

Taming Subgraph Isomorphism for RDF Query Processing

Jinha Kim ^{†#} Hyungyu Shin [†] Wook-Shin Han ^{†*}
jinha.kim@oracle.com hgshin@dblab.postech.ac.kr wshan@postech.ac.kr

Sungpack Hong [#] Hassan Chafi [#]
{sungpack.hong, hassan.chafi}@oracle.com

[†]POSTECH, South Korea [#]Oracle Labs, USA

ABSTRACT

RDF data are used to model knowledge in various areas such as life sciences, Semantic Web, bioinformatics, and social graphs. The size of real RDF data reaches billions of triples. This calls for a framework for efficiently processing RDF data. The core function of processing RDF data is subgraph pattern matching. There have been two completely different directions for supporting efficient subgraph pattern matching. One direction is to develop specialized RDF query processing engines exploiting the properties of RDF data for the last decade, while the other direction is to develop efficient subgraph isomorphism algorithms for general, labeled graphs for over 30 years. Although both directions have a similar goal (i.e., finding subgraphs in data graphs for a given query graph), they have been independently researched without clear reason. We argue that a subgraph isomorphism algorithm can be easily modified to handle the graph homomorphism, which is the RDF pattern matching semantics, by just removing the injectivity constraint. In this paper, based on the state-of-the-art subgraph isomorphism algorithm, we propose an in-memory solution, `TurboHOM++`, which is tamed for the RDF processing, and we compare it with the representative RDF processing engines for several RDF benchmarks in a server machine where billions of triples can be loaded in memory. In order to speed up `TurboHOM++`, we also provide a simple yet effective transformation and a series of optimization techniques. Extensive experiments using several RDF benchmarks show that `TurboHOM++` consistently and significantly outperforms the representative RDF engines. Specifically, `TurboHOM++` outperforms its competitors by up to *five* orders of magnitude.

1. INTRODUCTION

The Resource Description Framework (RDF) is a standard for representing knowledge on the web. It is primarily designed for building the Semantic web and has been widely adopted in database and data mining communities. RDF models a fact as a triple which consists of a subject (S), a predicate (P), and an object (O). Due to its simple structure, many practitioners materialize their data in

*corresponding author

This work is licensed under the Creative Commons Attribution-NonCommercial-NoDerivs 3.0 Unported License. To view a copy of this license, visit <http://creativecommons.org/licenses/by-nc-nd/3.0/>. Obtain permission prior to any use beyond those covered by the license. Contact copyright holder by emailing info@vldb.org. Articles from this volume were invited to present their results at the 41st International Conference on Very Large Data Bases, August 31st - September 4th 2015, Kohala Coast, Hawaii.

Proceedings of the VLDB Endowment, Vol. 8, No. 11
Copyright 2015 VLDB Endowment 2150-8097/15/07.

an RDF format. For example, RDF datasets are now pervasive in various areas including life sciences, bioinformatics, and social networks. The size of real RDF data reaches billions of triples. Such billion-scale RDF data are fully loaded in main memory of today's server machine (The cost of a 1TB machine is less than \$40,000).

The SPARQL query language is a standard language for querying RDF data in a declarative fashion. Its core function is subgraph pattern matching, which corresponds to finding all graph homomorphisms in the data graph for a query graph [19].

In recent years, there have been significant efforts to speed up the processing of SPARQL queries by developing novel RDF query processing engines. Many engines [1, 18, 19, 25, 26, 29] model RDF data as tabular structures and process SPARQL queries using specialized join methods. For example, RDF-3X [19] treats RDF data as an edge table, `EDGE(S,P,O)`, and materializes six different orderings for this table, so that it can support many SPARQL queries just by using merge based join. Note that this approach is efficient for both disk-based and in-memory environments since merge join exploits only sequential scans. Some engines [2, 30, 35] treat RDF data as graphs (or matrices) and develop specialized graph processing methods for processing SPARQL queries. For example, `gStore` [35] uses specialized index structures to process SPARQL queries. Note that these index structures are based on `gCode` [34], which was originally proposed for graph indexing.

Subgraph isomorphism, on the other hand, has been studied since the 1970s. The representative algorithms are VF2 [20], QuickSI [21], GraphQL [11], GADDI [32], SPATH [33], and `TurboISO` [9]. In order to speed up performance, these algorithms exploit good matching orders and effective pruning rules. A recent study [14] shows that good subgraph isomorphism algorithms significantly outperform graph indexing based ones. However, all of these algorithms use only small graphs in their experiments, and thus, it still remains unclear whether these algorithms can show good performance for billion-scale graphs such as RDF data.

Although subgraph isomorphism processing and RDF query processing have similar goals (i.e., finding subgraphs in data graphs for a given query graph), they have two inexplicably different directions. A subgraph isomorphism algorithm can be easily modified to handle the graph homomorphism, which is the RDF pattern matching semantics, just by removing the injectivity constraint.

In this paper, based on the state-of-the-art subgraph isomorphism algorithm [9], we propose an in-memory solution, `TurboHOM++`, which is tamed for the RDF processing, and we compare it with the representative RDF processing engines for several RDF benchmarks in a server machine where billions of triples can be loaded in memory. We believe that this approach opens a new direction for RDF processing so that both traditional directions can merge or benefit from each other.

By transforming RDF graphs into labeled graphs, we can apply subgraph homomorphism methods to RDF query processing. Extensive experiments using several benchmarks show that a direct modification of `TurboISO` outperforms the RDF processing engines for queries which require a small amount of graph exploration. However, for some queries which require a large amount of graph exploration, the direct modification is slower than some of its competitors. This poses an important research question: “*Is this phenomenon due to inherent limitations of the graph homomorphism (subgraph isomorphism) algorithm?*” Our profile results show that two major subtasks of `TurboISO` — 1) exploring candidate subgraphs in *ExploreCandidateRegion* and 2) enumerating solutions based on candidate regions in *SubgraphSearch* — require performance improvement. `TurboHOM++` resolves such performance hurdles by proposing the type-aware transformation and tailored optimization techniques.

First, in order to speed up *ExploreCandidateRegion*, we propose a novel transformation (Section 4.1), called *type-aware transformation*, which is simple yet effective in processing SPARQL queries. In type-aware transformation, by embedding the types of an entity (i.e., a subject or object) into a vertex label set, we can eliminate corresponding query vertices/edges from a query graph. With type-aware transformation, the query graph size decreases, its topology becomes simpler than the original query, and thus, this transformation improves performance accordingly by reducing the amount of graph exploration.

In order to optimize performance in depth, in both *ExploreCandidateRegion* and *SubgraphSearch*, we propose a series of optimization techniques (Section 4.3), each of which contributes to performance improvement significantly for such slow queries. In addition, we explain how `TurboHOM++` is extended to support 1) general SPARQL features such as `OPTIONAL`, and `FILTER`, and 2) parallel execution for `TurboHOM++` in a non-uniform memory access (NUMA) architecture [15, 16]. These general features are necessary to execute comprehensive benchmarks such as Berlin SPARQL benchmark (BSBM) [3]. Note also that, when the RDF data size grows large, we have to rely on the NUMA architecture.

