Fast Processing of SPARQL Queries on RDF Quadruples

V. Slavov A. Katib V. Nuchimaniyanda P. Rao S. Paturi
D. Barenkala

University of Missouri-Kansas City, Kansas City

Technical Report UMKC TR-DB-2014-01

Abstract

In this paper, we propose a new approach for fast processing of SPARQL queries on large RDF datasets containing RDF quadruples (or quads). Our approach called RIQ employs a *decrease-and-conquer* strategy: Rather than indexing the entire RDF dataset, RIQ identifies groups of similar RDF graphs and indexes each group separately. During query processing, RIQ uses a novel filtering index to first identify candidate groups that may contain matches for the query. On these candidates, it executes optimized queries using a conventional SPARQL processor to produce the final results. We evaluated RIQ using a synthetic and a real dataset, each containing about 1.4 billion quads, we show that RIQ can outperform RDF-3X and Jena TDB on a variety of SPARQL queries.

1 Introduction

The Resource Description Framework (RDF) is a standard model for representing data on the Web [5]. It enables the interchange and machine processing of data by considering its semantics. While RDF was first proposed with the vision of enabling the Semantic Web, it has now become popular in domain-specific applications and the Web. Through advanced RDF technologies, one can perform semantic reasoning over data and extract knowledge in domains such as healthcare, biopharmaceuticals, defense, and intelligence. Linked Data [12] is a popular use case of RDF on the Web; it has a large collection of different knowledge bases, which are represented in RDF (*e.g.*, DBpedia [10]).

With a growing number of new applications relying on Semantic Web technologies (*e.g.*, Pfizer [4]; Newsweek, BBC, The New York Times, and Best Buy [1]) and the availability of large RDF datasets (*e.g.*, Billion Triples Challenge (BTC) [6], Linking Open Government Data (LOGD) [2]), there is a need to advance the state-of-the-art in storing, indexing, and query processing of RDF datasets.

Today, datasets containing over a billion RDF quads are becoming popular on the Web (*e.g.*, BTC [6], LOGD [2]). Such datasets can be viewed as a collection of RDF graphs. Using SPARQL's `GRAPH` keyword [7], one can pose a query to match a specific graph pattern within any single RDF graph. While researchers in the database community have proposed scalable approaches for indexing and query processing of large RDF datasets [8, 27, 24, 9, 21, 13, 29, 30], they have designed these techniques for RDF datasets containing triples. In addition, none of them have investigated how large and complex graph patterns in SPARQL queries can be processed efficiently. Evidently, RDF-3X [24], a popular scalable approach for a local/centralized environment, yields poor performance when SPARQL queries containing large graph patterns are processed over large RDF datasets. This is because of the large number of join operations that must be performed to process a query.

We posit that, on RDF datasets containing billions of quads, any approach that first finds matches for subpatterns in a large graph pattern and then employs join operations to merge partial matches will face a similar limitation. Motivated by the aforementioned reasons, we propose a new approach called RIQ (**R**DF **I**ndexing on **Q**uads) and make the following contributions in this paper:

- We propose a new vector representation for RDF graphs and graph patterns in SPARQL queries. This representation enables us to group similar RDF graphs and index each group separately rather than

```
PREFIX movie: <http://data.linkedmdb.org/resource/movie/>
PREFIX rdfs: <http://www.w3.org/2000/01/rdf-schema#>
PREFIX foaf: <http://xmlns.com/foaf/0.1/>
SELECT ?g ?producer ?name ?label ?page ?film WHERE {
  GRAPH ?g { ?producer movie:producer_name ?name .
             ?producer rdfs:label ?label .
             OPTIONAL { ?producer foaf:page ?page . } .
             ?film movie:producer ?producer .  } }
```

Figure 1: An example of a SPARQL query

constructing an index on the entire dataset. We propose a novel filtering index, which employs a combination of Bloom Filters and Counting Bloom Filters to compactly store it.

- We propose a *decrease-and-conquer* approach to efficiently process a SPARQL query. Using the filtering index, we can methodically and quickly identify candidate groups of RDF graphs that may contain a match for the query. We can then execute optimized queries on the candidates using a conventional SPARQL processor that supports quads.
- We report the results from our performance evaluation using a synthetic and a real dataset, each containing about 1.4 billion quads. We observed that RIQ can outperform RDF-3X and Jena TDB on a variety of SPARQL queries.

2 Background

The RDF data model provides a simple way to represent any assertion as a (subject, predicate, object) triple. A collection of triples can be modeled as a directed, labeled graph. If each triple has a graph name (or context), it is called a quad. Below is an example of a quad from the BTC 2012 dataset [6] with its subject, predicate, object, and context: `<http://data.linkedmdb.org/resource/producer/10138>` `<http://data.linkedmdb.org/resource/movie/producer_name>` "`Mani Ratnam`" `<http://data.linkedmdb.org/data/producer/10138>`. Triples with the same context belong to the same RDF graph.

Using SPARQL, one can express complex graph pattern queries on RDF graphs. One of the fundamental operations in RDF query processing is *Basic Graph Pattern Matching* [7]. A Basic Graph Pattern (BGP) in a query combines a set of triple patterns. A triple pattern contains variables (prefixed by `?`) and constants. During query processing, the variables in a BGP are bound to RDF terms in the data, *i.e.*, the nodes in the same RDF graph, via subgraph matching [7]. Common variables within a BGP or across BGPs denote a join operation on the variable bindings of triple patterns. Consider the query shown in Figure 1. The bindings for the subject (variable) in `?producer movie:producer_name ?name` are joined with the bindings for the object (variable) in `?film movie:producer ?producer`. The variable `?g` will be bound to the names/contexts of those RDF graphs that contain a match for the graph pattern specified inside the `GRAPH` block. `OPTIONAL` allows certain patterns to have empty bindings; `UNION` combines bindings of multiple graph patterns; `FILTER EXISTS` / `NOT EXISTS` tests for existence or non-existence of certain graph patterns.

3 Related Work and Motivation

Several approaches have been developed for indexing and querying RDF data in a local/centralized environment. Early approaches employed an RDBMS to store and query RDF data (*e.g.*, Sesame [17], Oracle [18]). Unfortunately, the cost of self-joins on a single (triples) table became a serious bottleneck. Later, Abadi *et al.* proposed the idea of vertically partitioning the property tables [28] and used a column-oriented DBMS to achieve an order of magnitude performance improvement over previous techniques [8]. Recently, Neumann *et al.* developed RDF-3X [24] that builds exhaustive indexes on the six permutations of (s, p, o) triples. RDF-3X significantly outperformed the vertical partitioning approach. It uses a new join ordering method based on selectivity estimates and builds compressed indexes. Weiss *et al.* [27] developed Hexastore that also builds

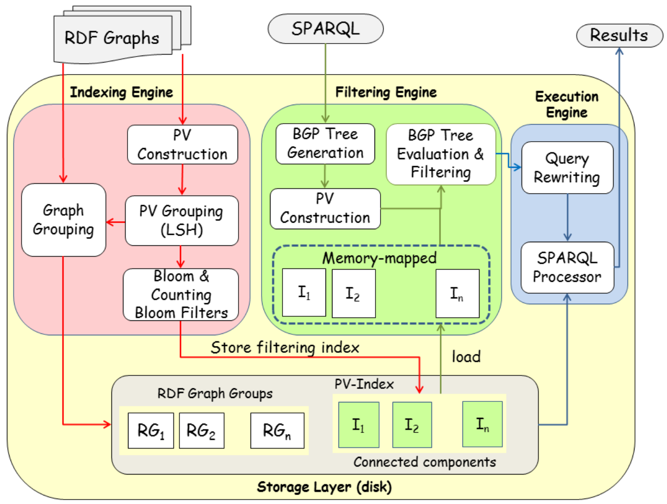

Figure 2: Overview of RIQ

exhaustive indexes. However, Hexastore suffers from large index sizes due to lack of compression. Atre *et al.* [9] developed BitMat to overcome the overhead of large intermediate join results for queries containing low selectivity triple patterns. BitMat performs in-memory processing of compressed bit matrices during query processing.

