constraints. Our optimisation 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 optimisation. We then define our optimisation method by exploiting information extracted from a ShEx schema. We finally report on evaluation results performed showing the advantages of applying our optimisation on the top of an existing state-of-the-art query evaluation system.

[Résumé]. Nous nous intéressons dans ce travail à l'optimisation de l'évaluation de requêtes SPARQL conjonctives sur des graphes RDF en présence d'un schema de ShEx (Shape Expressions). ShEx est un langage qui sert à exprimer des contraintes sur les graphes RDF et notre optimisation s'appuie sur ces contraintes pour classer les triple patterns composant la requête. Notre plan optimisé d'exécution de la requête est déduit de ce classement. Dans un premier temps, nous définissons un ensemble de schémas ShEx bien formés possédant des caractéristiques intéressantes pour l'optimisation de requêtes SPARQL. Ensuite, nous introduisons une méthode d'optimisation qui exploite les informations extraites du schéma ShEx. Enfin, nous présentons le résultat de nos expérimentations mettant en évidence les gains de performance obtenus par notre optimisation dans un évaluateur distribué performant.

KEYWORDS

SPARQL optimisation, SPARQL evaluation, ShEx, Shape Expressions, RDF constraints

1 INTRODUCTION

The Shape Expressions (ShEx) language [20] allows to describe constraints on RDF graph structures [22]. 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 [23] can be optimized in the presence of ShEx constraints. We propose a method for optimizing the order or evaluation of subqueries, by taking advantage of the information on the data described in ShEx. While SPARQL query optimisation 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. 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. Optimisation 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 $?x$.

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

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

Listing 1: SPARQL query example

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 optimised 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 optimisation method by exploiting information extracted from a ShEx schema. To the best of our knowledge, this is the first work addressing SPARQL query optimisation 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 [10]. SPARQLGX already implements various query optimisation techniques including reordering triple patterns [10], but without considering schema constraints. We show that our technique further improves the efficiency of query execution times from 5% up to 17%.

Outline. The rest of the paper is organized as follows: In Sect. 2 we review the closest related works. In Sect. 3 we introduce some preliminaries necessary for understanding the rest of the paper. In Sect. 4 we define well-formation rules for a data-schema pair that are of interest for our optimisation process. In Sect. 5 we define a graph representation of a ShEx schema. In Sect. 6 we formally define our optimisation process. In Sect. 7 we report on experimental results with our optimisation technique. Finally, we conclude in Sect. 8.

2 RELATED WORKS

Query optimisation 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 [9, 11, 15, 16, 18] are based on techniques that mainly focus on join optimisations 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 optimised query processing, is vertical data partitioning [1] (which is already used in SPARQLGX [10]).

The works in [6, 12, 17, 19] also provide query optimisation techniques, by proposing new data representations that are more compact after scanning the data. Additionally, in [13, 27, 28] 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 optimisations are computed by analysing only a schema (whose size is much smaller than the actual data); and our optimisations 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 optimisation for object-oriented databases is considered. In [21, 24] the authors, as in our work, consider query optimisation 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 [10] 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, optimised 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 [10]. For this reason we implemented and tested our optimisation technique on the top of this system. In the experimental part, we compare the results that we obtain with and without our optimisations to the results obtained with SPARQLGX.

A novelty of our work is that it provides a new optimisation technique that is based on a simple analysis of a given schema. Optimisations 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 optimisations techniques like [1, 10].

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 [14, 22]. SPARQL contains capabilities for querying required and optional graph patterns along with their conjunctions and disjunctions [23].