Extensive experiments using several representative benchmarks show that `TurboHOM++` consistently and significantly outperforms all its competitors for all queries tested. Specifically, our method outperforms the competitors by up to five orders of magnitude with only a single thread. This indicates that a subgraph isomorphism algorithm tamed for RDF processing can serve as an in-memory RDF accelerator on top of a commercial RDF engine for *real-time* RDF query processing.

Our contributions are as follows. 1) We provide the first direct comparison between RDF engines and the state-of-the-art subgraph isomorphism method tamed for RDF processing, `TurboHOM++`, through extensive experiments and analyze experimental results in depth. 2) In order to simplify a query graph, we propose a novel transformation method called type-aware transformation, which contributes to boosting query performance. 3) In order to speed up query performance further, we propose a series of performance optimizations as well as NUMA-aware parallelism for fast RDF query processing. 4) Extensive experiments using several benchmarks show that the optimized subgraph isomorphism method consistently and significantly outperforms representative RDF query processing engines.

The rest of the paper is organized as follows. Section 2 describes the subgraph isomorphism, its state-of-the-art algorithms, `TurboISO`, and their modification for the graph homomorphism. Section 3 presents how a direct modification of `TurboISO`, `TurboHOM`, handles the SPARQL pattern matching. Section 4 describes how we

obtain `TurboHOM++` from `TurboHOM` using the type-aware transformation and optimizations for the efficient SPARQL pattern matching. Section 5 reviews the related work. Section 6 presents the experimental result. Finally, Section 7 presents our conclusion. Note that due to the space limit, please refer [13] for how `TurboHOM++` handle `OPTIONAL`, `UNION`, `FILTER` keywords, and parallelize.

2. PRELIMINARY

2.1 Subgraph Isomorphism and RDF Pattern Matching Semantic

Suppose that a labeled graph is defined as $g(V, E, L)$, where V is a set of vertices, $E(\subseteq V \times V)$ is a set of edges, and L is a labeling function which maps from a vertex or an edge to the corresponding label set or label, respectively. Then, the subgraph isomorphism is defined as follows.

Definition 1. [14] Given a query graph $q(V, E, L)$ and a data graph $g(V', E', L')$, a *subgraph isomorphism* is an injective function $M : V \rightarrow V'$ such that 1) $\forall v \in V, L(v) \subseteq L'(M(v))$ and 2) $\forall (u, v) \in E, (M(u), M(v)) \in E'$ and $L(u, v) = L'(M(u), M(v))$.

If a query vertex, u , has a blank label set (or does not specify vertex label equivalently), it can match any data vertex. Here, $L(u) = \emptyset$, and thus, the subset condition, $L(u) \subseteq L'(M(u))$, is always satisfied. Similarly, if a query edge (u, v) has a blank label, it can match any data edge by generalizing the equality condition $L(u, v) = L'(M(u), M(v))$ to $L(u, v) \subseteq L'(M(u), M(v))$.

The graph homomorphism [6] is easily obtained from the subgraph isomorphism by just removing the injective constraint on M in Definition 1. Even though the RDF pattern matching semantics is based on the graph homomorphism, to answer SPARQL queries which have variables on predicates, a mapping from a query edge to an edge label is also required. We call such graph homomorphism the *e(extended)-graph homomorphism* and present a formal definition for it as follows.

Definition 2. Given a query graph $q(V, E, L)$ and a data graph $g(V', E', L')$, an *e(extended)-graph homomorphism* is a pair of two mapping functions, a query vertex to data vertex function $M_v : V \rightarrow V'$ such that 1) $\forall v \in V, L(v) \subseteq L'(M_v(v))$ and 2) $\forall (u, v) \in E, (M_v(u), M_v(v)) \in E'$, and $L(u, v) = L'(M_v(u), M_v(v))$, and a query edge to edge label function $M_e : V \times V \rightarrow L$ such that $\forall (u, v) \in E, M_e(u, v) = L'(M_v(u), M_v(v))$.

The subgraph isomorphism problem (resp. the e-graph homomorphism problem) is to find all distinct subgraph isomorphisms (resp. e-graph homomorphisms) of a query graph in a data graph.

Figure 1 shows a query q_1 and a data graph g_1 . In q_1 , $_$ means a blank vertex label set or blank edge label. In the subgraph isomorphism, there is only one solution – $M^1 = \{(u_0, v_0), (u_1, v_1), (u_2, v_2), (u_3, v_3), (u_4, v_4)\}$. In the e-graph homomorphism, there are three solutions – $M_v^1 = M^1$, $M_e^1 = \{((u_0, u_1), a), ((u_0, u_4), b), ((u_2, u_1), a), ((u_2, u_3), a), ((u_3, u_4), c)\}$, $M_v^2 = \{(u_0, v_2), (u_1, v_3), (u_2, v_2), (u_3, v_3), (u_4, v_5)\}$, $M_e^2 = M_e^1$, and $M_v^3 = \{(u_0, v_2), (u_1, v_1), (u_2, v_2), (u_3, v_3), (u_4, v_5)\}$, $M_e^3 = M_e^1$.

2.2 Turbo_{ISO}

In this subsection, we introduce the state-of-the-art subgraph isomorphism solution, `TurboISO`[9], and its modification for the e-graph homomorphism. Although we only describe the modification of `TurboISO` for the e-graph homomorphism, such modification is applicable to other subgraph isomorphism algorithms including




Figure 1: Example of subgraph isomorphism and e-graph homomorphism.

VF2 [20], QuickSI [21], GraphQL [11], GADDI [32], and SPATH [33], since all of the subgraph algorithms mentioned are instances of a generic subgraph isomorphism framework [14].

Turbo_{ISO} presents an effective method for the notorious *matching order* problem from which all the previous subgraph isomorphism algorithms have suffered [14]. Figure 2 illustrates an example of the matching order problem, where q_2 is the query graph, and g_2 is the data graph¹. Note that this example query results in no answers. However, the time to finish this query can differ drastically by how one chooses the matching order, as it leads to different number of comparisons. For instance, a matching order $\langle u_0, u_2, u_1, u_3 \rangle$ requires $1 + 10000 * 10 * 5$ comparisons while a different matching order $\langle u_0, u_3, u_1, u_2 \rangle$ requires only $1 + 5 * 10$ comparisons.

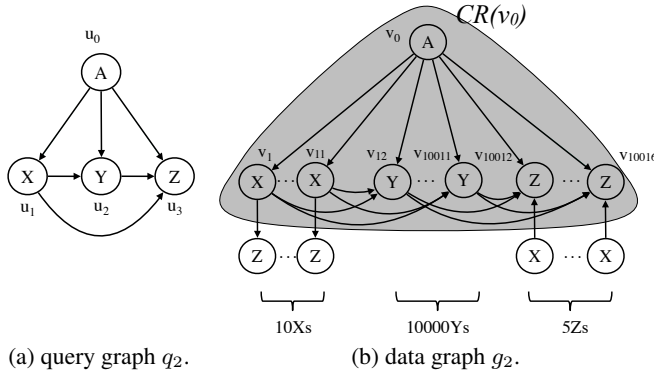


Figure 2: Example of showing the matching order problem.