More recently, Bornea *et al.* [13] developed DB2RDF by using an RDBMS to store and query RDF data. By storing the predicate-object pairs of each subject in the same row of the relational table, they reduced the number of joins required for star-shaped BGPs. DB2RDF maintains only subject and object indexes and employs a novel SPARQL-to-SQL translation technique for generating optimized queries. Yuan *et al.* [29] developed TripleBit, which uses a compact storage scheme for RDF data by representing triples via a Triple Matrix. For each predicate, TripleBit maintains SO and OS ordered buckets. Using a collection of indexes and optimal join ordering, it reduces the size of the intermediate results during query processing.

A few approaches exploit the graph properties of RDF data for indexing and query processing [26, 14, 31]. These techniques, however, have been tested only on small RDF datasets containing less than 50 million triples. Recently, a few schemes were proposed for distributed/parallel RDF query processing [21, 30]. Our work, however, focuses on RDF query processing in a local environment.

The motivation for our work stems from two key observations: First, the above approaches were designed to process RDF datasets containing triples. Simply ignoring the context in an RDF quad and using an existing approach designed for triples may produce incorrect results due to bindings for a BGP from different graphs. For example, consider the two quads: `<a> <b> <c> <g1> .` `<a> <b> <e> <g2> .` If we use a technique designed to process triples for the query `SELECT ?x WHERE { GRAPH ?g { ?x <b> <c> . ?x <b> <e> . } }` on these quads. Then, `?x` will be bound to `<a>` as triples '`<a> <b> <c>`' and '`<a> <b> <e>`' will be treated as part of the same graph. In reality, the correct evaluation of this query should produce no results.

Second, most of the queries tested by these approaches contain BGPs with a modest number of triples patterns (at most 8). None of them have investigated how to efficiently process SPARQL queries with large and complex BGPs (*e.g.*, containing undirected cycles[1]). A few examples are shown in Appendix A.

4 The Design of RIQ

In this section, we present the design of RIQ (RDF Indexing on Quadruples) and describe its three main components, *i.e.*, the Indexing Engine, the Filtering Engine, and the Execution Engine. (See Figure 2.) Our goal is to support a subset of the SPARQL grammar [7].

[1]Here is an example: {?a p ?b . ?b q ?c . ?a r ?c .}.

Transformation f_D	Transformation f_Q
f_D(SPO, (s,p,o)) = (s,p,o)	f_Q(‘s p o’) = (SPO,(s,p,o))
f_D(SP?, (s,p,o)) = (s,p,?)	f_Q(‘s p ?v_o’) = (SP?,(s,p,?))
f_D(S?O, (s,p,o)) = (s,?,o)	f_Q(‘s ?v_p o’) = (S?O,(s,?,o))
f_D(?PO, (s,p,o)) = (?,p,o)	f_Q(‘?v_s p o’) = (?PO,(?,p,o))
f_D(S??, (s,p,o)) = (s,?,?)	f_Q(‘s ?v_p ?o’) = (S??,(s,?,?))
f_D(?P?, (s,p,o)) = (?,p,?)	f_Q(‘?v_s p ?v_o’) = (?P?,(?,p,?))
f_D(??O, (s,p,o)) = (?,?,o)	f_Q(‘?v_s ?v_p o’) = (??O,(?,?,o))

Table 1: Transformations in RIQ

4.1 Indexing RDF Data

We introduce a new vector representation for RDF graphs and BGPs, which will allows us to capture the properties of the triples and triple patterns in them. This vector representation plays a key role in the construction of an effective filtering index, where similar RDF graphs will be grouped together.

4.1.1 Essential Transformations

To begin with, we define two transformations: one for a triple in an RDF graph and the other for a triple pattern in a BGP. Let $\mathbb{P} = \{SPO, SP?, S?O, ?PO, S??, ?P?, ??O\}$ be a set of canonical patterns. We denote the transformation on a triple (s,p,o) by $f_D : \mathbb{P} \times \{(s,p,o)\} \rightarrow O_D$, where the range O_D is shown Table 1 for each canonical pattern. Note that O_D resembles triple patterns (variable names excluded) that can appear in a BGP.

Next, we denote a transformation $f_Q : T \rightarrow \mathbb{P} \times O_Q$, where T denotes the set of triple patterns that can appear in a query. The range $\mathbb{P} \times O_Q$ is shown in Table 1 and identifies the canonical pattern for a given triple pattern. Although the triple pattern ‘s p o’ has no variables, it is still a valid triple pattern in a BGP.[2]

The transformations f_D and f_Q allow us to map a triple in the data and a triple pattern in a query to a common plane of reference. This will enable us to quickly test *if a triple pattern in a BGP has a match in the data*.

4.1.2 Pattern Vectors

Given an RDF graph with context c, we map it into a vector representation called a Pattern Vector (PV) and denote it by $\overline{V_c}$. Essentially, $\overline{V_c} = (V_{c,SPO}, V_{c,SP?}, V_{c,S?O}, V_{c,?PO}, V_{c,S??}, V_{c,?P?}, V_{c,??O})$, where each $V_{c,r}$ denotes the vector constructed for $r \in \mathbb{P}$. We assume a hash function $\mathbb{H} : B \rightarrow \mathbb{Z}^*$, where B denotes a bit string and the range is the set of non-negative integers. Now, we construct $\overline{V_c}$ as follows: Initially, each $V_{c,r}$ is empty. Given a quad (s,p,o,c) in the graph, for each $r \in \mathbb{P}$, we compute $\mathbb{H}(f_D(r,(s,p,o)))$ and insert it into $V_{c,r}$. We perform this computation on every quad in the graph to generate $\overline{V_c}$. Note that $\overline{V_c}$ requires space linear in the number of quads in the graph.

Our hash function $\mathbb{H}$ is based on Rabin's fingerprinting technique [25], which is efficient to compute. If we generate 32-bit hash values, the probability of collision is extremely low. Suppose the hash values are 32 bit unsigned integers and we use an irreducible polynomial of degree 31. If each triple pattern requires at most 2048 bits, then the probability of collision is less than 2^{-20} [15]. Thus, in practice, we can view $V_{c,SPO}$ as a set, because the quads/triples in a graph are always assumed to be unique. However, the remaining vectors of $\overline{V_c}$ should be viewed as multisets, because f_D can produce the same output for different triples due to the presence of ‘?’ in the output.

Given a BGP q, we map it into a PV, denoted by $\overline{V_q}$, and compute it slightly differently: Initially, each $V_{q,r}$ is empty. For each triple pattern t in q, we compute $f_Q(t)$ to produce a pair (r,o), where r denotes the canonical pattern for t. We then insert $\mathbb{H}(o)$ into $V_{q,r}$. As before, $V_{q,SPO}$ can be viewed as a set. The rest of the vectors of $\overline{V_q}$ should be viewed as multisets, because two different triple patterns (each containing at least one variable) in a BGP may hash to the same value. For example, if a BGP contains two triple

[2] SELECT ?g WHERE { GRAPH ?g { s p o . } }.

patterns ?s_1 `movie:producer` ?o_1 and ?s_2 `movie:producer` ?o_2, then f_Q('?s_1 `movie:producer` ?o_1') = f_Q('?s_2 `movie:producer` ?o_2') and therefore, the hash values produced by $\mathbb{H}$ will be identical. Note that we will ignore triple patterns of the type `?s ?p ?o` during the PV construction. These will be, however, processed when optimized queries are executed on candidate groups during query processing.

4.1.3 Operations on Pattern Vectors

Next, we define two operations on PVs, which will be used during the construction of the filtering index. Our goal is to group similar PVs (and as a result, similar RDF graphs) together so that candidate RDF graphs are identified and processed quickly during query processing.