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 $(IUBUV)(IUV)(IUBULUV)$ 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 [25]. 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 [25].

```
<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 built-in type (like `xsd:string` in the concrete syntaxes of ShEx).

3.2.2 Semantics of ShEx [25]. 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 3.1. 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 $e = e_1 \parallel e_2$, and $(p, s) \in e_1$ or $(p, s) \in e_2$

We also define the following shape expression optional condition.

Definition 3.2. 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
- $e = e_1 \parallel e_2$ 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 modeling 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 4.1 (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 is valid w.r.t. the shape s_1 .
 - O_2 is the set of nodes of G which occur as objects of p_b whose subject is valid w.r.t. 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 optimisation 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 4.2 (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 4.3. 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 4.3, 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 4.4. 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 4.4, 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 only on the relations existing between shapes in a ShEx document, discarding cardinalities.

Definition 5.1 (Shape Relation Graph). Given a ShEx document S , we define a shape relation graph $G = \mathcal{SRG}(S)$ 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 (appendix) 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. For visualisation reasons, we also replicate some type nodes in the mentioned figure.

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

Definition 5.2. 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 5.3. 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 4.3. 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 clear 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 5.2). 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.

$$\delta_P(p, s) = \frac{\delta_S(s)}{\delta_S(s_1) \times \delta_S(s_2) \times \dots \times \delta_S(s_n)}, \text{ if } s \in \{s_1, s_2, \dots, s_n\} \text{ shapes for the subject of } t.$$

$$\delta_P(p, s) = \frac{1}{\delta_S(s_1) \times \delta_S(s_2) \times \dots \times \delta_S(s_n)}, \text{ otherwise}$$

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 x occurring in the triple patterns we have a set $C(x)$ of candidate shapes for x . For convenience, given a triple pattern t , we define $C(t) = C(x)$ if x 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 two different two experiments with different setups that show the advantageous effects of the optimisation procedure described in this paper. Both experiments use real ShEx schema examples (Web Index [7] and LDCB SNB [5] respectively) and a benchmark data generator for each of them.