Turbo_{ISO} solves the matching order problem with *candidate region exploration*, a technique that accurately estimates the number of candidate vertices for a given query path [9]. In particular, Turbo_{ISO} first identifies candidate data subgraphs (i.e., candidate regions) from the starting vertices (e.g. the shaded area in Figure 2b), then explores each region by performing a depth-first search, which allows almost exact selectivity for each query path.

Algorithm 1 outlines the overall procedure of Turbo_{ISO} in detail. First, if a query graph has only one vertex u and no edge, it is sufficient to retrieve all data vertices which have u 's labels ($= V(g)_{L(u)}$) and to find a subgraph isomorphism for each of them (lines 2–4). Otherwise, it selects the starting query vertex from the query graph (line 6). Then, it transforms the query graph into its corresponding query tree (line 7). After getting the query tree, for each data vertex that contains the vertex label of the starting query vertex, the candidate region is obtained by exploring the data graph

¹For simplicity, we omit the edge labels and allow only one vertex label in the data graph.

(lines 9). If the candidate region is not empty, its matching order is determined (line 11). The data vertex, v_s , is mapped to the first query vertex u_s by assigning $M(u_s) = v_s$ and $F(v_s) = true$ where $F : V \rightarrow boolean$ is a function which checks whether a data vertex is mapped or not (line 12). Then, the remaining subgraph matching is conducted (line 13). Lastly, the mapping (u_s, v_s) is restored by removing the mapping for u_s and assigning $F(v_s) = false$ (line 14).

Algorithm 1 Turbo_{ISO}($g(V, E, L), q(V', E', L')$)

Require: q : query graph, g : data graph
Ensure: all subgraph isomorphisms from q to g .

- 1: **if** $V(q) = \{u\}$ and $E = \phi$ **then**
- 2: **for each** $v \in V(g)_{L(u)}$ **do**
- 3: report $M = \{(u, v)\}$
- 4: **end for**
- 5: **else**
- 6: $u_s \leftarrow ChooseStartQueryVertex(q, g)$
- 7: $q' \leftarrow WriteQueryTree(q, u_s)$
- 8: **for each** $v_s \in \{v | v \in V, L(u_s) \subseteq L(v)\}$ **do**
- 9: $CR \leftarrow ExploreCandidateRegion(u_s, v_s)$
- 10: **if** CR is not empty **then**
- 11: $order \leftarrow DetermineMatchingOrder(q', CR)$
- 12: $UpdateState(M, F, u_s, v_s)$
- 13: $SubgraphSearch(q, q', g, CR, order, 1)$
- 14: $RestoreState(M, F, u_s, v_s)$
- 15: **end if**
- 16: **end for**
- 17: **end if**

ChooseStartQueryVertex. *ChooseStartQueryVertex* tries to pick the starting query vertex which has the least number of candidate regions. First, as a rough estimation, the query vertices are ranked by their scores. The score of a query vertex u is $rank(u) = \frac{freq(g, L(u))}{deg(u)}$, where $freq(g, L(u))$ is the number of data vertices that have u 's vertex labels. The score function prefers lower frequencies and higher degrees. After obtaining the top- k least-scored query vertices, the number of candidate regions is more accurately estimated for each of them by using the degree filter and the neighborhood label frequency (NLF) filter. The degree filter qualifies the data vertices which have equal or higher degree than their corresponding query vertices. The NLF filter qualifies the data vertices which have equal or larger number of neighbors for all distinct labels of the query vertex. In Figure 2, for example, u_0 becomes the starting query vertex since it has the least number of candidate regions ($= 1$).

WriteQueryTree. Next, *WriteQueryTree* transforms the query graph to the query tree. From the starting query vertex obtained by *ChooseStartQueryVertex*, a breath-first tree traversal is conducted. Every non-tree edge (u, v) of the query graph also is recorded in the corresponding query tree. For example, when u_0 is the starting query vertex, the non-tree edges of q_2 's query tree are $(u_1, u_2), (u_1, u_3),$ and (u_2, u_3) .

ExploreCandidateRegion. Using the query tree and the starting query vertex, *ExploreCandidateRegion* collects the candidate regions. A candidate region is obtained by exploring the data graph from the starting query vertex in a depth-first manner following the topology of the query tree. During the exploration, the injectivity constraint should be enforced. The shaded area of Figure 2b is the candidate region $CR(v_0)$ based on q_2 's query tree. Note that the candidate region expansion is conducted only after the current data vertex satisfies the constraints of the degree filter and the NLF filter.

DetermineMatchingOrder. After obtaining the candidate regions for a starting data vertex, the matching order is determined for each candidate region. Using the candidate region, *DetermineMatchingOrder* can accurately estimate the number of candidate vertices for each query path. Then, it orders all query paths in the query tree by the number of candidate vertices. For example, from $CR(v_0)$, the ordered list of query paths is $[u_0.u_3, u_0.u_1, u_0.u_2]$. Thus, we can easily see that $\langle u_0, u_3, u_1, u_2 \rangle$ is the best matching order based on this ordered list.

SubgraphSearch. Exploiting the data structures obtained from the previous steps, *SubgraphSearch* (Algorithm 2) enumerates all distinct subgraph isomorphisms. It first determines the current query vertex u from a given matching order $order$ (line 1). Then, it obtains a set of data vertices, C_R from a candidate region CR (line 2). $CR(u, v)$ represents the candidate vertices of a query vertex u which are the children of v in CR , and $P(q', u)$ is the parent of u in a query tree q' . For each candidate data vertex v , if v has already been mapped, the current solution is rejected since it violates the injectivity constraint of the subgraph isomorphism (lines 4–6). Next, by calling *IsJoinable*, if the query vertex u of the current data vertex v has non-tree edges, the existence of the corresponding edges are checked in the data graph (line 7). For example, given $CR(v_0)$ and the matching order $\langle u_0, u_3, u_1, u_2 \rangle$, when making the embedding for u_1 , we must check whether there is an edge from $M(u_1)$ to $M(u_3)$. If the *IsJoinable* test is passed, the mapping information is updated by assigning $M(u) = v$ and $F(v) = true$ (line 8). After updating the mapping, if all query vertices are mapped, a subgraph isomorphism M is reported (lines 9–10). Otherwise, further subgraph search is conducted (line 12). Finally, all changes done by *UpdateState* are restored (line 14).