Definition 1 (Union) *Given two PVs, say $\overline{V_a}$ and $\overline{V_b}$, their union $\overline{V_a} \cup \overline{V_b}$ is a PV say $\overline{V_c}$, where $V_{c,r} \leftarrow V_{a,r} \cup V_{b,r}$ and $r \in \mathbb{P}$.*

Definition 2 (Similarity) *Given two PVs, say $\overline{V_a}$ and $\overline{V_b}$, their similarity is denoted by $sim(\overline{V_a}, \overline{V_b}) = \max_{r \in \mathbb{P}} sim(V_{a,r}, V_{b,r})$, where $sim(V_{a,r}, V_{b,r}) = \frac{|V_{a,r} \cap V_{b,r}|}{|V_{a,r} \cup V_{b,r}|}$.*

4.1.4 Index Construction

We begin by describing a key necessary condition, which forms the basis for indexing and query processing in RIQ. Because we map both the RDF graphs and BGPs into their PVs, we must characterize the relationship between them when processing a BGP via subgraph matching. We state the following theorem.

Theorem 1 *Suppose $\overline{V_c}$ and $\overline{V_q}$ denote the PVs of an RDF graph and a BGP, respectively. If the BGP has a subgraph match in the RDF graph, then $\bigwedge_{r \in \mathbb{P}} (V_{q,r} \subseteq V_{c,r}) = \texttt{TRUE}$.*

Proof. We assume that the BGP q denotes a connected graph. Because q has a subgraph match in the graph, every triple pattern in q has a matching triple in the graph. Consider a triple pattern t in q. Let $(r, o) \leftarrow f_Q(t)$. During the construction of V_q, we inserted $\mathbb{H}(o)$ into $V_{q,r}$. Suppose d denotes the matching triple pattern for t in the graph. During the construction of V_c, we had inserted $\mathbb{H}(f_D(r, d))$ into $V_{c,r}$. Also, $\mathbb{H}(o) = \mathbb{H}(f_D(r, d))$. Therefore, elements in $V_{q,r}$ have a one-to-one correspondence with a subset of elements in $V_{c,r}$. Hence, $V_{q,r} \subseteq V_{c,r}$. This is true for every $r \in \mathbb{P}$, and hence, $\bigwedge_{r \in \mathbb{P}} (V_{q,r} \subseteq V_{c,r}) = \texttt{TRUE}$. □

According to Theorem 1, given a BGP, if we can identify those RDF graphs in the database whose PVs satisfy the necessary condition, then we have a superset of RDF graphs that contain a subgraph match for the BGP. This also guarantees that there are no false dismissals.

Rather than testing every PV in the database – one-at-a-time – during query processing, we propose a novel filtering index called the PV-Index to effectively organize millions of PVs in the database. Using this index, we aim to quickly identify candidate RDF graphs in the early stages of query processing using Theorem 1. Our goal is to discard most of the non-matching RDF graphs without any false dismissals. As a result, the subsequent stages of query processing will process fewer candidates to obtain the final results, thereby speeding up query processing.

There are two issues that arise while designing the PV-Index: First, we want to group similar PVs together so that for a given BGP, we can quickly discard most of the non-matching RDF graphs. Second, we want to compactly store the PV-Index to minimize the cost of I/O during query processing. To address the first issue, we use the concept of locality sensitive hashing (LSH) [22]. For similarity on sets based on the Jaccard index, LSH on a set S, denoted by $\mathsf{LSH}_{k,l,m}(S)$ can be performed as follows [20]: Pick $k \times l$ random linear hash functions of the form $h(x) = (ax + b)\ mod\ u$, where u is a prime, and a and b are integers such that $0 < a < u$ and $0 \leq b < u$. Compute $g(S) = \min\{h(x)\}$ over all items in the set as the output hash value for S. Each group of l hash values is hashed (*e.g.*, using Rabin's fingerprinting) to the range $[0, m-1]$. This results in k hash values for S. It is known that given two sets S_1 and S_2 with similarity $p = \frac{|S_1 \cap S_2|}{|S_1 \cup S_2|}$, $\Pr[g(S_1) = g(S_2)] = p$. Also, the probability that $\mathsf{LSH}_{k,l,m}(S_1)$ and $\mathsf{LSH}_{k,l,m}(S_2)$ have at least one hash value identical is $1 - (1 - p^l)^k$. The above properties also hold for multisets.

To address the second issue, we employ Bloom filters (BFs) and Counting Bloom filters (CBFs) [16] to compactly represent the PV-Index. A Bloom filter is a popular data structure to compactly represent a set of items and process membership queries on it. A Counting Bloom filter maintains n-bit counters instead of single bits and can represent multisets. Both BFs and CBFs can be configured to achieve a false positive rate based on their capacities [16].

Algorithm 1 The PV-Index Construction

Input: a list of PVs; (k, l, m): LSH parameters; ϵ: false positive rate
Output: filters of all the groups of similar RDF graphs
1: Let $\mathbb{G}(\mathbb{V}, \mathbb{E})$ be initialized to an empty graph
2: **for** each PV $\overline{V}$ **do**
3:    Add a new vertex v_i to $\mathbb{V}$
4:    **for** each $r \in \mathbb{P}$ **do**
5:       $\{h_{i1}, ..., h_{ik}\} \leftarrow \mathsf{LSH}_{k,l,m}(V_r)$
6:       **for** every $v_j \in \mathbb{V}$ and $i \neq j$ **do**
7:          **if** $\exists\ o$ s.t. $1 \leq o \leq k$ and $h_{io} = h_{jo}$ **then**
8:             Add an edge (v_i, v_j) to $\mathbb{E}$ if not already present
9: Compute the connected components of $\mathbb{G}$. Let $\{C_1, ..., C_t\}$ denote these components.
10: **for** $i = 1$ to t **do**
11:    Compute the union U_i of all PVs corresponding to the vertices in C_i
12:    Construct a BF for $U_{i,SPO}$ with false positive rate ϵ given the capacity $|U_{i,SPO}|$
13:    Construct a CBF for each of the remaining vectors of U_i with false positive rate ϵ given the capacity $|U_{i,*}|$
14:    Store the ids of graphs belonging to C_i
15: **return**

In Algorithm 1, we outline the steps to construct the PV-Index. We build a graph $\mathbb{G}$, where each vertex of $\mathbb{G}$ represents a PV. For every PV, we apply LSH on each of its seven vectors. Suppose there are two PVs such that the application of LSH on their vectors for the same pattern r, produces at least one identical hash value, then we add an edge between the vertices representing these PVs (Lines 2 to 8). Essentially, a missing edge between two vertices indicates that their corresponding PVs are dissimilar with high probability. Once $\mathbb{G}$ is constructed, we compute (in linear time) the connected components in it. Each connected component represents RDF graphs whose corresponding PVs are similar with high probability. We treat these graphs as a group and compute the union of their PVs (Line 11). The union operation summarizes the PVs as well as preserves the condition stated in Theorem 1. (The individual vectors in a PV are kept sorted so that the union operation can be performed in linear time.)

To compactly represent the union computed for a connected component, we use a combination of one Bloom filter (BF) and six Counting Bloom filters (CBFs). The vector for the canonical pattern SPO is stored using a BF and the others are stored using CBFs. Each filter of a vector is configured for a false positive rate of ϵ and capacity equal to the cardinality of the vector (Lines 12 and 13). For each connected component, we also store the ids of graphs belonging to it. In summary, the BFs and CBFs for all the connected components constitute the PV-Index. Each group of graphs is separately indexed using a tool like Jena TDB.

4.2 Query Processing

Next, we propose a *decrease-and-conquer* approach for efficient SPARQL query processing in RIQ. That is, we first identify candidate groups of RDF graphs that may contain matches for a query using the PV-Index and then execute optimized SPARQL queries on these candidates.