In the first experiment, the queries are designed by us to utilise different combinations of the schema types. In addition to showing the results of this experiment, the main purpose of hand-crafting queries is to allow readers to investigate the queries in a clear way, where the variable names in the queries are indicative to the type they belong to.

In the second experiment, the queries are generated by a specific benchmarking tool (gMark [3]). This setup is more realistic and convincing concerning the obtained results. Queries generated by the gMark tool are of different

structures and sizes, and thus are suitable for benchmarking purposes.

7.1 Experiment 1: With Web Index Schema

In this experiment we generate data according to a real example ShEx schema, using a generator called wiGen (or WebIndex Data Generator) [8]. We make our own set of 12 queries which are of different sizes, and designed to utilise all schema shapes defined in the ShEx schema of wiGen.

7.1.1 Generated data. wiGen is a generator of random data that can be used to benchmark RDF schema languages. The data model used by it is inspired by the WebIndex data, which are data intended as a measure of the World Wide Web's contribution to social, economic and political progress in countries across the world [7]. The ShEx schema utilised by wiGen is represented as a *shape relation graph* in Fig. 6.

The generator allows the user to define the number of instances needed for each shape in the schema. We generated 15M nodes, and the resulting RDF dataset is about 80M RDF triples.

7.1.2 Queries. The set of queries we use for this experimentation is shown in Fig. 7. The variable names in the queries are hand-crafted to indicate a schema type from the schema of Fig. 6 (?d: DataSet, ?s: Slice, ?o: Observation, ?c: Country, ?comp: Computation, ?i: Indicator, ?org: Organization). The focus of our study is on BGP queries, and thus all of them are within this fragment. The queries of Fig. 7 are given in their optimised form i.e. with their order of triple patterns computed by our ranking system. Our purpose is to compare each query with counter part queries, which are just equivalent to the original ones, with different order of their triple patterns. The number of permutations for each query depends on the number of triples. For 4 triple patterns there are 24 different permutations, while for 7 triple patterns there are 5040. In our experiment we generate all the permutations if there are less than 50 of them, and otherwise we generate 50 random permutations.