Algorithm 2 *SubgraphSearch*($q, q', g, CR, order, d_c$)

```

1:  $u \leftarrow order[d_c]$ 
2:  $C_R \leftarrow CR(u, M(P(q', u)))$ 
3: for each  $v \in C_R$  such that  $v$  is not yet matched do
4:   if  $F(v) = true$  then
5:     continue
6:   end if
7:   if IsJoinable( $q, g, M, u, v, \dots$ ) then
8:     UpdateState( $M, F, u, v$ )
9:     if  $|M| = V(q)$  then
10:      report  $M$ 
11:     else
12:       SubgraphSearch( $q, q', g, CR, order, d_c + 1$ )
13:     end if
14:     RestoreState( $M, F, u, v$ )
15:   end if
16: end for

```

Modifying Turbo_{ISO} for e-Graph Homomorphism. We first explain how the generic subgraph isomorphism algorithm [14] can easily handle graph homomorphism. The generic subgraph isomorphism algorithm is implemented as a backtrack algorithm, where we find solutions by incrementing partial solutions or abandoning them when it is determined that they cannot be completed. Here, given a query graph q and its matching order $(u_{\sigma(1)}, u_{\sigma(2)}, \dots, u_{\sigma(|V(q)|)})$, a solution is modeled as a vector $\vec{v} = (M(u_{\sigma(1)}), M(u_{\sigma(2)}), \dots, M(u_{\sigma(|V(q)|)}))$ where each element in \vec{v} is a data vertex for the corresponding query vertex in the matching order. At each step in the backtrack algorithm, if a partial solution is given, we extend it by adding every possible candidate data vertex at the end. Here, any candidate data vertex that does not satisfy the following three conditions must be pruned.

- 1) $\forall u_i \in V(q), L(u_i) \subseteq L(M(u_i))$
- 2) $\forall (u_i, u_j) \in E(q), (M(u_i), M(u_j)) \in E(g)$ and $L(u_i, u_j) = L(M(u_i), M(u_j))$
- 3) $M(u_i) \neq M(u_j)$ if $u_i \neq u_j$

Note that the third condition ensures the injective condition, guaranteeing that no duplicate data vertex exists in each solution vector. Thus, by just disabling the third condition, the generic subgraph isomorphism algorithm finds all possible homomorphisms.

Now, we describe how to disable the third condition in Turbo_{ISO}, which is an instance of the generic subgraph isomorphism algorithm. Turbo_{ISO} uses pruning rules by applying filters in *ExploreCandidateRegion* and *SubgraphSearch*. First, the degree filter and the NLF filter should be modified since a data vertex can be mapped to multiple query vertices. The degree filter qualifies data vertices which have an equal number or more neighbors than distinct labels of their corresponding query vertices. The NLF filter qualifies data vertices which have at least one neighbor for all distinct labels of their corresponding query vertices. Second, lines 4–6 of *SubgraphSearch* ensuring the third condition should be removed in order to disable the injectivity test. As we see here, with minimal modification to Turbo_{ISO}, it can easily support graph homomorphism.

In order to make Turbo_{ISO} handle the e-graph homomorphism, the query edge to edge label mapping, M_e , should be additionally added in *SubgraphSearch*. For this, *UpdateState* assigns $M_e(P(q', u), u) = L(M_v(P(q', u)), M_v(u))$, and *RestoreState* removes such mapping. From here on, let us denote Turbo_{ISO} modified for the e-graph homomorphism as Turbo_{HOM}.

3. RDF QUERY PROCESSING BY E-GRAPH HOMOMORPHISM

In this section, we discuss how RDF datasets can be naturally viewed as graphs (Section 3.1), and thus how an RDF dataset can be directly transformed into a corresponding labeled graph (Section 3.2). After such a transformation, henceforth, the subgraph isomorphism algorithms modified for the e-graph homomorphism such as Turbo_{HOM} can be applied for processing SPARQL queries.

3.1 RDF as Graph

An RDF dataset is a collection of triples each of which consists of a subject, a predicate, and an object. By considering triples as directed edges, an RDF dataset naturally becomes a directed graph: the subjects and the objects are vertices while the predicates are edges. Figure 3 is a graph representation of triples that captures type relationships between university organizations. Note that we use rectangles to represent vertices in RDF graphs to distinguish them from the labeled graphs.

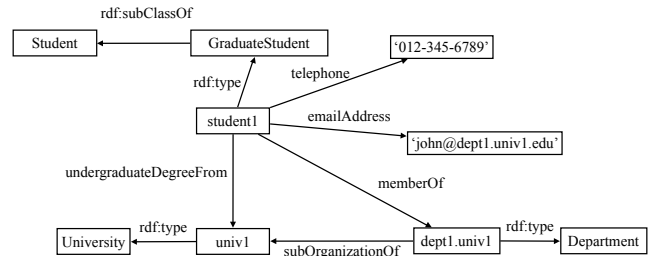


Figure 3: RDF graph.

3.2 Direct Transformation

To apply subgraph isomorphism algorithms modified for e-graph homomorphism (e.g. Turbo_{HOM}) for RDF query processing, RDF graphs have to be transformed into labeled graphs first.


The most basic way to transform RDF graphs is (1) to map subjects and objects to vertex IDs and (2) to map predicates to edge labels. We call such transformation the *direct transformation* because the topology of the RDF graph is kept in the labeled graph after the transformation. The vertex label function $L(v)(v \in V(g))$ is the identity function (i.e. $L(v) = \{v\}$).

Figure 4 shows the result of the direct transformation of Figure 3 – Figures 4a, 4b, and 4c are the vertex mapping table, the edge label mapping table, and the transformed graph, respectively.

Subject/Object	Vertex	Predicate	Edge Label
GraduateStudent	v_0	rdf:type	a
Student	v_1	rdf:subClassOf	b
University	v_2	undergradDegreeFrom	c
Department	v_3	memberOf	d
student1	v_4	subOrganizationOf	e
univ1	v_5	telephone	f
dept1.univ1	v_6	emailAddress	g
'012-345-6789'	v_7		
'john@dept1.univ1.edu'	v_8		

(a) vertex mapping table.

(b) edge label mapping table.




(c) graph.

Figure 4: Direct transformation of RDF graph (Vertex label function $L(v) = \{v\}$).

A query graph is obtained from a SPARQL query. A query vertex may hold the vertex label which corresponds to the subject or object specified in the SPARQL query. If the query vertex corresponds to a variable, the vertex label is left blank. For example, the SPARQL query of Figure 5a is transformed into the query graph of Figure 5b. Here the query vertex u_0 , which corresponds to *Student*, holds the vertex label $\{v_1\}$; To the contrary, the query vertex u_3 , which corresponds to the variable X , has blank ($_$) as the vertex label. Similarly, a query edge may hold the edge label which corresponds to the predicate. For example, the edge label of (u_3, u_4) is c as the edge corresponds to the *undergradDegreeFrom* predicate.

```
SELECT ?X, ?Y, ?Z WHERE
{?X rdf:type Student .
 ?Y rdf:type University .
 ?Z rdf:type Department .
 ?X undergradDegreeFrom ?Y .
 ?X memberOf ?Z .
 ?Z subOrganizationOf ?Y.}
```

(a) SPARQL query.



(b) query graph.