Given a query, the first step is to parse its `GRAPH` block according to the SPARQL grammar and generate a tree-representation, which we call the BGP Tree. This tree serves as an execution plan for processing individual BGPs in the query. For example, the query in Figure 3(a) is represented by the BGP Tree in Figure 3(b). We maintain a Boolean variable $eval[n]$ for each node n in the tree to denote the status of the evaluation on a connected component of the PV-Index. With $eval[n]$ = `FALSE` for every node in the tree, we invoke Algorithm 2 on each connected component, starting from the root of the BGP Tree in depth-first order. When a child of `GroupGraphPatternSub` evaluates to `FALSE`, we skip processing the remaining children (Line 4), because the RDF graphs belonging to that connected component will not produce a match for the

```
SELECT * WHERE {
  GRAPH ?g {
    { ?city onto:areaLand ?area .
      ?city onto:areaCode ?code . }
    UNION
    { ?city onto:timeZone ?zone .
      ?city onto:abstract ?abstract . }
    ?city onto:country res:United_States .
    ?city onto:postalCode ?postal .
    FILTER EXISTS { ?city onto:utcOffset ?offset . }
    OPTIONAL { ?city onto:populationTotal ?popu . }
  }}
```

BGP_1 BGP_2 BGP_3 BGP_4 BGP_5

(a) Original query

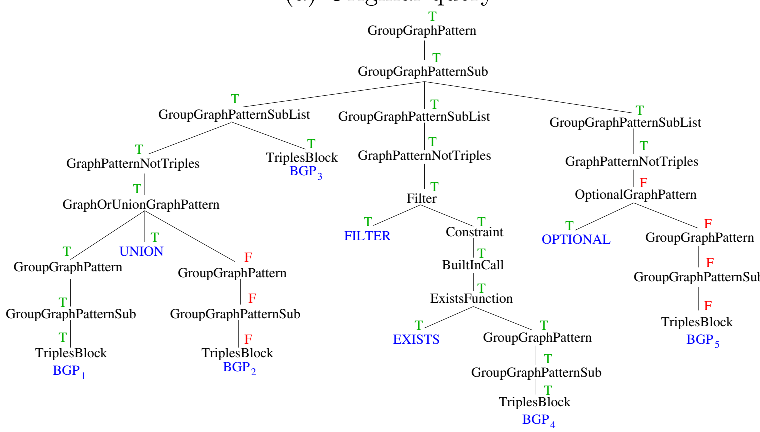

(b) A processed BGP Tree

```
SELECT ?g ?city ?area ?code ?zone ?abstract ?postal ?offset ?popu
WHERE {
  GRAPH ?g {
    { ?city onto:areaLand ?area .
      ?city onto:areaCode ?code . }
    ?city onto:country res:United_States .
    ?city onto:postalCode ?postal .
    FILTER EXISTS { ?city onto:utcOffset ?offset . }
  }}
```

(c) Optimized Query

Figure 3: Query processing

subexpression rooted at `GroupGraphPatternSub`. For `GroupOrUnionGraphPattern`, however, at least one of its children *i.e.*, `GroupGraphPattern`, should evaluate to `TRUE` to produce a match (Line 7).

When a BGP is encountered (Line 15), we test the necessary condition stated in Theorem 1 by calling Algorithm 3. This involves the processing of membership queries on the BF and CBFs constructed for that connected component. If `OptionalGraphPattern` evaluates to `FALSE`, we return `TRUE` because of the semantics of `OPTIONAL` in SPARQL. If $eval[root] =$ `TRUE`, then the group of RDF graphs belonging to that connected component is a candidate for further processing.

For the candidate, an optimized SPARQL query can be generated by traversing the BGP Tree and checking the evaluation status of each node. We prune the BGP Tree to produce a pruned BGP Tree from which the optimized query can be constructed. Algorithm 4 shows the steps involved during the pruning process starting from the root of the BGP Tree. The result modifiers and predicates within `FILTER` are included in the optimized query. All the projected variables in the original query are projected in the optimized query. Based on pruning the BGP Tree in Figure 3(b), we can generate the optimized query as shown in Figure 3(c). In this query, the `OPTIONAL` block and one block in the `UNION` are absent. The optimized query can then be executed on the candidate using a tool like Jena TDB. The results from all the candidates are combined to produce the final output.

Algorithm 2 EvalBGPTree(node n, conn. component j)

Let $c_1, ..., c_\tau$ denote the child nodes of n (left-to-right) ignoring those corresponding to braces
for $i = 1$ to τ **do**
  $eval[c_i] \leftarrow$ EvalBGPTree(c_i, j)
  if n is `GroupGraphPatternSub` & $eval[c_i]$ = `FALSE` **then**
    $eval[n] \leftarrow$ `FALSE`
    return `FALSE` {//skip rest of the nodes}
if n is `GroupOrUnionGraphPattern` **then**
  $eval[n] \leftarrow \bigvee_{i=1}^{\tau} eval[c_i]$
else if n is `ExistsFunction` **then**
  $eval[n] \leftarrow eval[c_\tau]$
else if n is `NotExistsFunction` **then**
  $eval[n] \leftarrow$ `TRUE`
else if n is `Expression` **then**
  $eval[n] \leftarrow$ `TRUE` {//skip processing predicates}
else if n is `TriplesBlock` **then**
  Let q denote the basic graph pattern
  $eval[n] \leftarrow$ IsMatch(q, j)
else if n is not a leaf **then**
  $eval[n] \leftarrow eval[c_\tau]$
else
  $eval[n] \leftarrow$ `TRUE` {//leaf nodes like `UNION`, `FILTER`}
if n is `OptionalGraphPattern` **then**
  return `TRUE`
return $eval[n]$

Algorithm 3 IsMatch(BGP q, conn. component j)

For connected component j, let $\mathbb{F}_{j,r}$ denote the BF or CBF constructed for pattern r
Construct $\mathbb{F}_{q,r}$ with the same capacity and false positive rate as $\mathbb{F}_{U_j,r}$
if (1) for each bit in $\mathbb{F}_{q,SPO}$ set to 1, the corresponding bit in $\mathbb{F}_{U_j,SPO}$ is 1, and (2) for each of the remaining patterns, given a non-zero counter in $\mathbb{F}_{q,r}$, the corresponding counter in $\mathbb{F}_{U_j,r}$ is greater than or equal to it **then**
  return `TRUE`, otherwise **return** `FALSE`

5 Performance Evaluation

In this section, we report the initial performance evaluation of RIQ and have compared it with the latest version of RDF-3X and Apache Jena 2.11.1 (TDB). RDF-3X and Jena TDB readily index datasets with more than a billion triples. Also, Jena supports RDF quads. We ran all the experiments on a 64-bit Ubuntu 12.04 machine with 4 Intel Xeon 2.4GHz cores and 16GB RAM. RIQ, a C++ codebase, uses popular open-source libraries for parsing RDF data [11] and constructing BFs and CBFs [19]. All the three approaches were single-threaded.

5.1 Datasets and Queries

We used one synthetic and one real dataset in our experiments. The synthetic dataset was generated using the Lehigh University Benchmark (LUBM) [3] and contained 1.38 billion triples, 18 unique predicates, and 10,000 universities. The triples were divided across 200,004 files and each file was treated as one RDF graph. The real dataset was BTC 2012 [6], which is widely used in the Semantic Web community. It contained 1.36 billion RDF quads with 57,000 unique predicates and 9.59 million RDF graphs.

For LUBM, the query set included 3 SPARQL queries with large, complex BGPs (L1-L3) and 9 others (L4-L12) that are variations of the queries in the LUBM benchmark. For BTC 2012, the query set also included 2 SPARQL queries with large, complex BGPs (B1, B2) and 5 others (B3-B7). (The queries are listed in Appendix A. Each query has a single BGP.) The number of triples patterns in each query and the number of results obtained for each query using the three approaches are shown in Table 2.

Table 2: Queries for LUBM and BTC-2012.

Dataset	Query	Type	# of BGPs	# of triple patterns	# of results
LUBM	L1	large	1	18	24
	L2	large	1	11	7,082
	L3	large	1	22	0
	L4	small	1	6	2,462
	L5	small	1	1	25,205,352
	L6	small	1	6	468,047
	L7	small	1	1	79,163,972
	L8	small	1	2	10,798,091
	L9	small	1	6	440,834
	L10	small	1	5	8,341
	L11	small	1	4	172
	L12	small	1	6	0
BTC-2012	B1	large	1	19	6
	B2	large	1	21	5
	B3	small	1	4	47,493
	B4	small	1	6	146,012
	B5	small	1	7	1,460,748
	B6	small	1	5	0
	B7	small	1	5	12,101,709
	B8	multi-bgp	5	8	249,318
	B9	multi-bgp	4	7	149,306
	B8	multi-bgp	5	8	249,318
	B9	multi-bgp	4	7	149,306

Table 3: Query processing times for LUBM. Best times are shown in bold within shaded cells. $X^{\dagger}$ indicates that the query ran for more than X seconds and was terminated. Aborted query (L7) was not included in the Geometric Mean calculation.