7.1.3 SPARQLGX. An advantage of our optimisation technique is that it can be applied on the top of query systems like SPARQLGX [10]. This system in turns is based on SPARK coding with Hadoop underlying infrastructure for evaluating SPARQL queries [26]. SPARQLGX is known to outperform many competitors in the field concerning conjunctive queries [10]. 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.

7.1.4 Results. The results of our query evaluations are presented in Fig. 1 which is 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 of 200 seconds. Some queries times out, and the query run time is considered 200 seconds for calculating the average. In our given chart if the average is shown to be 80 seconds (the top of the chart), then this means it is ≥ 80 seconds.

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.




Figure 1: Comparing ranking-optimised query evaluation with other systems (WebIndex data)

The average of the improvement of query executions by our system (**Optimised**) compared to **avg(S2)** ranged between 3.7% and 18%. The mean improvement of all test queries

is 9.8%. We notice that in this calculation we did not consider query Q12 which has a theoretical 98% improvement average. The purpose of Q12 was to show that if the query is unsatisfiable according to the ShEx schema (it violates it), then we do not execute it and the time considered is the validation time of the query. In Q12 the last triple pattern has `foaf:homePage` in the predicate position and the string `"homepageOrg988"` in the object position, while the object of `foaf:homePage` should be an IRI according the ShEx schema.

7.2 Experiment 2: With LDBC SNB Schema and gMark Queries

In this experiment we generate data according to the Social Network Benchmark (SNB) schema of the Linked Data Benchmark Council (LDBC). The data and queries are generated by gMark [3]. gMark is a graph and query workload generator based on an input schema. Technically the setup is comparable to that of Experiment 1 except that in Experiment 2 queries are generated by the benchmarking tool rather than being hand-crafted. In addition Experiment 2 is applied on 4 different dataset of different sizes, thus we show 4 different charts corresponding to the datasets, and therefore allowing to further comment on the effect of data size.

7.2.1 Generated data. Using the gMark tool, we generated 4 datasets, all according to the LDBC SNB schema (check the schema in Fig. 8 and the corresponding well-formed shape relation graph in Fig. 9). 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.

7.2.2 Queries. Using gMark, we also generated a set of 12 SPARQL queries based on the LDBC SNB schema. We setup the the sizes of 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.

As in Experiment 1, we generated 50 random triple pattern permutations of each query.

7.2.3 Results. The results of our query evaluation with the 4 datasets are presented in 4 charts (Figs. 2, 3, 4, and 5), where each is similar to the chart described in Experiment 1.

