

Recursion in SPARQL

Juan L. Reutter, Adrián Soto, and Domagoj Vrgoč

PUC Chile and Center for Semantic Web Research

Abstract. The need for recursive queries in the Semantic Web setting is becoming more and more apparent with the emergence of datasets where different pieces of information are connecting by complicated patterns. This was acknowledged by the W3C committee by the inclusion of property paths in the latest SPARQL standard. However, it is becoming clear that property paths alone are not enough to capture all recursive queries that the users are interested in. In this paper we propose a general purpose recursion operator to be added to SPARQL, formalize its syntax and develop algorithms for evaluating it in practical scenarios. We also show how to implement it as a plug-in on top of existing systems and test its performance on several real world datasets.

1 Introduction

The Resource Description Framework (RDF) has emerged as the standard for describing Semantic Web data and SPARQL as the main language for querying RDF [13]. After the initial proposal of SPARQL, and with more data becoming available in the RDF format, users found use cases that required exploring the structure of the data in more detail. In particular queries that are inherently recursive, such as traversing paths of arbitrary length, have lately been in demand. This was acknowledged by the W3C committee with the inclusion of property paths in the latest SPARQL 1.1. standard [12], allowing queries to navigate paths connecting two objects in an RDF graph.

However, in terms of expressive power, several authors have noted that property paths fall short when trying to express a number of important properties related to navigating RDF documents (cf. [5,7,23]), and that a more powerful form of recursion needs to be added to SPARQL to address this issue. Consequently, this realization has brought forward a good number of extensions of property paths that offer more expressive recursive functionalities (see e.g. [3,6,15,18,23]), but to the best of our knowledge no attempt to add a general recursion operator to the language has been made.

To illustrate the need for such an operator we consider the case of tracking provenance of Wikipedia articles presented by Missier and Chen in [20]. They use the PROV standard [25] to store information about how a certain article was edited, whom was it edited by and what this change resulted in. Although they store the data in a graph database, all PROV data is easily representable as RDF using the PROV-O ontology [27]. The most common type of information in this RDF graph tells us when an article A_1 is a revision of an article A_2 . This fact is represented by adding a triple of the form $(A_1, \text{prov:wasRevisionOf}, A_2)$

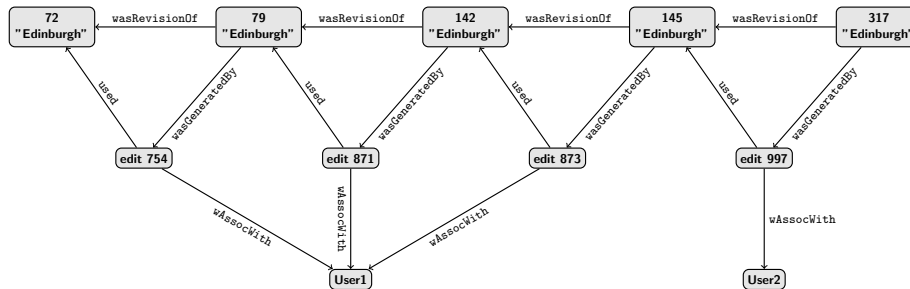


Fig. 1. RDF database of Wikipedia traces. The abbreviation `wAssocWith` is used instead of `wasAssociatedWith` and the `prov:` prefix is omitted from all the properties.

to the database. These revisions are associated to user’s edits with the predicate `prov:wasGeneratedBy` and the edits can specify that they used a particular article with a `prov:used` link. Finally, there is a triple $(E, \text{prov:wasAssociatedWith}, U)$ if the edit E was made by the user U . A snapshot of the data, showing provenance of articles about Edinburgh, is depicted in Figure 1.

A natural query to ask in this context is the history of revisions that were made by the same user: that is all pairs of articles (A, A') such that A is linked to A' by a path of `wasRevisionOf` links and where all of the revisions along the way were made by the same user. For instance, in the graph of Figure 1 we have that the article 145 ”Edinburgh” is a revision of the article 72 ”Edinburgh” and all the intermediate edits were made by User1. Such queries abound in any version control system (note that the PROV traces of Wikipedia articles are the same as tracking program development in svn or Git) and can be used to detect which user introduced errors or bugs, when the data is reliable, or to find the latest stable version of the data. Since these queries can not be expressed with property paths [5, 18], nor by using standard SPARQL functionalities (as provenance traces can contain links of arbitrary length), the need for a general purpose recursive operator seems like a natural addition to the language.

One possible reason why recursion was not previously considered as an integral operator of SPARQL could be the fact that in order to compute recursive queries we need to apply the query to the result of a previous computation. However, typical SPARQL queries do not have this capability as their inputs are RDF graphs but their outputs are mappings. This hinders the possibility of a fixed point recursion as the result of a SPARQL query cannot be subsequently queried. One can avoid this by using CONSTRUCT queries, which output RDF graphs as well, and indeed [16] has proposed a way of defining a fixed point like extension for SPARQL based on this idea.

In this paper we extend the recursion operator of [16] to function over a more widely used fragment of SPARQL and study how this operator can be implemented in an efficient and non-intrusive way on top of existing SPARQL engines. We begin by showing what the general form of recursion looks like and how to evaluate it. We then argue that any implementation of this general form of recursion is unlikely to perform well on real world data, so we restrict it

to the so called *linear recursion*, which is well known in the relational context [1,10]. We then demonstrate that this class of queries can express almost any use case for recursion found in practice. Next, we develop an elegant algorithm for evaluating this class of recursive queries and show how it can be implemented on top of an existing SPARQL system. For our implementation we use Apache Jena framework [14] and we implement recursive queries as an add-on to the ARQ SPARQL query engine. We use Jena TDB version 2.12.1, which allows us not to worry about queries whose intermediate results do not fit into main memory, thus resulting in a highly reliable system. Lastly, we test how this implementation performs on YAGO, LMBD and PROV records of Wikipedia revision history using several natural queries over these datasets.¹

Related work. There are now various systems supporting property paths either in a full capacity [11], or with some limitations that ensure they can be efficiently evaluated, most notable amongst them being Virtuoso [26]. The systems that support full property paths are capable of returning all pairs of nodes connected by a property path specified by the query, while Virtuoso needs a starting point in order to execute the query. We would like to note that recursive queries we introduce are capable of expressing the transitive closure of any binary operator [16] and can thus be used to express property paths and any of their extensions [3, 15, 23]. Regarding attempts to implement a full-fledged recursion as a part of SPARQL, there have been none as far as we are aware. There were some attempts to use SQL recursion to implement property paths [29], or to allow recursion as a programming language construct [4, 21], however none of them view recursion as a part of the language, but as an outside add-on.