Query	**Cold cache** Time taken (in secs)			**Warm cache** Time taken (in secs)		
	RIQ	**RDF-3X**	**Jena TDB**	RIQ	**RDF-3X**	**Jena TDB**
L1	**16.47**	68.15	263.31	**1.75**	64.28	4.22
L2	**874.35**	$77,315^{\dagger}$	77,315	**74.33**	$64,637^{\dagger}$	64,637
L3	**18.26**	1692.95	179.19	2.51	1698.01	**1.97**
Geo. Mean (large)	**64.07**	2073.92	1539.39	**6.89**	1917.93	81.3
L4	**211.01**	1967.96	698.08	**133.98**	1898.87	664.75
L5	**555.96**	1392.1	1130.43	**499.2**	657.55	1127.37
L6	**437.76**	1156.04	1119.31	**374.13**	849.13	1144.11
L7	**844.51**	1455.9	aborted	**826.2**	1298.34	aborted
L8	**477.83**	1212.67	1346.17	408.79	**68.42**	1395.48
L9	**467.01**	1262	1137.38	**400.59**	1236.12	1142.73
L10	11.57	9.03	**7.15**	**0.56**	6.38	3.19
L11	11.01	**1.7**	5.79	0.32	**0.26**	1.13
L12	68.77	**20.92**	725.93	34.95	**19.75**	703.26
Geo. Mean (small)	**163.77**	235.94	282.57	**65.17**	116.04	207.72
Geo. Mean (all)	**129.52**	406.26	448.65	**37.16**	233.97	160.83

Algorithm 4 PruneBGPTree(node n)

Let $c_1, ..., c_\tau$ denote the child nodes of n (left-to-right) ignoring those corresponding to braces
if $eval[n]$ = `FALSE` **then**
  if n's parent is `NotExistsFunction` **then**
    return `TRUE`
  else if n is `OptionalGraphPattern` **then**
    return `FALSE`
  else if n is `GroupGraphPattern` & left-sibling is `UNION` **then**
    Prune away the subtree rooted at the left-sibling of n
  Prune away the subtree rooted at n from the BGP Tree
else
  for $i = 1$ to τ **do**
    $status \leftarrow$ PruneBGPTree(i)
    if $status$ = `FALSE` **then**
      Prune away the subtree rooted at i from the BGP Tree
return `TRUE`

5.2 Index Size

Here we report the size of the indexes built by the three approaches. Note that the size of the LUBM and BTC 2012 datasets were 217 GB and 218 GB, respectively. For LUBM, the size of the index built by RDF-3X and Jena TDB were 77 GB and 121 GB, respectively. The filtering index of RIQ was 5.7 GB in size and had 487 unions. For BTC 2012, the size of the index built by RDF-3X and Jena TDB were 87 GB and 110 GB, respectively. RIQ's filtering index was 6.5 GB in size and had 526 unions. When constructing the filters of LUBM and BTC 2012 in RIQ, we set the false positive rate ϵ equal to 1% and 5%, respectively.

Table 4: Query processing times for BTC-2012. Best times are shown in bold within shaded cells.

Query	**Cold cache** Time taken (in secs)			**Warm cache** Time taken (in secs)		
	RIQ	**RDF-3X**	**Jena TDB**	RIQ	**RDF-3X**	**Jena TDB**
B1	**5.21**	419.02	12.5	**1.16**	414.53	10.09
B2	**9.23**	365.44	17.18	**5.41**	365.34	15.21
Geo. Mean (large)	**6.93**	391.31	14.65	**2.51**	389.16	12.39
B3	**37.46**	194.93	357.34	1.82	**0.8**	13.42
B4	**36.12**	161.72	290.34	3.6	**2.35**	17.07
B5	**59.55**	188.78	319.41	**25.55**	28.43	46.16
B6	16.13	**0.41**	11.98	0.52	**0.16**	10.3
B7	**214.37**	286.64	245.3	183.61	**180.22**	124.47
Geo. Mean (small)	**48.87**	58.74	157.65	6.93	**4.34**	26.7
Geo. Mean (all)	**27.97**	100.98	79.97	**5.18**	15.68	21.44

Table 5: Multi-BGP Queries for BTC-2012.

Query	**# of BGPs**	**# of triple patterns**	**# of results**
B8	5	8	249,318
B9	4	7	149,306
B10	7	12	196
B11	2	4	347

Table 6: Query processing times for queries with multiple BGPs on BTC 2012. Best times are shown in bold within shaded cells.

Query	**Cold cache** Time taken (in secs)			**Warm cache** Time taken (in secs)		
	RIQ	RIQ **(w/o opt)**	**Jena TDB**	RIQ	RIQ **(w/o opt)**	**Jena TDB**
B8	**159.35**	180.78	663.31	76.76	85.56	**38.77**
B9	**154.97**	168.92	648.93	69.08	86.4	**33.36**
B10	**18.4**	724.13	3,564.44	**7.12**	452.42	369.81
B11	**262.08**	269.13	3,631.35	**81.89**	251.02	1671.1
Geo. Mean	**104.46**	277.75	1,536.36	**41.93**	170.22	168.14

5.3 Query Processing

We measured the wall-clock time taken to process each query in both cold and warm cache settings, and report the average over 3 runs. Jena TDB was executed with its default statistics-based optimization.

The results for LUBM and BTC 2012 are reported in Tables 3 and 4, respectively. For LUBM, RIQ processed queries with large, complex BGPs (L1-L3) significantly faster than RDF-3X and Jena TDB in both cold and warm cache settings. For BTC 2012, RIQ was significantly faster than RDF-3X in processing queries B1 and B2. This demonstrates that the *decrease-and-conquer* approach of RIQ is more effective than the popular join-based processing (by first matching individual triple patterns) on queries with large, complex BGPs. All of the large, complex queries had at least one undirected cycle. RIQ identified a maximum of 22 candidate groups for queries L1-L3 and 4 candidate groups for queries B1 and B2.

Next, we report the performance of RIQ on queries with (small) BGPs containing less than 8 triple patterns (L4-L12 and B3-B7). Interestingly, on LUBM, RIQ was faster than RDF-3X and Jena TDB for six out of the nine queries in both cold and warm cache settings. On BTC 2012, RIQ was the fastest in the cold cache setting for four out of the five queries. However, RDF-3X was the fastest in the warm cache setting for four out of the five queries. Finally, we compared the three approaches based on the geometric mean of their query processing times. Across all queries, RIQ was the winner for both LUBM and BTC 2012. For large BGP queries, RIQ was the winner in terms of the geometric mean. Only in one case, for the geometric mean across queries with small BGPs in the warm cache setting, RIQ lost to RDF-3X.

5.4 Queries with multiple BGPs

We also tested the query processing strategy of RIQ on SPARQL queries containing multiple BGPs and keywords like `OPTIONAL` and `UNION`. We ran the queries listed in Table 5 (B8-B11) on BTC 2012. This table shows the number of BGPs and triple patterns in each query along with the number of results produced. (The queries are listed in Appendix A.) Note that B10 and B11 are based on the DBpedia SPARQL Benchmark [23]. Since RDF-3X cannot process such queries, we did not use it for the comparison. Jena TDB indexed the context information in each triple of BTC 2012.

Table 6 shows the performance results. For baseline comparison, we also turned off the optimization feature in RIQ by (a) selecting a candidate group if any of the BGPs in the query have a match in the group, and (b) executing the original query on each candidate. This is shown as 'RIQ (w/o opt)' in the table. Clearly, RIQ was faster when the optimization was enabled. Also, RIQ outperformed Jena TDB for all the queries in the cold cache setting. For the warm cache setting, RIQ and Jena TDB ran faster than the other on two queries each. However, RIQ beat Jena TDB in terms of the geometric mean for the tested queries.

6 Conclusions

We presented RIQ, a new approach for indexing large RDF datasets containing quads. RIQ employs a *decrease-and-conquer* approach to efficiently process SPARQL queries. Through our experiments, we demonstrate that RIQ enables efficient SPARQL query processing on large RDF datasets with more than a billion quads. It significantly outperformed popular tools like RDF-3X and Jena TDB for queries with large, complex BGPs. It also achieved competitive performance for queries with small BGPs.