Figure 5: Direct transformation of SPARQL query.

Note that, when a variable is declared on a predicate in a SPARQL query, a query edge has a blank edge label. An e-graph homomorphism algorithm can answer such SPARQL queries since an

e-graph homomorphism has edge label mapping from query edges to their corresponding edge labels.

Consequently, the direct transformation makes it possible to apply conventional e-graph homomorphism algorithms for processing SPARQL queries. In order to evaluate the performance of such an approach, we applied Turbo_{HOM} on LUBM8000, a billion-triple RDF dataset of Lehigh University Benchmark (LUBM) [8], after applying direct transformation. We compared the performance of Turbo_{HOM} against two existing RDF engines: RDF-3X [19], and System-X². Figure 6 depicts the measured execution time of these three systems in log scale. (See Section 6.1 for the details of the experiment setup)




Figure 6: Comparison between original Turbo_{HOM} with the direct transformation graph and other RDF engines.

Although there is no clear winner among them, the figure reveals that Turbo_{HOM} performs as good as the existing RDF engines. For short-running queries (i.e. Q1, Q3-Q5, Q7, Q8, Q10-Q13), Turbo_{HOM} shows faster elapsed time. As those queries specify a data vertex ID, Turbo_{HOM} only needs a small amount of graph exploration from one candidate region with an optimal matching order, while RDF-3X and System-X require expensive join operations. For long-running queries (i.e., Q2, Q6, Q9, and Q14), Turbo_{HOM} is slower than some of its competitors. The performance of Turbo_{HOM} largely relies on 1) graph exploration by *ExploreCandidateRegion* and 2) subgraph enumeration by *SubgraphSearch*. Moreover, when a query graph has non-tree edges, *IsJoinable* constitutes a large portion of *SubgraphSearch*. The profiling results of long running queries confirmed that 1) *ExploreCandidateRegion* and *SubgraphSearch* are the dominating factors and 2) for queries which have non-tree edges (Q2 and Q9), *IsJoinable* is the dominating factor of *SubgraphSearch*. Specifically, Turbo_{HOM} spent the most time on *ExploreCandidateRegion* (e.g. 46% for Q2, 70% for Q6, 72% for Q9, and 69% for Q14) and *SubgraphSearch* (e.g. 54% for Q2, 30% for Q6, 28% for Q9, and 31% for Q14). Moreover, for queries which have non-tree edges, the most of *SubgraphSearch* time was spent on *IsJoinable* (e.g. 81.4% for Q2 and 77.6% for Q9). In order to speed up *ExploreCandidateRegion*, we propose a novel transformation (Section 4.1). Tailored optimization techniques are proposed for improving performance for both functions (Section 4.3).

4. TURBOHOM++

In this section, we propose an improved e-graph homomorphism algorithm, Turbo_{HOM++}. Introduced first is the *type-aware transformation*, which can result in faster pattern matching than direct transformation (Section 4.1). Turbo_{HOM++} processes the labeled graph transformed by the type-aware transformation (Section 4.2). Furthermore, for efficient RDF query processing, four optimizations are applied to Turbo_{HOM++} (Section 4.3).

²We anonymize the product name to avoid any conflict of interest.

4.1 Type-aware Transformation

To enable the type-aware transformation, we devise the *two-attribute vertex model* which makes use of the type information specified by the `rdf:type` predicate. Specifically, this model assumes that each vertex is associated with a set of labels (the label attribute) in addition to its ID (the ID attribute). The label attribute is obtained by following the `rdf:type` predicate – if a subject has one or more `rdf:type` predicates, its types can be obtained by following the `rdf:type` (as well as `rdf:subClassOf` predicates transitively). For example, `student1` in Figure 3 has the label attribute, $\{\text{GradStudent}, \text{Student}\}$.

The above two-attribute vertex model naturally leads to our new RDF graph transformation, the *type-aware transformation*. Here, subjects and objects are transformed to two-attribute vertices by utilizing `rdf:type` predicates as described above. Then, the ID attribute corresponds to the vertex ID, and the label attribute corresponds to the vertex label. Figure 7 shows an example of the mapping tables and the data graph, which is the result of type-aware transformation applied to Figure 3. Now, we formally define the type-aware transformation as follows.

Definition 3. The type-aware transformation $(F_V, F_{ID}, F_E, F_{VL}, F_{EL})$ converts a set of triples $T(S, P, O)$ to a type-aware transformed graph $G(V, E, ID, L)$. Let us divide T into three disjoint subsets whose union is $T = T' \cup T_{sc}' \cup T_{sc}''$, $T'_i(S'_i, P'_i, O'_i) = \{(s, \text{rdf:type}, o) \in T\}$, and $T_{sc}'(S'_{sc}, P'_{sc}, O'_{sc}) = \{(s, \text{rdf:subClassOf}, o) \in T\}$.

1. A vertex mapping $F_V : S' \cup O' \cup S'_i \rightarrow V$, which is bijective, maps a subject in $S' \cup S'_i$ or an object in O' to a vertex.
2. A vertex ID mapping $F_{ID} : S' \cup O' \cup S'_i \rightarrow N \cup \{_ \}$, which is bijective, maps a subject in $S' \cup S'_i$ or an object in O' to a vertex ID or blank. Here, $F_{ID}(x) = _$ if x is a variable.
3. An edge mapping $F_E : T' \rightarrow E$, which is bijective, maps a triple of T' into an edge, $F_E(s, p, o) = (F_V(s), F_V(o))$.
4. A vertex label mapping $F_{VL} : O'_i \cup O'_{sc} \rightarrow VL \cup \{_ \}$, which is bijective, maps an object of $O'_i \cup O'_{sc}$ into a vertex label. Here, $F_{VL}(x) = _$ if x is a variable.
5. An edge label mapping $F_{EL} : P' \rightarrow EL \cup \{_ \}$, which is bijective, maps a predicate of P' into an edge label. Here, $F_{EL}(x) = _$ if x is a variable.
6. A vertex ID mapping function $ID : V \rightarrow N$ maps a vertex to a vertex ID where $ID(v) = F_{ID} \circ F_V^{-1}(v)$.
7. A labeling function L 1) maps a vertex to a set of vertex labels such that $v \in V, L(v) = \{F_{VL}(o) \mid \text{there is a path from } F_V^{-1}(v) \text{ to } o \text{ using triples in } T'_i \cup T_{sc}'\}$ and 2) maps an edge e to an edge label such that $e \in E, L(e) = F_{EL}(Pred(F_E^{-1}(e)))$ where $Pred(s, p, o) = p$.

After finding a type-aware transformation $(F'_V, F'_{ID}, F'_E, F'_{VL}, kF'_{EL})$ for a data graph $g(V', E', L', ID')$, we can also convert a SPARQL query into a type-aware transformed query graph $q(V, E, L, ID)$ by using another type-aware transformation $(F_V, F_{ID}, F_E, F_{VL}, F_{EL})$ such that $F_{ID} = F'_{ID}, F_{VL} = F'_{VL}$, and $F_{EL} = F'_{EL}$. For example, Figure 8 is the query graph type-aware transformed from the SPARQL query in Figure 5a. Note that a query vertex may have multiple vertex labels like a data vertex.