2 Preliminaries

RDF Graphs and Datasets. RDF graphs can be seen as edge-labeled graphs where edge labels can be nodes themselves, and an RDF dataset is a collection of RDF graphs. Formally, let \mathbf{I} be an infinite set of *IRIs*². An *RDF triple* is a tuple (s, p, o) from $\mathbf{I} \times \mathbf{I} \times \mathbf{I}$, where s is called the *subject*, p the *predicate*, and o the *object*. An *RDF graph* is a finite set of RDF triples, and an *RDF dataset* is a set $\{G_0, \langle u_1, G_1 \rangle, \dots, \langle u_n, G_n \rangle\}$, where G_0, \dots, G_n are RDF graphs and u_1, \dots, u_n are distinct IRIs. The graph G_0 is called the *default graph*, and G_1, \dots, G_n are called *named graphs* with *names* u_1, \dots, u_n , respectively. For a dataset D and IRI u we define $\text{gr}_D(u) = G$ if $\langle u, G \rangle \in D$ and $\text{gr}_D(u) = \emptyset$ otherwise. Given two datasets D and D' with default graphs G_0 and G'_0 , we define the union $D \cup D'$ as the dataset with the default graph $G_0 \cup G'_0$ and $\text{gr}_{D \cup D'}(u) = \text{gr}_D(u) \cup \text{gr}_{D'}(u)$ for any IRI u . Union of datasets without default graphs is defined in the same way, i.e., as if the default graph was empty.

SPARQL Syntax. We assume the reader is familiar with the syntax and semantics of SPARQL 1.1 query language. However, we do recall two particular

¹ The implementation, test data and complete formulation of used queries can be found in the online appendix available at <http://web.ing.puc.cl/~jreutter/Recsparql.html>.

² For clarity of presentation we do not include literals or blank nodes in our definitions.

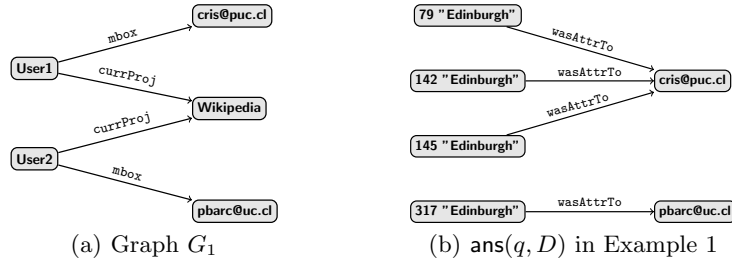


Fig. 2. Graphs used for Example 1. The prefixes foaf: and prov: are omitted.

features that are of importance to us: the GRAPH operator and the CONSTRUCT result form. First we fix some conventional notation.

We assume all variables come from an infinite set $\mathbf{V} = \{?x, ?y, \dots\}$ of *variables*. The official syntax for SPARQL 1.1 queries considers several operators such as OPTIONAL, UNION, FILTER, GRAPH and concatenation via the point symbol (.) to construct what is known as *graph patterns*. Users then use a result form such as SELECT or CONSTRUCT to form either result sets or RDF graphs from the matchings of a graph pattern. We assume that readers are familiar with graph patterns, we just note the syntax of the GRAPH operator: if P is a graph pattern and $g \in \mathbf{I} \cup \mathbf{V}$ then $(\text{GRAPH } g \ P)$ is a graph pattern, called a *GRAPH-pattern*. The expression $(\text{GRAPH } g \ P)$ allows us to determine which graph from the dataset we will be matching the pattern P to. For instance if we use an IRI in place of g the pattern will be matched against the named graph with the corresponding name (if such a graph exists in the dataset), and in the case that g is a variable, P will be matched against all the graphs in the dataset.

The fragment of SPARQL graph patterns, as well as its generalisation to SELECT queries, has drawn most of the attention in the Semantic Web community. However, as the results of such queries need not be RDF graphs, we shall use the CONSTRUCT operator in order to obtain a base for recursion. A SPARQL CONSTRUCT query, or *c-query* for short, is an expression

CONSTRUCT $H \ DS$ WHERE P ,

where H is a set of triples from $(\mathbf{I} \cup \mathbf{V}) \times (\mathbf{I} \cup \mathbf{V}) \times (\mathbf{I} \cup \mathbf{V})$, called a *template*³; DS is a set of expressions of the form FROM NAMED $u_1, \dots, \text{FROM NAMED } u_n$, with each $u_i \in \mathbf{I}$ and $i \geq 0$, called a *dataset clause*⁴; and P is a graph pattern.

The idea behind the CONSTRUCT operator is that the mappings matched to the pattern P are used to construct an RDF graph according to the template H . Since all the patterns in the template are triples we will end up with an RDF graph as desired. Before giving the formal semantics of the operators we illustrate how they work by means of an example.

Example 1. Let G and G_1 be the graphs depicted in Figure 1 and Figure 2(a), respectively. Suppose we want to query both graphs to obtain a new graph where each article is linked to the email of a user who modified it. Assuming

³ We leave the study of blanks in templates as future work.