Acknowledgement

This work was supported by the National Science Foundation under Grant No. 1115871.

References

[1] Have semantic technologies crossed the chasm yet? https://semanticweb.com/have-semantic-technologies-crossed-the-chasm-yet_b16484.

[2] Linking Open Gov. Data. http://logd.tw.rpi.edu/.

[3] LUBM. http://swat.cse.lehigh.edu/projects/lubm/.

[4] Pfizer. https://semanticweb.com/tag/pfizer.

[5] Resource Descrip. Framework. http://www.w3.org/RDF.

[6] Seman. Web Challenge. http://challenge.semanticweb.org/.

[7] SPARQL 1.1. http://www.w3.org/TR/sparql11-query/.

[8] D. Abadi, A. Marcus, S. Madden, and K. Hollenbach. SW-Store: A vertically partitioned DBMS for semantic web data management. *VLDB Journal*, 18(2):385–406, 2009.

[9] M. Atre, V. Chaoji, M. J. Zaki, and J. A. Hendler. Matrix "Bit" loaded: A scalable lightweight join query processor for RDF data. In *Proc. of the 19th WWW Conference*, pages 41–50, 2010.

[10] S. Auer, C. Bizer, G. Kobilarov, J. Lehmann, and Z. Ives. DBpedia: A nucleus for a web of open data. In *Proc. of ISWC '07*, pages 11–15, 2007.

[11] D. Beckett. Raptor. http://librdf.org/raptor/.

[12] C. Bizer, T. Heath, and T. Berners-Lee. Linked Data - The story so far. *Int. Journal on Semantic Web and Information Systems*, 5(3):1–22, 2009.

[13] M. A. Bornea, J. Dolby, A. Kementsietsidis, K. Srinivas, P. Dantressangle, O. Udrea, and B. Bhattacharjee. Building an efficient RDF store over a relational database. In *Proc. of 2013 SIGMOD Conference*, pages 121–132, 2013.

[14] M. Bröcheler, A. Pugliese, and V. S. Subrahmanian. DOGMA: A disk-oriented graph matching algorithm for RDF databases. In *Proc. of ISWC '09*, pages 97–113, 2009.

[15] A. Broder. On the resemblance and containment of documents. In *Proc. of the Compress. and Complex. of Sequences*, pages 21–29, 1997.

[16] A. Broder and M. Mitzenmacher. Network applications of bloom filters: A survey. *Internet Mathematics*, 1(4):485–509, 2003.

[17] J. Broekstra, A. Kampman, and F. van Harmelen. Sesame: A generic architecture for storing and querying RDF and RDF Schema. In *Proc. of ISWC '02*, pages 54–68.

[18] E. I. Chong, S. Das, G. Eadon, and J. Srinivasan. An efficient SQL-based RDF querying scheme. In *Proc. of the 31st VLDB Conference*, pages 1216–1227, 2005.

[19] Dablooms. https://github.com/bitly/dablooms.

[20] T. H. Haveliwala, A. Gionis, D. Klein, and P. Indyk. Evaluating strategies for similarity search on the Web. In *Proc. of the 11th WWW Conference*, pages 432–442, 2002.

[21] J. Huang, D. J. Abadi, and K. Ren. Scalable SPARQL querying of large RDF graphs. *Proc. of VLDB Endow.*, 4(11):1123–1134, 2011.

[22] P. Indyk and R. Motwani. Approximate nearest neighbors: Towards removing the curse of dimensionality. In *Proc. of the 13th ACM STOC*, pages 604–613, 1998.

[23] M. Morsey, J. Lehmann, S. Auer, and A.-C. N. Ngomo. Dbpedia sparql benchmark: Performance assessment with real queries on real data. In *Proceedings of the 10th International Conference on The Semantic Web*, pages 454–469, Bonn, Germany, 2011.

[24] T. Neumann and G. Weikum. The RDF-3X engine for scalable management of RDF data. *The VLDB Journal*, 19(1):91–113, 2010.

[25] M. O. Rabin. Fingerprinting by random polynomials. Technical Report TR 15-81, Harvard University, 1981.

[26] O. Udrea, A. Pugliese, and V. S. Subrahmanian. GRIN: a graph based RDF index. In *Proc. of the 22nd National Conf. on Artificial Intelligence*, pages 1465–1470, 2007.

[27] C. Weiss, P. Karras, and A. Bernstein. Hexastore: Sextuple indexing for Semantic Web data management. *Proc. VLDB Endow.*, 1(1):1008–1019, 2008.

[28] K. Wilkinson. Jena property table implementation. In *SSWS 2006*, pages 35–46, Athens, GA, 2006.

[29] P. Yuan, P. Liu, B. Wu, H. Jin, W. Zhang, and L. Liu. TripleBit: A fast and compact system for large scale RDF data. *Proc. VLDB Endow.*, 6(7):517–528, 2013.

[30] K. Zeng, J. Yang, H. Wang, B. Shao, and Z. Wang. A distributed graph engine for Web Scale RDF data. *Proc. VLDB Endow.*, 6(4):265–276, Feb. 2013.

[31] L. Zou, J. Mo, L. Chen, M. T. Özsu, and D. Zhao. gStore: Answering SPARQL queries via subgraph matching. *Proc. VLDB Endow.*, 4:482–493, May 2011.

A Appendix

A.1 LUBM Queries

Large BGP queries L1, L2, and L3 are shown below. Figure 4 shows the visual representation of each query's BGP.

```
PREFIX ub: <http://www.lehigh.edu/~zhp2/2004/0401/univ-bench.owl#>
PREFIX rdf: <http://www.w3.org/1999/02/22-rdf-syntax-ns#>
```

L1:

```
SELECT ?p ?c ?e ?ph ?res ?uguni ?msuni ?phduni ?s1n ?s2n ?s1 ?s2 ?pub WHERE {
  GRAPH ?g {
    ?s1 ub:advisor ? .
    ?s1 ub:name ?s1n .
    ?s1 rdf:type ub:UndergraduateStudent .
    ?s2 ub:advisor ?p .
    ?s2 ub:name ?s2n .
    ?s2 rdf:type ub:GraduateStudent .
    ?p rdf:type ub:FullProfessor .
    ?p ub:name "FullProfessor7" .
    ?p ub:teacherOf ?c .
    ?p ub:undergraduateDegreeFrom ?uguni .
    ?p ub:mastersDegreeFrom ?msuni .
    ?p ub:doctoralDegreeFrom ?phduni .
    ?p ub:worksFor <http://www.Department17.University1001.edu> .
    ?p ub:emailAddress ?e .
    ?p ub:telephone ?ph .
    ?p ub:researchInterest ?res .
    ?pub ub:publicationAuthor ?p .
    ?pub ub:publicationAuthor ?s2 .
  }
}
```

L2:

```
SELECT ?s1 ?s2 ?pub ?uguni ?dept WHERE {
  GRAPH ?g {
```

```
    ?s1 rdf:type ub:GraduateStudent .
    ?s1 ub:undergraduateDegreeFrom ?uguni .
    ?s1 ub:memberOf ?dept .
    ?s2 rdf:type ub:GraduateStudent .
    ?s2 ub:undergraduateDegreeFrom ?uguni .
    ?dept rdf:type ub:Department .
    ?dept ub:subOrganizationOf <http://www.University1167.edu> .
    ?uguni rdf:type ub:University .
    ?pub rdf:type ub:Publication .
    ?pub ub:publicationAuthor ?s1 .
    ?pub ub:publicationAuthor ?s2 .
  }
}
```

L3:

```
SELECT ?p1 ?uni ?n1 ?e1 ?ph1 ?res1 ?c ?pub1 ?pub2 ?p2 ?n2 ?e2 ?ph2 ?res2 WHERE {
  GRAPH ?g {
    ?p1 rdf:type ub:FullProfessor .
    ?p1 ub:undergraduateDegreeFrom <http://www.University584.edu> .
    ?p1 ub:mastersDegreeFrom <http://www.University584.edu> .
    ?p1 ub:doctoralDegreeFrom <http://www.University429.edu> .
    ?p1 ub:worksFor ?uni .
    ?p1 ub:name ?n1 .
    ?p1 ub:emailAddress ?e1 .
    ?p1 ub:telephone ?ph1 .
    ?p1 ub:researchInterest ?res1 .
    ?p1 ub:teacherOf ?c .
    ?p2 rdf:type ub:AssociateProfessor .
    ?p2 ub:undergraduateDegreeFrom <http://www.University584.edu> .
    ?p2 ub:mastersDegreeFrom <http://www.University584.edu> .
    ?p2 ub:doctoralDegreeFrom <http://www.University9999.edu> .
    ?p2 ub:worksFor ?uni .
    ?p2 ub:name ?n2 .
    ?p2 ub:emailAddress ?e2 .
    ?p2 ub:telephone ?ph2 .
    ?p2 ub:researchInterest ?res2 .
    ?p2 ub:teacherOf ?course2 .
    ?pub1 ub:publicationAuthor ?p1 .
    ?pub2 ub:publicationAuthor ?p2 .
  }
}
```

Small BGP queries L4-L12 are shown below.