Now, we explain how the generic e-graph homomorphism algorithm works for type-aware transformed query/data graphs. When appending a candidate data vertex to the current partial solution, we additionally check the following condition for the ID attribute of the two-attribute vertex model.

Subject/Object	Vertex ID
student1	0
univ1	1
dept1.univ1	2
'012-345-678'	3
'john@dept1.univ1.edu'	4


(a) vertex ID mapping table.

Type	Vertex Label
GraduateStudent	A
Student	B
University	C
Department	D

(b) vertex label mapping table.

Predicate	Edge Label
undergradDegreeFrom	a
memberOf	b
subOrganizationOf	c
telephone	d
emailAddress	e

(c) edge label mapping table.



(d) data graph.

Figure 7: Type-aware transformation of an RDF graph.




Figure 8: Type-aware transformation of SPARQL query of Figure 5a.

$$\forall u \in \{u \mid ID(u) \neq _ \text{ for } u \in V\}, ID(u) = ID'(M_v(u)).$$

The virtue of the type-aware transformation is that it can improve the efficiency of RDF query processing. Since the type-aware transformation eliminates certain vertices and edges by embedding *type* information into the vertex label, the resulting data/query graphs have smaller size and simpler topology than those transformed by the direct transformation.

As an example, let us consider the SPARQL query in Figure 5a. After direct transformation, it becomes the query graph in Figure 5b that has a relatively complex topology consisting of six vertices and six edges. On the other hand, the type-aware transformation produces the query graph in Figure 8 that has a simple triangle topology. This reduced number of vertices and edges has a positive effect on efficiency because it results in less graph exploration.

In general, the effect of the type-aware transformation can be described in terms of the number of data vertices in all candidate regions. Consider a SPARQL query which consists of a set of triples T , its direct transformed query graph $q(V, E, L)$, and its type-aware transformed query graph $q'(V', E', ID', L')$. Let $O_{type} = \{o \mid (s, \text{rdf:type}, o) \in T \text{ or } (s, \text{rdf:subClassOf}, o) \in T\}$. In the direct transformation, $o \in O_{type}$ is transformed to a query vertex. Let V_{type} a set of direct transformed query vertices from O_{type} . However, in the type-aware transformation, $o \in O_{type}$ is not transformed to a query vertex, which satisfies $|V'| = |V| - |V_{type}|$. Therefore, the type-aware transformation leads to less graph exploration in *ExploreCandidateRegion* and *SubgraphSearch*. Formally, using the type-aware transformation, the number of data vertices in all candidate regions is reduced by

$$\sum_{v_s} \sum_{u \in V_{type}} |CR_{v_s}(u)|$$

where v_s represents the starting data vertex for each candidate region, and $CR_{v_s}(u)$ represents a set of data vertices in a candidate region $CR(v_s)$ that correspond to u .

4.2 Implementation

Turbo_{HOM++} maintains two in-memory data structures – the inverse vertex label list and the adjacency list. Figure 9a shows the inverse vertex label list of Figure 7d. The ‘end offsets’ records the exclusive end offset of the ‘vertex IDs’ for each vertex label. Figure 9b shows the adjacency list of Figure 7d for the outgoing edges. The adjacency list stores the adjacent vertices for each data vertex in the same way as the inverse vertex label list. One difference is that the adjacency list has an additional array (‘end offsets’) to group the adjacent vertices of a data vertex for each neighbor type. Here, the neighbor type refers to the pair of the edge label and the vertex label. For example, v_0 in Figure 7d, has four different neighbor types – (a, C) , (b, D) , $(d, _)$ and $(e, _)$. Those four neighbor types are stored in ‘end offsets,’ and each entry points to the exclusive end offset of the ‘adjacent vertex ID’. **Turbo_{HOM++}** maintains another adjacency list for the incoming edges.

We assume that graphs in our system are periodically updated from an underlying RDF source. For efficient graph update, a transactional graph store is definitely required. We leave this exploration to future work since it is beyond the scope of the paper.

Note also that **Turbo_{HOM++}** can also handle SPARQL queries under the simple entailment regime correctly. In order to deal with the simple entailment regime in the type-aware transformed graph, **Turbo_{HOM++}** distinguishes $L_{simple}(v) = \{F_{L_V}(o) \mid \text{there is an edge from } F_V^{-1}(v) \text{ to } o \text{ using triples in } T'\}$ from $L(v)$. **Turbo_{HOM++}** can process a SPARQL query under the simple entailment regime using $L_{simple}(v)$ instead of $L(v)$.

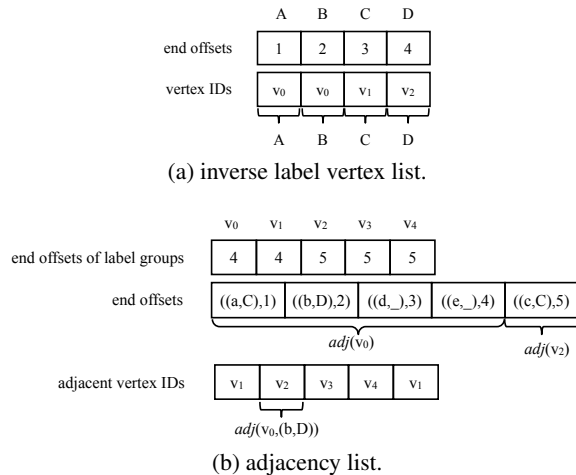


Figure 9: In-memory data structures for type-aware transformed data graph of Figure 7d ($adj(v)$: adjacent vertices of v , $adj(v, (el, vl))$: adjacent vertices v , which have vertex label vl and are connected with edge label el).

As the overall behavior of **Turbo_{HOM++}** is similar to **Turbo_{HOM}**, here, we describe how **Turbo_{HOM++}** uses the data structures in *ChooseStartQueryVertex* (line 6 of Algorithm 1), *ExploreCandidateRegion* (line 9 of Algorithm 1), and *IsJoinable* (line 7 of Algorithm 2).

ChooseStartQueryVertex. When computing $rank(u)$ for a query vertex u , the inverse vertex list is used to get $freq(g, L(u))$ ($= |\bigcap_{l \in L(u)} V(g)_l|$) where $V(g)_l$ is the set of vertices having vertex label l . When $|L(u)| = 1$, Getting the start and end offset of a specific vertex label is enough. When $|L(u)| > 1$, for each $l \in |L(u)|$, all data vertices having l , $V(g)_l$, are retrieved from the inverse vertex list, and $freq(g, L(u))$ is obtained by intersecting

all $V(g)_l$. Additionally, when a data vertex ID v is specified in u , $freq(g, L(u)) = 1$ if $v \in V(g)_l$ for each $l \in L(u)$. Otherwise, $freq(g, L(u)) = 0$.