⁴ For readability we assume the default graph as given.

we have a dataset with default graph G and that the IRI identifying G_1 is <http://db.ing.puc.cl/mail>, this would be achieved by the following SPARQL CONSTRUCT query q :

```
PREFIX prov: <http://www.w3.org/ns/prov#>
PREFIX foaf: <http://xmlns.com/foaf/0.1>
CONSTRUCT {?article prov:wasAttributedTo ?mail}
FROM NAMED <http://db.ing.puc.cl/mail>
WHERE {
  ?article prov:wasGeneratedBy ?comment .
  ?comment prov:wasAssociatedWith ?usr .
  GRAPH <http://db.ing.puc.cl/mail> {?usr foaf:mbox ?mail}}
```

The result $\text{ans}(\mathbf{q}, D)$ of evaluating q over D is depicted in Figure 2(b). The construct FROM NAMED is used to specify that the dataset needs to include the graph G_1 associated with the IRI <http://db.ing.puc.cl/mail>.

SPARQL Semantics. The semantics of graph patterns is defined in terms of *mappings* [12]; that is, partial functions from variables \mathbf{V} to IRIs \mathbf{I} . Given a dataset D , and a graph G amongst the graphs of D , we denote the *evaluation* of a graph pattern P over D with respect to G as $\llbracket P \rrbracket_G^D$. The evaluation $\llbracket P \rrbracket^D$ of a pattern P over a dataset D with default graph G_0 is $\llbracket P \rrbracket_{G_0}^D$. The full definition of $\llbracket P \rrbracket^D$ and $\llbracket P \rrbracket_G^D$ can be found in the SPARQL standard, here we just note the semantics of GRAPH-patterns, for which we need some notation.

The *domain* $\text{dom}(\mu)$ of a mapping μ is the set of variables on which μ is defined. Two mappings μ_1 and μ_2 are *compatible* (written as $\mu_1 \sim \mu_2$) if $\mu_1(?x) = \mu_2(?x)$ for all variables $?x$ in $\text{dom}(\mu_1) \cap \text{dom}(\mu_2)$. If $\mu_1 \sim \mu_2$, then we write $\mu_1 \cup \mu_2$ for the mapping obtained by extending μ_1 according to μ_2 on all the variables in $\text{dom}(\mu_2) \setminus \text{dom}(\mu_1)$. Given two sets of mappings M_1 and M_2 , the *join* and *union* between M_1 and M_2 are defined respectively as follows:

$$\begin{aligned} M_1 \bowtie M_2 &= \{\mu_1 \cup \mu_2 \mid \mu_1 \in M_1, \mu_2 \in M_2 \text{ and } \mu_1 \sim \mu_2\}, \\ M_1 \cup M_2 &= \{\mu \mid \mu \in M_1 \text{ or } \mu \in M_2\} \end{aligned}$$

Let us now define the semantics of GRAPH-patterns. Consider a GRAPH-pattern $P = (\text{GRAPH } g \ P')$. Then

$$\llbracket P \rrbracket_G^D = \begin{cases} \llbracket P' \rrbracket_{\text{gr}_D(g)}^D & \text{if } g \in \mathbf{I} \\ \bigcup_{u \in \mathbf{I}} \left(\llbracket P' \rrbracket_{\text{gr}_D(u)}^D \bowtie \{\mu_{g \rightarrow u}\} \right) & \text{if } g \in \mathbf{V} \end{cases}$$

where $\mu_{g \rightarrow u}$ is the mapping with domain $\{g\}$ and where $\mu_{g \rightarrow u}(g) = u$.

Next we recall the semantics of SPARQL queries. Let \mathbf{q} be a SPARQL query and D a dataset. The answer $\text{ans}(\mathbf{q}, D)$ of q over D depends on the form of q :

- If q is a SELECT query, then $\text{ans}(\mathbf{q}, D)$ is the answer to q as defined in the SPARQL standard [12].
- If q is a c-query $\mathbf{q} = \text{CONSTRUCT } H \ DS \ \text{WHERE } P$, then let u_1, \dots, u_n be the IRIs in DS and G_1, \dots, G_n the graphs associated to these IRIs; and consider the dataset $D' = D \cup \{\langle u_1, G_1 \rangle, \dots, \langle u_n, G_n \rangle\}$. We define

$$\text{ans}(\mathbf{q}, D) = \{\mu(t) \mid \mu \in \llbracket P \rrbracket^{D'}, t \text{ is a triple in } H \text{ and } \mu \text{ is defined on variables in } t\}.$$

3 Adding Recursion to SPARQL

The most basic example of a recursive query in RDF is reachability: given a resource x , compute all the resources that are reachable from x via a path of arbitrary length. These queries, amongst others, motivated the inclusion of property paths into the recent SPARQL 1.1 standard [12]. However, as several authors subsequently pointed out, property paths fall short when trying to express queries that involve more complex ways of navigating RDF documents (cf. [3, 7, 8, 23]) and as a result several extensions have been brought forward to combat this problem [2, 6, 15, 18, 23]. Almost all of these extensions are also based on the idea of computing paths between nodes in a recursive way, and thus share a number of practical problems with property paths. Most importantly, these queries need to be implemented using algorithms that are not standard in SPARQL databases, as they are based on automata-theoretic techniques, or clever ways of doing Breadth-first search on the graph structure of RDF documents.

3.1 A Fixed Point Based Recursive Operator

We have decided to implement a different approach: a much more widespread recursive operator that allows us compute the fixed point of a wide range of SPARQL queries. Before proceeding with the formal definition we illustrate the idea behind such queries by means of an example.

Example 2. Recall graph G from Figure 1. In the Introduction we made a case for the need of a query that could compute all pairs of articles (A, A') such that A is linked to A' by a path of `wasRevisionOf` links and where all of the revisions along the way were made by the same user. We can compute this with the following recursive query.

```
PREFIX prov: <http://www.w3.org/ns/prov#>
WITH RECURSIVE http://db.ing.puc.cl/temp AS {
  CONSTRUCT {?x ?u ?y}
  FROM NAMED <http://db.ing.puc.cl/temp>
  WHERE{{
    ?x prov:wasRevisionOf ?y .
    ?x prov:wasGeneratedBy ?w .
    ?w prov:used ?y .
    ?w prov:wasAssociatedWith ?u}
  UNION{
    GRAPH <http://db.ing.puc.cl/temp> {?x ?u ?z} .
    GRAPH <http://db.ing.puc.cl/temp> {?z ?u ?y}}}
}
SELECT ?x ?y
FROM <http://db.ing.puc.cl/temp>
WHERE {?x ?u ?y}
```

Let us explain how this query works. The second line specifies that a temporal graph named `http://db.ing.puc.cl/temp` is to be constructed according to the query below which consists of a `UNION` of two subpatterns. The first pattern does not use the temporal graph and it simply extracts all triples (A, U, B) such that A was a revision of B and U is the user generating this revision. All these triples should be added to the temporal graph.