L4:

```
SELECT ?x ?y ?z WHERE {
  GRAPH ?g {
    ?z ub:subOrganizationOf ?y .
    ?y rdf:type ub:University .
    ?z rdf:type ub:Department .
    ?x ub:memberOf ?z .
    ?x rdf:type ub:GraduateStudent .
    ?x ub:undergraduateDegreeFrom ?y .
  }
```

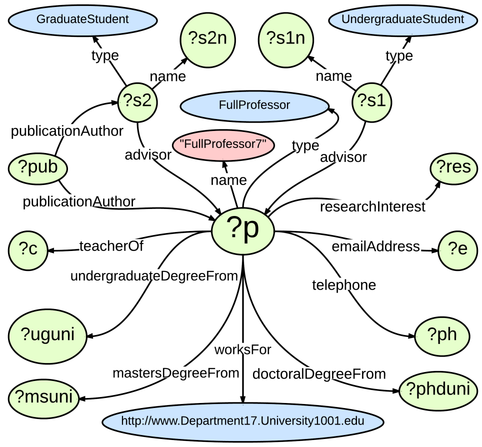

(a)

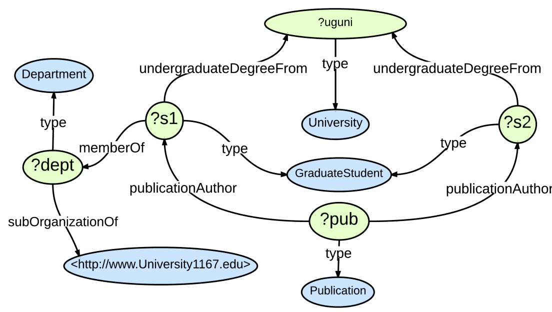

(b)

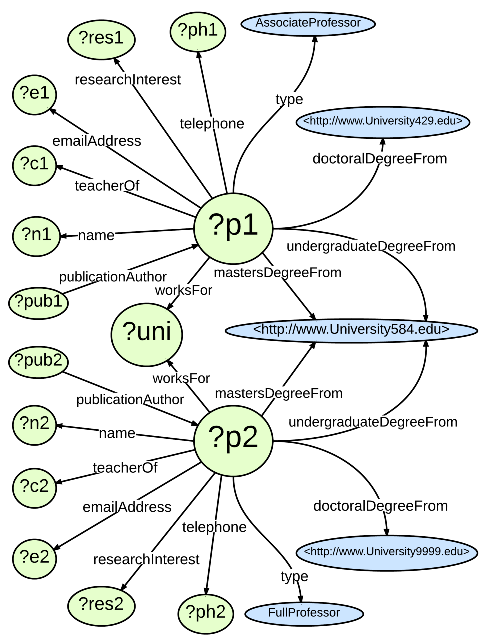

(c)

Figure 4: Visual representation of LUBM queries with large, complex BGPs

```
}
```

L5:

```
SELECT ?x WHERE {
  GRAPH ?g {
    ?x rdf:type ub:GraduateStudent .
  }
}
```

L6:

```
SELECT ?x ?y ?z WHERE {
  GRAPH ?g {
    ?x rdf:type ub:GraduateStudent .
    ?y rdf:type ub:AssistantProfessor .
    ?z rdf:type ub:GraduateCourse .
    ?x ub:advisor ?y .
    ?y ub:teacherOf ?z .
    ?x ub:takesCourse ?z .
  }
}
```

L7:

```
SELECT ?x WHERE {
  GRAPH ?g {
    ?x rdf:type ub:UndergraduateStudent .
  }
}
```

L8:

```
SELECT ?x WHERE {
  GRAPH ?g {
    ?x rdf:type ub:Course .
    ?x ub:name ?y .
  }
}
```

L9:

```
SELECT ?x ?y ?z WHERE {
  GRAPH ?g {
    ?y ub:teacherOf ?z .
    ?y rdf:type ub:FullProfessor .
    ?z rdf:type ub:Course .
    ?x ub:advisor ?y .
    ?x rdf:type ub:UndergraduateStudent .
    ?x ub:takesCourse ?z .
  }
}
```

L10:

```
SELECT ?x ?y ?z WHERE {
  GRAPH ?g {
    ?x rdf:type ub:UndergraduateStudent .
    ?y rdf:type ub:Department .
```

```
    ?x ub:memberOf ?y .
    ?y ub:subOrganizationOf <http://www.University0.edu> .
    ?x ub:emailAddress ?z .
  }
}
```

L11:

```
SELECT ?x ?y WHERE {
  GRAPH ?g {
    ?x rdf:type ub:FullProfessor .
    ?y rdf:type ub:Department .
    ?x ub:worksFor ?y .
    ?y ub:subOrganizationOf <http://www.University0.edu> .
  }
}
```

L12:

```
SELECT ?x ?y ?z WHERE {
  GRAPH ?g {
    ?x rdf:type ub:UndergraduateStudent .
    ?y rdf:type ub:University .
    ?z rdf:type ub:Department .
    ?x ub:memberOf ?z .
    ?z ub:subOrganizationOf ?y .
    ?x ub:undergraduateDegreeFrom ?y .
  }
}
```

A.2 BTC-2012 Queries

Large BGP queries B1 and B2 are shown below. Figure 5 shows the visual representation of each query's BGP.

```
PREFIX rdf: <http://www.w3.org/1999/02/22-rdf-syntax-ns#>
PREFIX xsd: <http://www.w3.org/2001/XMLSchema#>
PREFIX geo: <http://aims.fao.org/aos/geopolitical.owl#>
PREFIX collect: <http://purl.org/collections/nl/am/>
PREFIX ore: <http://www.openarchives.org/ore/terms/>
PREFIX fbase: <http://rdf.freebase.com/ns/>
```

B1:

```
SELECT ?s1 ?o1 ?s2 WHERE {
  GRAPH ?g {
    ?s1 collect:acquisitionDate "1980-05-16" .
    ?s1 collect:acquisitionMethod collect:t-14382 .
    ?s1 collect:associationSubject ?o1 .
    ?s1 collect:contentMotifGeneral collect:t-8782 .
    ?s1 collect:creditLine collect:t-14773 .
    ?s1 collect:material collect:t-3249 .
    ?s1 collect:objectCategory collect:t-15606 .
    ?s1 collect:objectName collect:t-10444 .
    ?s1 collect:objectNumber "KA 17150" .
    ?s1 collect:priref "23182" .
    ?s1 collect:productionDateEnd "1924" .
    ?s1 collect:productionDateStart "1924" .
    ?s1 collect:productionPlace collect:t-624 .
    ?s1 collect:title "Plate commemorating the first Amsterdam-Batavia flight"@en .
    ?s1 ore:proxyFor collect:physical-23182 .
    ?s1 ore:proxyIn collect:aggregation-23182 .
    ?s1 collect:relatedObjectReference ?s2 .
    ?s2 collect:relatedObjectReference ?s1 .
  }
}
```