One last case is when a SPARQL query has a query vertex which has no label or ID at all. In order to handle such queries, we maintain an index called the predicate index where a key is a predicate, and a value is a pair of a list of subject IDs and a list of object IDs. This index is used to compute $freq(g, L(u))$.

ExploreCandidateRegion. After a query tree is generated, candidate regions are collected by exploring the data graph in an inductive way. In the base case, all data vertices that correspond to the start query vertex are gathered in the same way of computing $freq(g, L(u))$. In the inductive case, once the starting data vertices are identified, the candidate region exploration continues by exploiting the adjacency information stored in the adjacency list. If one vertex label and one edge label are specified in the query graph, we can get the adjacent data vertices directly from the adjacency list. If multiple vertex labels and one edge label are specified, we collect the adjacent data vertices for each vertex label using the adjacency list, and intersect them. In a case where the vertex label or edge label is blank, **Turbo_{HOM++}** finds the correct adjacent data vertices by 1) collecting all adjacent vertices which match available information (either vertex label or edge label) and 2) unioning them. Additionally, if the current query vertex has the data vertex ID attribute, we check whether the specified data vertex is included in the data vertices collected from the adjacency list.

IsJoinable. The *IsJoinable* test is equivalent to the inductive case of *ExploreCandidateRegion* when a data vertex ID (previously matched data vertex) is specified.

4.3 Optimization

In this subsection, we introduce optimizations that we apply to improve the efficiency of **Turbo_{HOM++}**. Even though these optimizations do not change **Turbo_{HOM++}** severely, they could improve the query processing efficiency quite significantly.

Use intersection on *IsJoinable* test (+INT). We optimize the *IsJoinable* test in *SubgraphSearch*. *SubgraphSearch* calls the *IsJoinable* test by multiple membership operations. However, the optimization allows a bulk of *IsJoinable* tests with one k -way intersection operation where k is the number of edges between the current query vertex, u in line 1 of Algorithm 2, and the previously matched query vertices connected by non-tree edges.

SubgraphSearch checks the existence of the edges between the current candidate data vertex and the already bounded data vertices by calling *IsJoinable* (line 7 of Algorithm 2) when the corresponding query graph has non-tree edges. Let us consider the query graph (Figure 8), the query tree (Figure 10) and a data graph (Figure 11). Suppose that, for a given matching order $u_1 \rightarrow u_2 \rightarrow u_0$, the vertex v_1 is bound to u_1 , and the vertex v_2 is bound to u_2 . Then, the next step is to bind a data vertex to u_0 . Because there is a non-tree edge between u_0 and u_2 , to bind a data vertex of ID v_i ($i = 0, 3, 4, \dots, 1001$) to u_0 , we need to check whether there exists an edge $v_i \rightarrow v_2$.




Figure 10: A query tree of the query graph of Figure 8.




Figure 11: An example data graph for illustrating +INT.

IsJoinable checks for the existence of the edge between the current data vertex and already matched data vertices by repeatedly calling *IsJoinable*. Let us consider the above example. For each $v_i (i = 0, 3, 4, \dots, 1001)$, *IsJoinable* tests whether the edge $v_i \rightarrow v_2$ exists. If v_2 is a member of v_i 's outgoing adjacency list, the test succeeds, and the graph matching continues.

Instead, our modified *IsJoinable* tests all the edge occurrences between the current candidate vertices (C_R in line 3 of Algorithm 2) and the adjacency lists of the already matched data vertices by one k -way intersection operation. Let us consider the above example again. The modified *IsJoinable* finds the edge between v_2 and the candidate data vertices $v_0, v_3, \dots, v_{1001}$ at once. For this, it is enough to perform one intersection operation between the v_2 's incoming adjacency vertices and the candidate data vertices. Since the modified *IsJoinable* takes C_R as a parameter, the lines 3 and 7 of Algorithm 2 are merged into one statement.

Note that this optimization can improve the performance significantly. In the above example, since only v_0 and v_{1001} pass the test, we can avoid calling the original *IsJoinable* 998 times. Formally speaking, let us denote 1) the candidate data vertex set for the current query vertex u as C_R , 2) the previously matched query vertex set, which is connected to the current query vertex by non-tree query edges, as $\{u'_i\}_{i=1}^k$ and 3) the adjacent vertex set of $v'_i (= M_v(u'_i))$ where u'_i is connected to u with the vertex label vl_i and the edge label el_i , as $adj(v'_i, vl_i, el_i)$. Suppose that C_R and $adj(v'_i, vl_i, el_i)$ are stored in ordered arrays. Then, the complexity of the original *IsJoinable* test is

$$C_{original} = O(|C_R| \cdot \sum_{i=1}^k \log |adj(v'_i, vl_i, el_i)|)$$

, since *IsJoinable* is called for each $v \in C_R$, and $O(\log |adj(v'_i, vl_i, el_i)|)$ time is required to conduct a binary search for $|adj(v'_i, vl_i, el_i)|$ elements. On the contrary, the complexity of the modified *IsJoinable* test is

$$\min(O(|C_R| + \sum_{i=1}^k |adj(v'_i, vl_i, el_i)|), C_{original})$$

since the modified *IsJoinable* can choose the k -way intersections strategy between scanning $(k + 1)$ sorted lists and performing binary searches.

Disable NLF Filter (-NLF). The second optimization is to disable the NLF filter in *ExploreCandidateRegion*. The NLF filter may be effective when the neighbor type are very irregular. However, in practice, most RDF datasets are structured [7, 17]. For example, in our sample RDF dataset (Figure 3), in most case, a vertex corresponding to a graduate student has *telephone*, *emailAddress*, *memberOf*, and *undergraduateDegreeFrom* predicates. Accordingly, the NLF filter is not helpful for such structured RDF datasets.

Disable Degree Filter (-DEG). The third optimization is to disable the degree filter in *ExploreCandidateRegion*. Similar to

the NLF filter, the degree filter is effective when the degree is very irregular while RDF datasets typically are not.

Reuse Matching Order (+REUSE). The last optimization is to reuse the matching order of the first candidate region for all the other candidate regions. That is, *DetermineMatchingOrder* (line 6 of Algorithm 1) is called only once throughout the $TURBO_{ISO}$ execution, and the same matching order is used throughout the query processing. $TURBO_{HOM++}$ uses a different matching order for each candidate region, because each candidate region could have a very different number of candidate vertices for a given query path in the e-graph homomorphism problems. However, typical RDF datasets are regular at the schema level, i.e. well structured in practice, and generating the matching order for each candidate region is ineffective, especially when the size of each candidate region is small. We also performed experiments with more heterogeneous datasets, including Yet Another Great Ontology (YAGO) [23], and Billion Triples Challenge 2012 (BTC2012) [10]. This optimization technique still shows good matching performance as we will see in our extensive experiments in Section 6, since these heterogeneous datasets do not show extreme irregularity at the schema level.