Then comes the recursive part: if (A, U, B) and (B, U, C) are triples in the temporal graph, then we also add (A, U, C) to the temporal graph. We continue

iterating until a fixed point is reached, and finally we obtain a graph that contains all the triples (A, U, A') such that A is linked to A' via a path of revisions of arbitrary length but always generated by the same user U . Finally, the **SELECT** query extracts all such pairs of articles from the constructed graph.

As hinted in the example, the following is the syntax for recursive queries. It is based on the recursive operator that is part of SQL.

Definition 1 (Syntax of recursive queries). A recursive SPARQL query, or just recursive query, is either a SPARQL query or an expression of the form

$$\text{WITH RECURSIVE } t \text{ AS } \{q_1\} q_2, \quad (1)$$

where t is an IRI from \mathbf{I} , q_1 is a c-query, and q_2 is a recursive query. The set of all recursive queries is denoted *rec-SPARQL*.

Note that in this definition q_1 is allowed to use the temporal graph t , which leads to recursive iterations. Furthermore, the query q_2 could be recursive itself, which allows us to compose recursive definitions. As usual with this type of queries, semantics is given via a fixed point iteration.

Definition 2 (Semantics of recursive queries). Let q be a recursive query of the form (1) and D an RDF dataset. If q is a non recursive query then $\text{ans}(q, D)$ is defined as usual. Otherwise the answer $\text{ans}(q, D)$ is equal to $\text{ans}(q_2, D_{LFP})$, where D_{LFP} is the least fixed point of the sequence D_0, D_1, \dots with $D_0 = D$ and

$$D_{i+1} = D \cup \{t, \text{ans}(q_1, D_i)\}, \text{ for } i \geq 0.$$

In other words, D_1 is the union of D with a temporal graph t that corresponds to the evaluation of q_1 over D , D_2 is the union of D with a temporal graph t that corresponds to the evaluation of q_1 over D_1 , and so on until $D_{i+1} = D_i$. Note that the temporal graph is completely rewritten after each iteration. This definition suggests the following pseudocode for computing the answers of a recursive query q of the form (1) over a dataset D ⁵:

1. Initialize a temporal RDF graph named after the IRI t as $G_{\text{Temp}} = \emptyset$.
2. While $\text{ans}(q_1, D \cup \{t, G_{\text{Temp}}\}) \neq G_{\text{Temp}}$ do:
 - Set $G_{\text{Temp}} = \text{ans}(q_1, D \cup \{t, G_{\text{Temp}}\})$
3. Output $\text{ans}(q_2, D \cup \{t, G_{\text{Temp}}\})$

Obviously this definition only makes sense as long as such fixed point exists. From the Knaster-Tarski theorem [17] it easily follows that the fixed point exists as long as queries used to define recursion are monotone. For the sake of presentation, here we ensure this condition by disallowing explicit negation (such as **NOT EXISTS** or **MINUS**) and optional matching from our c-queries (note that under construct queries, this fragment is known to be equivalent to queries defined by union of well designed graph patterns [16]). It was also shown in [16] that the existence of a fixed point can be guaranteed even when q_1 belongs to a rather technical fragment that does allow a limited form of negation and optional matching that extends beyond the use of unions of well designed patterns.

⁵ For readability we assume that t is not a named graph in D . If this is not the case then the pseudocode needs to be modified to meet the definition above

3.2 Complexity Analysis

Recursive queries can use either the SELECT or the CONSTRUCT result form, so there are two decision problems we need to analyze. For SELECT queries we define the problem SELQUERYANS, that receives as an input a recursive query Q using the SELECT result form, a tuple \bar{a} of IRIs from \mathbf{I} and a dataset D , and asks whether \bar{a} is in $\text{ans}(Q, D)$. For CONSTRUCT queries the problem CONQUERYANS receives a recursive query Q with a CONSTRUCT result form, a triple (s, p, o) over $\mathbf{I} \times \mathbf{I} \times \mathbf{I}$ and a dataset D , and asks whether this triple belongs to $\text{ans}(Q, D)$.

Proposition 1. *SELQUERYANS is PSPACE-complete and CONQUERYANS is NP-complete. The complexity of SELQUERYANS drops to Π_2^p if one only considers SELECT queries given by unions of well-designed graph patterns.*

Thus, at least from the point of view of computational complexity, our class of recursive queries are not more complex than standard select queries [22] or construct queries [16]. We also note that the complexity of similar recursive queries in most data models is typically complete for exponential time; what lowers our complexity is the fact that our temporal graphs are RDF graphs themselves, instead of arbitrary sets of mappings or relations.

For databases it is also common to study the data complexity of the query answering problem, that is, the same decision problems as above but considering the input query to be fixed. We denote these problems as SELQUERYANS(Q) and CONQUERYANS(Q), for select and construct queries, respectively. As we see, the problem remains in polynomial time for data complexity, albeit in a higher class than for non recursive queries (see again [22] or [16]).

Proposition 2. *Both the problem SELQUERYANS(Q) and the problem CONQUERYANS(Q) are PTIME-complete. They remain PTIME-hard even for queries without negation or optional matching.*

However, even if theoretically the problems have the same combined complexity as queries without recursion and are polynomial in data complexity, any implementation of the above algorithm is likely to run excessively slow due to a high demand on computational resources (computing the temporal graph over and over again) and would thus not be useful in practice. For this reason, instead of implementing full-fledged recursion, we decided to support a fragment of recursive queries based on what is commonly known as *linear recursive queries* [1, 10]. This restriction is common when implementing recursive operators in other database languages, most notably in SQL [24], but also in graph databases [8], as it offers a wider option of evaluation algorithms while maintaining the ability of expressing almost any recursive query that one could come up with in practice. For instance, as demonstrated in the following section, linear recursion captures all the examples we have considered thus far and it can also define any query that uses property paths. Furthermore, it can be implemented in an efficient way on top of any existing SPARQL engine using a simple and easy to understand algorithm. Next we formally define this fragment.