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 optimisation is less evident when the 5M nodes dataset is compared to the bigger datasets. Compared to system **S2** in Fig. 2, the optimised orderings of queries Q1, Q10, and Q12 showed a slight improvement. In Fig. 3 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%.

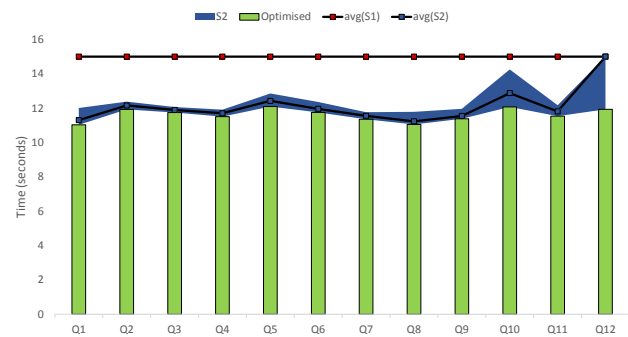


Figure 2: Comparing ranking-optimised query evaluation with other systems (SNB data 5M nodes)

8 CONCLUSION

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




Figure 3: Comparing ranking-optimised query evaluation with other systems (SNB data 30M nodes)




Figure 4: Comparing ranking-optimised query evaluation with other systems (SNB data 50M nodes)

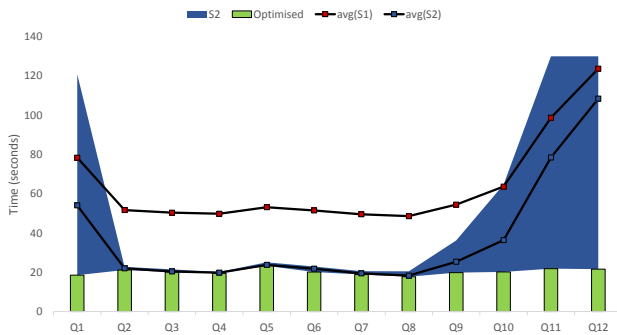


Figure 5: Comparing ranking-optimised query evaluation with other systems (SNB data 100M nodes)

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 optimisation.

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

A EXPERIMENT 1: SHAPE RELATION GRAPH

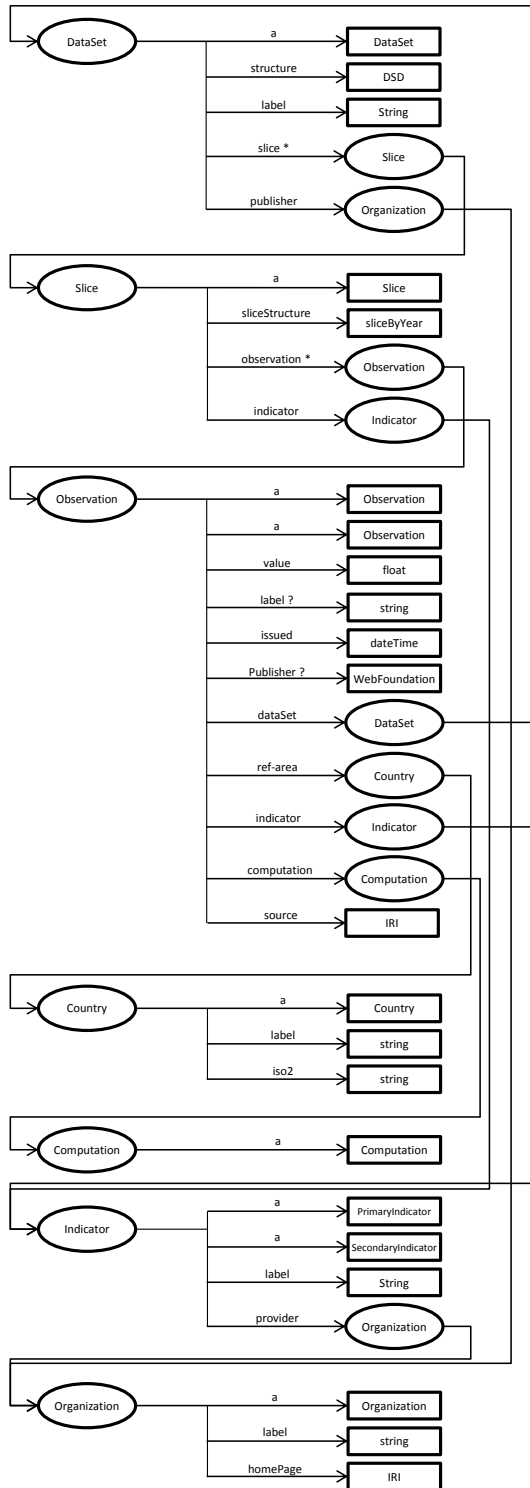


Figure 6: Shape relation graph (WebIndex)

B EXPERIMENT 1: QUERIES

<p>Q1 (optimized ordering)</p> <pre>SELECT ?o, ?d, ?struct WHERE { ?d rdfs:label "dataSet1576" . ?d qb:structure ?struct . ?o qb:dataSet ?d . ?o rdf:type qb:Observation . }</pre>	<p>Q8 (optimized ordering)</p> <pre>SELECT ?d, ?org WHERE { ex:obs4830 qb:dataSet ?d . ?d dct:publisher ?org . ?d rdf:type qb:DataSet . }</pre>
<p>Q2 (optimized ordering)</p> <pre>SELECT ?o, ?d, ?org WHERE { ?org foaf:homepage ex:homepageOrg988 . ?d dct:publisher ?org . ?o qb:dataSet ?d . ?o rdf:type qb:Observation . }</pre>	<p>Q9 (optimized ordering)</p> <pre>SELECT ?d, ?type, ?org WHERE { ex:obs4830 qb:dataSet ?d . ?d dct:publisher ?org . ?d rdf:type ?type . }</pre>
<p>Q3 (optimized ordering)</p> <pre>SELECT ?o, ?d, ?c, ?i, ?comp, ?org WHERE { ?org foaf:homepage ex:homepageOrg988 . ?i wf:provider ?org . ?o cex:indicator ?i . ?o cex:computation ?comp . ?o cex:ref-area ?c . ?o qb:dataSet ?d . ?o rdf:type qb:Observation . }</pre>	<p>Q10 (optimized ordering)</p> <pre>SELECT ?s, ?org, ?i, ?hp WHERE { ?s wf:provider ?org . ?org foaf:homepage ?hp . ?s cex:indicator ?i . ?org rdf:type org:Organization . ?s rdf:type qb:Slice . }</pre>
<p>Q4 (optimized ordering)</p> <pre>SELECT ?s, ?i WHERE { ?i wf:provider ex:org825 . ?s cex:indicator ?i . ?s rdf:type qb:Slice . }</pre>	<p>Q11 (optimized ordering)</p> <pre>SELECT ?s, ?d, ?struct1, ?struct2 WHERE { ?s qb:data ?d . ?d qb:structure ?struct2 . ?s qb:sliceStructure ?struct1 . ?d rdf:type qb:DataSet . ?s rdf:type qb:Slice . }</pre>
<p>Q5 (optimized ordering)</p> <pre>SELECT ?s, ?i WHERE { ?i rdfs:label "indicator2006" . ?s cex:indicator ?i . ?s rdf:type qb:Slice . }</pre>	<p>Q12 (violating)</p> <pre>SELECT ?o, ?d, ?c, ?i, ?comp, ?org WHERE { ?o rdf:type qb:Observation . ?o qb:dataSet ?d . ?o cex:ref-area ?c . ?o cex:indicator ?i . ?o cex:computation ?comp . ?i wf:provider ?org . ?org foaf:homepage "homepageOrg988" . }</pre>