B2:

```
SELECT ?u ?un ?cnt1 ?ctry1 ?on1 ?cnt2 ?ctry2 ?on2 WHERE {
  GRAPH ?g {
    ?u geo:nameShortEN ?un .
    ?u geo:hasMember ?ctry1 .
    ?u rdf:type geo:economic_region .
    ?cnt1 geo:hasMember ?ctry1 .
    ?cnt1 rdf:type geo:geographical_region .
    ?cnt1 geo:nameShortEN "Africa"^^xsd:string .
    ?cnt2 geo:hasMember ?ctry2 .
    ?cnt2 rdf:type geo:geographical_region .
    ?cnt2 geo:nameShortEN "Asia"^^xsd:string .
    ?ctry1 geo:nameOfficialEN ?on1 .
    ?ctry1 geo:isInGroup ?u .
    ?ctry1 geo:isInGroup ?cnt1 .
    ?ctry1 geo:isInGroup geo:World .
    ?ctry1 rdf:type geo:self_governing .
    ?ctry1 geo:hasBorderWith ?ctry2 .
```

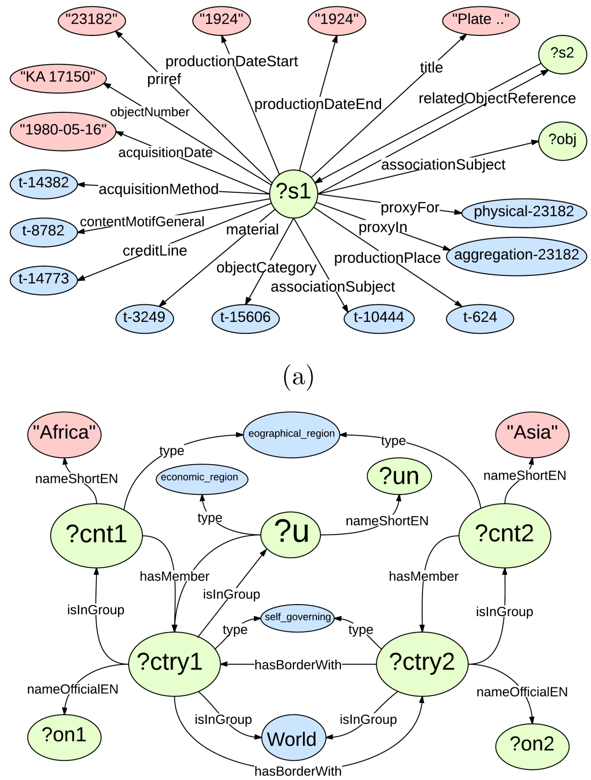

Figure 5: Visual representation of BTC queries with large, complex BGPs

```
    ?ctry2 geo:nameOfficialEN ?on2 .
    ?ctry2 geo:isInGroup ?cnt2 .
    ?ctry2 geo:isInGroup geo:World .
    ?ctry2 rdf:type geo:self_governing .
    ?ctry2 geo:hasBorderWith ?ctry1 .
  }
}
```

Small BGP queries B3-B7 are shown below.

B3:

```
SELECT ?fperf ?actor ?film ?name ?rel WHERE {
  GRAPH ?g {
    ?fperf fbase:film.performance.actor ?actor .
    ?fperf fbase:film.performance.film ?film .
    ?film fbase:type.object.name ?name .
    ?film fbase:film.film.initial_release_date ?rel .
  }
}
```

B4:

```
SELECT ?fperf ?actor ?film ?name ?rel ?lang ?gen WHERE {
```

```
  GRAPH ?g {
    ?fperf fbase:film.performance.actor ?actor .
    ?fperf fbase:film.performance.film ?film .
    ?film fbase:type.object.name ?name .
    ?film fbase:film.film.initial_release_date ?rel .
    ?film fbase:film.film.language ?lang .
    ?film fbase:film.film.genre ?gen .
  }
}
```

B5:

```
SELECT ?fperf ?actor ?film ?name ?rel ?lang ?gen ?star WHERE {
  GRAPH ?g {
    ?fperf fbase:film.performance.actor ?actor .
    ?fperf fbase:film.performance.film ?film .
    ?film fbase:type.object.name ?name .
    ?film fbase:film.film.initial_release_date ?rel .
    ?film fbase:film.film.language ?lang .
    ?film fbase:film.film.genre ?gen .
    ?film fbase:film.film.starring ?star .
  }
}
```

B6:

```
SELECT ?p1 ?p2 ?p1n ?p2n ?loc WHERE {
  GRAPH ?g {
    ?p1 fbase:people.place_lived.person ?p1n .
    ?p1 fbase:people.place_lived.location ?loc .
    ?p2 fbase:people.place_lived.person ?p2n .
    ?p2 fbase:people.place_lived.location ?loc .
    ?loc fbase:location.location.containedby fbase:en.iraq .
  }
}
```

B7:

```
SELECT ?s ?x ?y ?z ?w ?t WHERE {
  GRAPH ?g {
    ?s fbase:location.location.events ?x .
    ?s fbase:location.location.geolocation ?y .
    ?s fbase:location.location.people_born_here ?z .
    ?s fbase:location.location.people_born_here ?w .
    ?s fbase:location.location.containedby ?t .
  }
}
```

B8:

```
PREFIX resource: <http://dbpedia.org/resource/>
PREFIX ontology: <http://dbpedia.org/ontology/>


SELECT ?city ?area ?code ?zone ?abstract ?postal ?water ?popu ?g
WHERE {
  GRAPH ?g {
```

```
    { ?city ontology:areaLand ?area .
      ?city ontology:areaCode ?code . }
    UNION
    { ?city ontology:timeZone ?zone .
      ?city ontology:abstract ?abstract . }
    ?city ontology:country resource:United_States .
    ?city ontology:postalCode ?postal .
    OPTIONAL { ?city ontology:areaWater ?water . }
    OPTIONAL { ?city ontology:populationTotal ?popu . }
  }
}
```

B9:

```
PREFIX res: <http://dbpedia.org/resource/>
PREFIX onto: <http://dbpedia.org/ontology/>

SELECT ?city ?area ?code ?zone ?abstract ?postal ?popu ?g
WHERE {
GRAPH ?g {
    { ?city onto:areaLand ?area .
    ?city onto:areaCode ?code . }
    UNION
    { ?city onto:timeZone ?zone .
    ?city onto:abstract ?abstract . }
    ?city onto:country res:United_States .
    ?city onto:postalCode ?postal .
    OPTIONAL { ?city onto:populationTotal ?popu . }
    }
}
```

B10:

```
PREFIX rdfs: <http://www.w3.org/2000/01/rdf-schema#>
PREFIX geo: <http://www.w3.org/2003/01/geo/wgs84_pos#>
PREFIX foaf: <http://xmlns.com/foaf/0.1/>

SELECT *
WHERE {
GRAPH ?g {
    ?var6 a <http://dbpedia.org/ontology/PopulatedPlace> .
    ?var6 <http://dbpedia.org/ontology/abstract> ?var1 .
    ?var6 rdfs:label ?var2 .
    ?var6 geo:lat ?var3 .
    ?var6 geo:long ?var4 .
    {
        ?var6 rdfs:label "Brunei"@en .
    }
    UNION
    {
        ?var5 <http://dbpedia.org/property/redirect> ?var6 .
        ?var5 rdfs:label "Brunei"@en .

    }
    OPTIONAL { ?var6 foaf:depiction ?var8 }
    OPTIONAL { ?var6 foaf:homepage ?var10 }
```

```
    OPTIONAL { ?var6 <http://dbpedia.org/ontology/populationTotal> ?var12 }
    OPTIONAL { ?var6 <http://dbpedia.org/ontology/thumbnail> ?var14 }
    }
}
```

B11:

```
PREFIX foaf: <http://xmlns.com/foaf/0.1/>
SELECT ?var5 ?var6 ?var8 ?var9 ?var4
WHERE {
    GRAPH ?g {
        {
            <http://my.opera.com/Green%20Aya/xml/foaf#me> ?var5 ?var6 .
            ?var6 foaf:name ?var8 .
        } UNION {
            ?var9 ?var5 <http://my.opera.com/Green%20Aya/xml/foaf#me> .
            ?var9 foaf:name ?var4 .
        }
    }
}
```