4 Realistic Recursion in SPARQL

The concept of *linear recursion* has become popular in the industry as a restriction for fixed point operators in relational query languages, because it presents a good tradeoff between the expressive power of recursive operators and their practical applicability. Let Q be the query WITH RECURSIVE t AS $\{q_1\} q_2$, where t is an IRI from \mathbf{I} , q_1 is a c-query, and q_2 is a recursive query. We say that Q is *linear* if for every dataset D , the answer $\text{ans}(Q, D)$ of the query corresponds to the least fixed point of the sequence given by

$$\begin{aligned} D_0 &= D, & D_{-1} &= \emptyset, \\ D_{i+1} &= D_i \cup \{(t, \text{ans}(q_1, (D \cup D_i \setminus D_{i-1})))\}. \end{aligned}$$

In other words, a recursive query is linear if, in order to compute the i -th iteration, we only need the original dataset plus the tuples that were added to the temporal graph t in the previous iteration. Considering that the final size of t might be comparable to the original dataset, linear queries save us from evaluating the query several times over an ever increasing dataset.

Most of the recursive extensions proposed for SPARQL are linear: from property paths [12] to nSPARQL [23], SPARQLeR [15] or Trial [18], and even our example. Unfortunately it is undecidable to check if a recursive query is linear (under usual complexity-theoretic assumptions) [9], so one needs to impose syntactic restrictions to enforce this condition. This is what we do next.

4.1 Linear recursive queries

Our queries are made from the union of a graph pattern that does not use the temporal IRI, denoted as p_{base} and a graph pattern p_{rec} that does mention the temporal IRI. Formally, a *linear recursive query* is an expression of the form

$$\begin{aligned} &\text{WITH RECURSIVE } t \text{ AS } \{ \\ &\quad \text{CONSTRUCT } H \text{ } DS \text{ WHERE } p_{\text{base}} \text{ UNION } p_{\text{rec}} \} q_{\text{out}} \end{aligned} \quad (2)$$

with H and DS a construct template and dataset clause as usual, with p_{base} and p_{rec} graph patterns such that only p_{rec} is allowed to mention the IRI t and with q_{out} a linear recursive query. We further require that the recursive part p_{rec} mentions the temporal IRI only once. In order to describe our algorithm, we shall abuse the notation and speak of q_{base} to denote the query CONSTRUCT H DS WHERE p_{base} and q_{rec} to denote the query CONSTRUCT H DS WHERE p_{rec} , respectively.

This simple yet powerful syntax resembles the design choices taken in most SQL commercial systems supporting recursion [24] and even graph databases [8].

For example, the query in example 2 is not linear, because the temporal iri is used twice in the pattern. Nevertheless, it can be restated as the following query that uses one level of nesting:

```

PREFIX prov: <http://www.w3.org/ns/prov#>
WITH RECURSIVE http://db.ing.puc.cl/temp1 AS {
  CONSTRUCT {?x ?u ?y}
  FROM NAMED <http://db.ing.puc.cl/temp1>
  WHERE{
    {?x prov:wasRevisionOf ?z .
     ?x prov:wasGeneratedBy ?w .
     ?w prov:used ?z .
     ?w prov:wasAssociatedWith ?u}
  UNION
  {}}
}
WITH RECURSIVE http://db.ing.puc.cl/temp2 AS {
  CONSTRUCT {?x ?u ?y}
  FROM NAMED <http://db.ing.puc.cl/temp1>
  FROM NAMED <http://db.ing.puc.cl/temp2>
  WHERE{
    GRAPH <http://db.ing.puc.cl/temp1> {?x ?u ?y}
  UNION{
    GRAPH <http://db.ing.puc.cl/temp1> {?x ?u ?z}.
    GRAPH <http://db.ing.puc.cl/temp2> {?z ?u ?y}}
}
SELECT ?x ?y
FROM <http://db.ing.puc.cl/temp>
WHERE {?x ?u ?y}

```

We wrote the union in the first query for clarity, but in general either \mathbf{p}_{base} or \mathbf{p}_{rec} can be empty. The idea of this query is to first dump all meaningful triples from the graph into a new graph `http://db.ing.puc.cl/temp1`, and then use this graph as a basis for computing the required reachability condition, that will be dumped into a second temporal graph `http://db.ing.puc.cl/temp2`⁶.

Note that these queries are indeed linear, and thus we can perform the incremental evaluation that we have described above. The separation between base and recursive query also allows us to keep track of changes made in the temporal graph without the need of computing the difference of two graphs. We have decided to implement what is known as *seminaive evaluation*, although several other alternatives have been proposed for the evaluation of these types of queries (see [10] for a good survey). Our algorithm is presented in Algorithm 1.

So what have we gained? By looking at Algorithm 1 one realizes that in each iteration we only evaluate the query over the union of the dataset and the intermediate graph G_{temp} , instead of the previous algorithm where one needed the whole graph being constructed (in this case G_{ans}). Furthermore, \mathbf{q}_{base} is evaluated only once, using \mathbf{q}_{rec} in the rest of the iterations. Considering that the temporal graph may be large, and that no indexing scheme could be available, this often results in a considerable speedup for query computation. As we see next, the computational complexity is also reduced.

Complexity Analysis. We can find some explanation of why linear recursive queries behave better in practice when revisiting the computational complexity of the query answering problem. And although the combined complexity does not change, we can show the following about data complexity:

Theorem 1. *If Q is a linear recursive query, $\text{SELQUERYANS}(Q)$ and $\text{CONQUERYANS}(Q)$ are NLOGSPACE-complete.*

⁶ One can show that in this case the nesting in this query can be avoided, and indeed an equivalent non-nested recursive query is given in the online appendix.