<p>Q6 (optimized ordering)</p> <pre>SELECT ?s, ?i, ?org WHERE { ?org foaf:homepage ex:homepageOrg988 . ?i wf:provider ?org . ?s cex:indicator ?i . ?s rdf:type qb:Slice . }</pre>	<p>Q12 (optimized)</p> <pre>SELECT ?o, ?d, ?c, ?i, ?comp, ?org WHERE { }</pre>
<p>Q7 (optimized ordering)</p> <pre>SELECT ?s, ?type, ?i, ?org WHERE { ?org foaf:homepage ex:homepageOrg988 . ?i wf:provider ?org . ?s cex:indicator ?i . ?s rdf:type ?type . }</pre>	

Figure 7: Hand-crafted experiment queries

C EXPERIMENT 2: SCHEMA

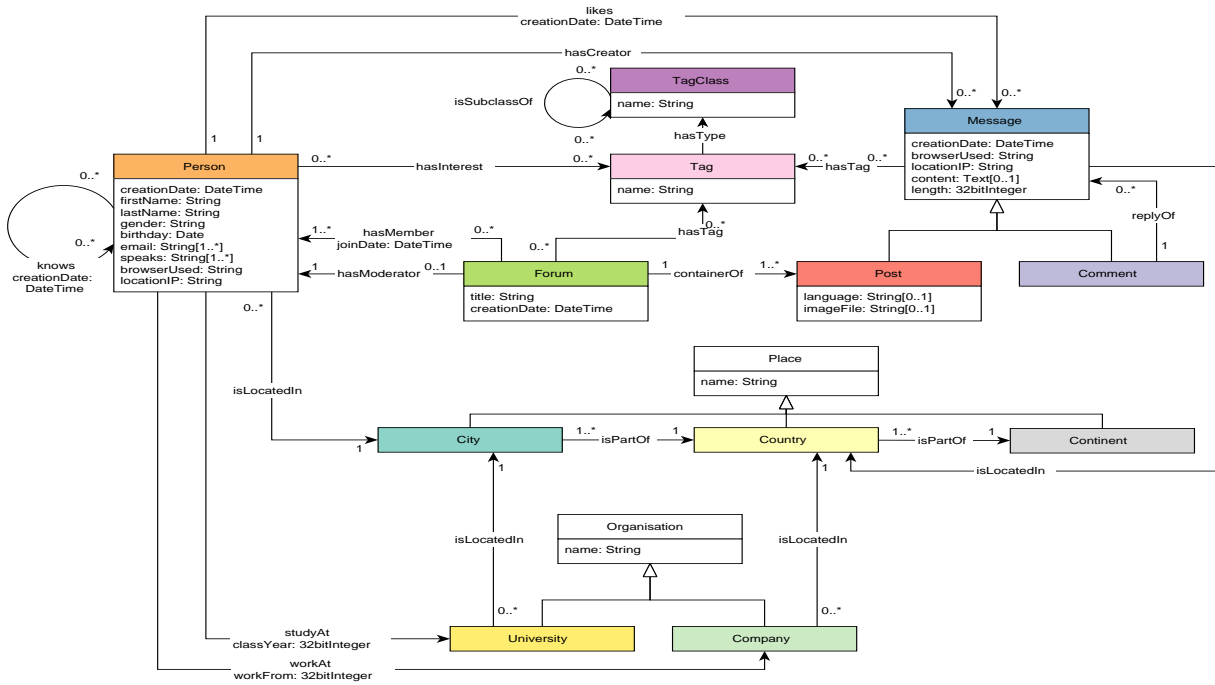


Figure 8: Schema (LDBC SNB)
 Source: https://github.com/ldbc/ldbc-snb_docs

D EXPERIMENT 2: SHAPE RELATION GRAPH




Figure 9: Shape relation graph (LDBC SNB)

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] World Wide Web Foundation. 2014. WEB INDEX. <http://thewebindex.org/>. (2014). [Online; accessed 2-May-2017].
- [8] Jose Emilio Labra Gayo, Eric Prud'Hommeaux, Harold Solbrig, and Jose María Alvarez Rodríguez. 2014. Validating and Describing Linked Data Portals using RDF Shape Expressions. In *1st Workshop on Linked Data Quality (LDQ) (CEUR Workshop Proceedings)*, Magnus Knuth, Dimitris Kontokostas, and Harald Sack (Eds.). Aachen. <http://ceur-ws.org/Vol-1215/#paper-06>
- [9] 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*.
- [10] 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
- [11] Jiewen Huang, Daniel J. Abadi, and Kun Ren. 2011. Scalable SPARQL Querying of Large RDF Graphs. *PVLDB*, 4(21). (August 2011), 1123–1134 pages.
- [12] 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
- [13] 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>
- [14] 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/>.
- [15] 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>
- [16] 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>
- [17] Jeff Z. Pan, José Manuel Gómez Pérez, Yuan Ren, Honghan Wu, Haofen 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
- [18] 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>
- [19] 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>
- [20] 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>
- [21] 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>
- [22] Guus Schreiber and Yves Raimond. 2014. *RDF 1.1 Primer*. W3C Note. W3C. <http://www.w3.org/TR/2014/NOTE-rdf11-primer-20140624/>.
- [23] Andy Seaborne and Steven Harris. 2013. *SPARQL 1.1 Query Language*. W3C Recommendation. W3C. <http://www.w3.org/TR/2013/REC-sparql11-query-20130321/>.
- [24] 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
- [25] 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>
- [26] 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>
- [27] 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>
- [28] 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>