Algorithm 1 Computing the answer for linear recursive queries of the form (2)

Input: Query Q of the form (2), dataset D

Output: Evaluation $\text{ans}(Q, D)$ of Q over D

```
1: Set  $G_{\text{temp}} = \text{ans}(q_{\text{base}}, D)$  and  $G_{\text{ans}} = G_{\text{temp}}$ 
2: Set  $\text{size} = |G_{\text{ans}}|$ 
3: loop
4:   Set  $G_{\text{temp}} = \text{ans}(q_{\text{rec}}, D \cup \{(t, G_{\text{temp}})\})$ 
5:   Set  $G_{\text{ans}} = G_{\text{ans}} \cup G_{\text{temp}}$ 
6:   if  $\text{size} = |G_{\text{ans}}|$  then
7:     break
8:   else
9:      $\text{size} = |G_{\text{ans}}|$ 
10:  end if
11: end loop
12: return  $\text{ans}(q_{\text{out}}, D \cup \{(t, G_{\text{ans}})\})$ 
```

5 Experimental Evaluation

Our implementation of linear recursive queries was carried out using the Apache Jena framework [14] as an add-on to the ARQ SPARQL query engine. The version used was Jena TDB 2.12.1 as it allows the user to run queries either in main memory, or using disk storage when needed. As previously mentioned, since the query evaluation algorithms we develop make use of the same operations that already exist in current SPARQL engines, we can use those as a basis for the recursive extension to SPARQL we propose. In fact, as we show by implementing recursion on top of Jena, this capability can be added to an existing engine in an elegant and non-intrusive way⁷.

We test our implementation using three different datasets. The first one is Linked Movie Database (LMDb) [19], an RDF dataset containing information about movies and actors⁸. The second dataset we use is a part of the YAGO ontology [28] and consists of all the facts that hold between instances. For the experiments the version from March 2015 was used. The last dataset is based on Missier and Chen’s database of Wikipedia traces [20] we described previously. We chose 3 of their datasets, but since they are very small we enlarge them by taking disjoint copies of the same data until it reached the desired size. Since these datasets contain only the traces and nothing else we also added 30% of random unrelated triples to simulate the database containing other pieces of information. We grew 4 different datasets out of the provenance traces, of 50, 100, 150 and 200 Mb of size approximately. We refer to these datasets as PROV1, PROV2, PROV3 and PROV4, respectively⁹. All the experiments were run on a MacBook Air with an Intel Core i5 1.3 GHz processor and 4GB of main memory.

⁷ The implementation we use is available at <http://web.ing.puc.cl/~jreutter/Recsparql.html>.

⁸ We use the data dump available at <http://queens.db.toronto.edu/~oktie/linkedmdb/>.

⁹ The datasets are available at <http://web.ing.puc.cl/~jreutter/Recsparql.html>.

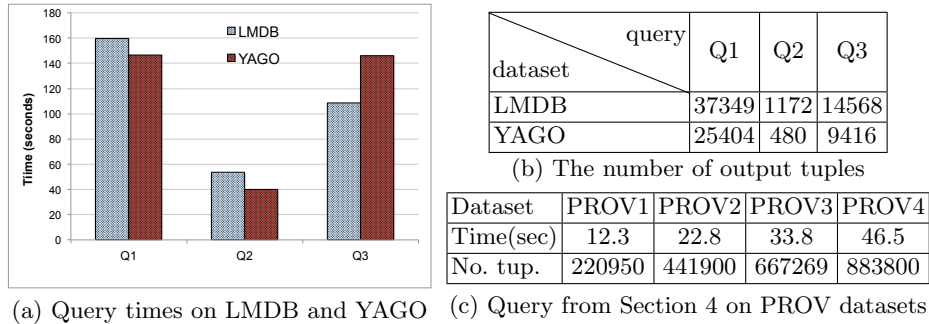


Fig. 3. Running times and the number of output tuples for the three datasets.

5.1 Query evaluation

Because of the novelty of our approach it was impossible to compare our times against other implementations, or run standard benchmarks to test the performance of our queries. Furthermore, while our formalism is similar to that of recursive SQL, all of the RDF systems that we checked were either running RDF natively, or running on top of a relational DBMS that did not support recursion as mandated by the SQL standard. OpenLink Virtuoso does have a *transitive closure* operator, but this operator can only compute transitivity when starting in a given IRI. Our queries were more general than this, and thus we could not compare them. For this reason we invented several queries that are very natural over the considered datasets and tested their performance. As all property paths can be expressed by linear recursive queries we will also test our implementation against current SPARQL systems in the following subsection.

We start our round of experiments with movie-related queries over both LMDB and YAGO. Since YAGO also contains information about movies, we have the advantage of being able to test the same queries over different real datasets (only the ontology differs). We use three different queries, all of them similar to that of Example 2 (precise queries are given in the online appendix). The first query Q1 returns all the actors in the database that have a finite Bacon number¹⁰, meaning that they co-starred in the same movie with Kevin Bacon, or another actor with a finite Bacon number. A similar notion, well known in mathematics, is the Erdős number. Note that Q1 is a property path query. To test recursive capabilities of our implementation we use another two queries, Q2 and Q3, that apply various tests along the paths computing the Bacon number. The query Q2 returns all actors with a finite Bacon number such that all the collaborations were done in movies with the same director. Finally the query Q3 tests if an actor is connected to Kevin Bacon through movies where the director is also an actor (not necessarily in the same movie). The structure of queries Q2 and Q3 is similar to the query from Example 2 and cannot be expressed using property paths either. The results of the evaluation can be found in Figure 3(a). As we can see the running times, although high, are reasonable considering the size of the datasets and the number of output tuples (Figure 3(b)).

¹⁰ See http://en.wikipedia.org/wiki/Six_Degrees_of_Kevin_Bacon.

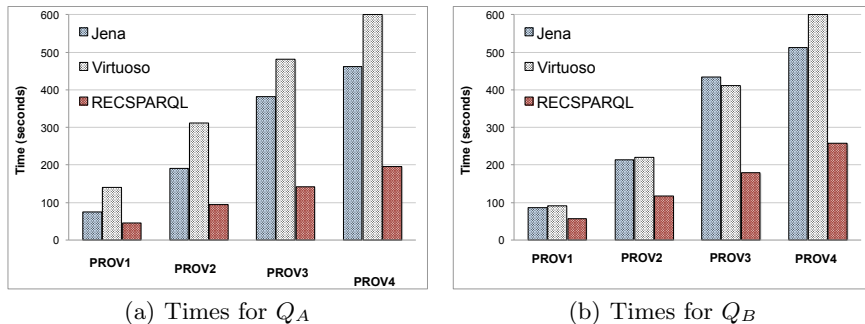


Fig. 4. Evaluation time of Q_A and Q_B in our implementation is comparable to that of Jena or Virtuoso. For PROV4 both queries reported more than 600 seconds in Virtuoso.

The next round of experiments pushes our implementation to compute inherently recursive queries. For this we use the query from Example 2 that finds all pairs of Wikipedia articles whose revision history can be attributed to the same user. As we implement linear recursion, the version of the query presented in Section 4 is used. Figure 3(c) shows the running time of this query on the datasets derived from Wikipedia traces described before; it illustrates that running times are quite low when we take the number of computed tuples into consideration.

5.2 Comparison with Property Paths

Since to the best of our knowledge no SPARQL engine implements general recursive queries, we cannot really compare the performance of our implementation with the existing systems. The only form of recursion mandated by the latest language standard are property paths, so in this section we test how our implementation stacks against popular systems when executing property paths.

Every property path query is easily expressible using linear recursion. However, it is not fair to compare our recursive implementation of property paths to the one in current systems, as they specialize in executing this type of recursive queries, while the recursive operator we introduced is aimed at expressing a wide variety of queries that lie beyond the scope of property paths. For this reason highly efficient systems like Virtuoso will run queries they are optimized for much faster. For instance to run the query Q1 from Subsection 5.1 that computes all actors with a finite Bacon number in LMDB or YAGO Virtuoso takes less than 10 seconds, while our implementation takes much longer. Part of the difference in running times could be attributed to the fact that in this particular case our implementation runs queries on disc, while Virtuoso can perform them in main memory, but the main detractor is the fact that Virtuoso is designed to be efficient at property paths that are given a starting point, while recursive queries are not since they can express more general queries.

To have a somewhat fair comparison we will use property path queries that compute all pairs of IRIs connected by a specified property path. We use the PROV datasets introduced above and in Figure 1 and test for the existence of property paths `wasRevisionOf*` and `(wasGeneratedBy/used)*`. We refer to these

queries as Q_A and Q_B .¹¹ Figure 4 presents the time each of the queries takes on the four PROV datasets of increasing size. We test the recursive implementation of property paths against the one in Jena and Virtuoso. As we can see our implementation is quite competitive with systems that specialize in property paths when we need to compute the entire relation. We can also see that Jena runs faster than Virtuoso in this case and we believe that this is due to the fact that Jena implements property paths in a way that returns all pairs of nodes that are connected by the specified query, while for Virtuoso we need to run the query from every possible starting point.

5.3 Limiting the number of iterations

In practical scenarios users are often interested in running recursive queries only for a predefined number of iterations. For instance, very long paths between nodes are seldom of interest and in a many use cases we will be interested in using property paths only up to depth four or five. For this reason we propose the following syntax to restrict the depth of recursion to a user specified number:

$$\begin{aligned}
 & \text{WITH RECURSIVE } t \text{ AS } \{ \\
 & \quad \text{CONSTRUCT } H \text{ DS WHERE } \mathbf{p}_{\text{base}} \text{ UNION } \mathbf{p}_{\text{rec}} \\
 & \quad \} \text{ MAXRECURSION } k \mathbf{q}_{\text{out}}
 \end{aligned} \tag{3}$$

Here all the keywords are the same as when defining linear recursion, and $k \geq 1$ is a natural number. The semantics of such queries is defined using Algorithm 1, where the loop between steps 4 and 12 is executed precisely $k - 1$ times.

It is straightforward to see that every query defined using recursion with predefined number of iterations can be rewritten in SPARQL by explicitly specifying each step of the recursion and joining them using the union operator. The question then is, why is specifying the recursion depth beneficial?

One apparent reason is that it makes queries much easier to write and understand (as a reference we include the rewritings of the query Q1, Q2 and Q3 from Subsection 5.1 using only SPARQL operators in the online appendix). The second reason we would like to argue for is that, when implemented using Algorithm 1, recursive queries with a predetermined number of steps result in faster query evaluation times than evaluating an equivalent query with lots of joins. The intuitive reason behind this is that computing \mathbf{q}_{base} , although expensive initially, acts as a sort of index to iterate upon, resulting in fast evaluation times as the number of iterations increases. On the other hand, for even a moderately complex query using lots of joins, the execution plan will seldom be optimal and will often resort to simply trying all the possible matchings to the variables, thus recomputing the same information several times.

We substantiate this claim by running two rounds of experiments on LMDB and YAGO using queries Q1, Q2 and Q3 from Subsection 5.1 and running them

¹¹ Complete recursive formulation of the queries can be found in the online appendix. Note that in Virtuoso we need to specify the starting point of a property path. This is done by extracting each node from a unique triple containing it.

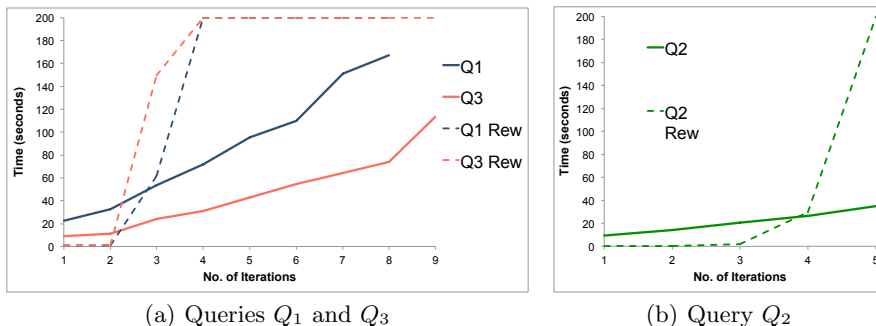


Fig. 5. Limiting the number of iterations for Q_1 , Q_2 and Q_3 over LMDB. Recursion dominates manually written SPARQL joins when several iterations are required.

for an increasing number of steps. In the first round we evaluate each of the queries using Algorithm 1 and run it for a fixed number of steps until the algorithm saturates. In the second round we use a SPARQL rewriting of a recursive query where the depth of recursion is fixed and evaluate it in Jena.

Figure 5 shows the results over LMDB. The results for YAGO show the same trend, so we do not include them. As we can see, the initial cost is much higher if we are using recursive queries, however as the number of steps increases we can see that they show much better performance and in fact, the queries that use only SPARQL operators time out after a small number of iterations.

6 Conclusion

As illustrated by several use cases, there is a need for recursive functionalities in SPARQL that go beyond the scope of property paths. To tackle this issue we propose a recursive operator to be added to the language and show how it can be implemented efficiently on top of existing SPARQL systems. We concentrated on linear recursive queries which have been well established in SQL practice and cover almost all interesting use cases and show how to implement them as an extension to Jena framework. Our tests show that, although very expressive, these queries run in reasonable time even on a machine with limited computational resources. We also include a command that allows to run recursive queries for a limited number of steps and show that the proposed implementation outperforms equivalent queries specified using only SPARQL operators. We believe all of this to be a good indicator of the usefulness of the recursion operator and why it should be a potential candidate for inclusion in the next SPARQL standard.

As for future work, we are interested in studying the capabilities of recursion using blank nodes in the templates of c-queries, and how this influences the capability of recursive SPARQL queries in terms of computing ontological knowledge. From the results of [23] we already know that recursive queries can compute the closure of a graph under RDFS entailment, however a detailed study on performance and expressivity is needed to fully understand the underlying issues. We are also interested in seeing if recursive algorithms can be parallelized. A good starting point for this would be try and use the algorithms developed in [21] which deal with similar problems.

References

1. S. Abiteboul, R. Hull, and V. Vianu. *Foundations of Databases*. A-W, 1995.
2. F. Alkhateeb, J.-F. Baget, and J. Euzenat. Extending SPARQL with regular expression patterns (for querying RDF). *J. Web Sem.*, 7(2):57–73, 2009.
3. K. Anyanwu and A. Sheth. ρ -Queries: enabling querying for semantic associations on the semantic web. In *WWW*, 2003.
4. M. Atzori. Computing recursive SPARQL queries. In *ICSC*, pages 258–259, 2014.
5. P. Barceló, J. Pérez, and J. L. Reutter. Relative Expressiveness of Nested Regular Expressions. In *AMW*, pages 180–195, 2012.
6. M. Bienvenu, D. Calvanese, M. Ortiz, and M. Simkus. Nested regular path queries in description logics. In *KR 2014, Vienna, Austria, July 20-24, 2014*, 2014.
7. P. Bourhis, M. Krötzsch, and S. Rudolph. How to best nest regular path queries. In *Proceedings of the 27th International Workshop on Description Logics*, 2014.
8. M. Consens and A. Mendelzon. Graphlog: A visual formalism for real life recursion. In *PODS*, pages 404–416, 1990.
9. H. Gaifman, H. Mairson, Y. Sagiv, and M. Y. Vardi. Undecidable optimization problems for database logic programs. *JACM*, 40(3):683–713, 1993.
10. T. J. Green, S. S. Huang, B. T. Loo, and W. Zhou. Datalog and recursive query processing. *Foundations and Trends in Databases*, 5(2):105–195, 2013.
11. A. Gubichev, S. J. Bedathur, and S. Seufert. Sparqling kleene: fast property paths in RDF-3X. In *GRADES*, 2013.
12. S. Harris and A. Seaborne. SPARQL 1.1 query language. *W3C Recommendation*, 21, 2013.
13. P. Hitzler, M. Krotzsch, and S. Rudolph. *Foundations of semantic web technologies*. CRC Press, 2011.
14. The Apache Jena Manual. <http://jena.apache.org>, 2015.
15. K. J. Kochut and M. Janik. Sparqler: Extended sparql for semantic association discovery. In *The SW: Research and Applications*, pages 145–159. Springer, 2007.
16. E. V. Kostylev, J. L. Reutter, and M. Ugarte. CONSTRUCT queries in SPARQL. In *ICDT*, pages 212–229, 2015.
17. L. Libkin. *Elements of Finite Model Theory*. Springer, 2004.
18. L. Libkin, J. Reutter, and D. Vrgoč. Trial for RDF: adapting graph query languages for RDF data. In *PODS*, pages 201–212. ACM, 2013.
19. Linked movie database. <http://linkedmdb.org/>.
20. P. Missier and Z. Chen. Extracting prov provenance traces from wikipedia history pages. In *EDBT/ICDT 2013 Workshops*, pages 327–330, 2013.
21. B. Motik, Y. Nenov, R. Piro, I. Horrocks, and D. Olteanu. Parallel materialisation of datalog programs in centralised, main-memory RDF systems. In *AAAI*, 2014.
22. J. Pérez, M. Arenas, and C. Gutierrez. Semantics and complexity of SPARQL. *ACM Transactions on Database Systems*, 34(3), 2009.
23. J. Pérez, M. Arenas, and C. Gutierrez. nSPARQL: A navigational language for RDF. *J. Web Sem.*, 8(4):255–270, 2010.
24. PostgreSQL. <http://www.postgresql.org/docs/current/interactive/queries-with.html>.
25. PROV Model Primer. <http://www.w3.org/TR/2013/NOTE-prov-primer-20130430/>, 2013.
26. Open Link Virtuoso. <http://virtuoso.openlinksw.com/>, 2015.
27. W3C. PROV-O: The PROV Ontology. <http://www.w3.org/TR/prov-o/>, 2013.
28. YAGO: A High-Quality Knowledge Base. <http://www.mpi-inf.mpg.de/departments/databases-and-information-systems/research/yago-naga/yago/>.
29. N. Yakovets, P. Godfrey, and J. Gryz. Evaluation of SPARQL property paths via recursive SQL. In *AMW*, 2013.