5. RELATED WORK

With the increasing popularity of RDF, the demand for SPARQL support in relational databases is also growing. To meet such demand, most open-source and commercial relational databases support the RDF store and the RDF query processing. RDF datasets are stored into relational tables with a set of indexes. After that, SPARQL queries are processed by translating them into the equivalent join queries or by using special APIs.

To support RDF query processing, many specialized stores for RDF data were proposed [2, 4, 18, 19, 26, 29]. Similar to RDBMS, RDF-3X [18, 19] treats RDF triples as a big three-attribute table, but boosts the RDF query processing by building exhaustive indexes and maintaining statistics. RDF-3X processes many SPARQL queries by using merge based join, which is efficient for disk-based and in-memory environments. Different from RDF-3X, Jena [26] exploits multiple-property tables, while BitMat [2] exploits 3-dimensional bit cube, so that it can also support 2D matrices of SO, PO, and PS. H-RDF-3X [12] is a distributed RDF processing engine where RDF-3X is installed in each cluster node.

Several graph stores support RDF data in their native graph storages [30, 35]. gStore [35] performs graph pattern matching using the filter-and-refinement strategy. It first finds promising subgraphs using the VS^* -tree index. After that, the exact subgraphs are enumerated in the refinement step. Trinity.RDF [30] is a sub-system of a distributed graph processing engine, Trinity [22]. The RDF triples are stored in Trinity's key-value store. When processing RDF queries, Trinity.RDF implements special query processing methods for RDF data.

In 1976, Ullmann [24] published his seminal paper on the subgraph isomorphism solution based on backtracking. After his work, many subgraph isomorphism methods were proposed to improve the efficiency by devising their own matching order selection algorithms and filtering constraints [9, 11, 20, 21, 32, 33]. Among those improved methods, $TURBO_{ISO}$ [9] solves the notorious matching order problem by generating the matching order for each candidate region and by grouping the query vertices which have the same neighbor information. The method shows the most efficient performance among all representative methods.

Along with the backtracking based methods, the index-based subgraph isomorphism methods were also proposed [5, 27, 28, 31, 34]. All of those methods first prune out unpromising data graphs using low-cost filters based on the graph indexes. After filtering,

any subgraph isomorphism methods can be applied to those unfiltered data graphs. This technique is only useful when there are many small data graphs. Thus, these index-based subgraph isomorphism methods do not enhance RDF graph processing since there is only one big graph in an RDF database.

6. EXPERIMENTS

We perform extensive experiments on large-scale real and synthetic datasets in order to show the superiority of a tamed subgraph isomorphism algorithm for RDF query processing. In the experiment, we use Turbo_{HOM++}. We assume that Turbo_{HOM} uses direct transformation, while Turbo_{HOM++} uses type-aware transformation along with all optimizations. The specific goals of the experiments are 1) We show the superior performance of Turbo_{HOM++} over the state-of-the-art RDF engines (Section 6.2), 2) We analyze the effect of the type-aware transformation and the series of optimizations (Section 6.3), and 3) We show the linear speed-up of the parallel Turbo_{HOM++} with an increasing number of threads (Due to the space limit, please refer [13] for the detailed result).

6.1 Experiment Setup

Competitors. We choose three representative RDF engines as competitors of Turbo_{HOM++} – RDF-3X, TripleBit, and System-X. Note that these three systems are publicly available. RDF-3X [19] is a well-known RDF store, showing good performance for various types of SPARQL queries. TripleBit [29] is a very recent RDF engine efficiently handling large-scale RDF data. System-X is a popular RDF engine exploiting bitmap indexing. We exclude BitMat [2] from performance evaluation since it is clearly inferior to TripleBit [29]. gStore is excluded since it is not publicly available.

Datasets. We use four RDF datasets in the experiment – LUBM [8], YAGO [23], BTC2012 [10], and BSBM [3]. LUBM is a de-facto standard RDF benchmark which provides a synthetic data generator. Using the generator, we create three datasets – LUBM80, LUBM800, and LUBM8000 where the number represents the scaling factor. YAGO is a real dataset which consists of facts from Wikipedia and the WordNet. BTC2012 is a real dataset crawled from multiple RDF web resources. Lastly, BSBM is an RDF benchmark which provides a synthetic data generator and benchmark queries. BSBM uses more general SPARQL query features such as FILTER, OPTIONAL, and UNION. Due to the space limit, please refer [13] for the experimental results for YAGO since the performance trends of YAGO are similar to those for BTC2012.

In order to support the original benchmark queries in LUBM, we load the original triples as well as inferred triples into databases. In order to obtain inferred triples, we use the state-of-the-art RDF inference engine. For example, LUBM8000 contains 1068394687 original triples and 869030729 inferred triples. Note that this is the *standard* way to perform the LUBM benchmark. However, regarding BTC2012, we use the original triples only for database loading. This is because the BTC2012 dataset contains many triples that violate the RDF standard, and thus the RDF inference engine refuses to load and execute inference for the BTC2012 dataset. BSBM contains 986410726 original triples and 11412064 inferred triples.

Table 1 shows the number of vertices and edges of the graphs transformed by the direct transformation and the type-aware transformation. The reduced number of edges in the type-aware transformed graph directly affects the amount of graph exploration in e-graph homomorphism matching.

Queries. Regarding LUBM, we use the 14 original benchmark queries provided in the website³. Previous work such as [29] and

³<http://swat.cse.lehigh.edu/projects/lubm/>

Table 1: Graph size statistics (direct: direct transformation, type-aware: type-aware transformation).

	V direct	E direct	V type-aware	E type-aware
LUBM80	2644579	19461754	2644573	12357312
LUBM800	26304872	193691328	26304863	122994224
LUBM8000	263133301	1937425416	263133295	1230263406
BTC2012	367728453	1436545556	367459811	1185887764
BSBM	223938701	997822791	1937425416	893575906

[30] modified some of the original queries because executing those original queries without the inferred triples returns an empty result set. Regarding BTC2012, we use the same query sets proposed in [29], because they do not have official benchmark queries. Regarding BSBM, we used 12 queries in the explore use case⁴ which contain OPTIONAL, FILTER, and UNION keywords which test the capability of more general SPARQL query support.

In order to measure the pure subgraph matching performance, (1) we omit modifiers which reorganize the subgraph pattern matching results (e.g. DISTINCT and ORDER BY) in all queries and (2) we measure the elapsed time excluding the dictionary look-up time.

Running Environment. We conduct the experiments in a server running Linux four Intel Xeon E5-4640 CPUs and 1.5TB RAM. The server has the NUMA [15, 16] architecture with 4 sockets in which each socket has its own CPU and local memory.

We measure the elapsed times with a warm cache. To do that, we set up the competitors’ running environment as follows. For RDF-3X and TripleBit, as done in [30], we put the database files in the *tmpfs* in-memory filesystem, which is a kind of RAM disk. For System-X, we set the memory buffer size to 400GB, which is sufficient for loading the entire database in memory. We execute every query five times, exclude the best and worst times, and compute the average of the remaining three.

6.2 Comparison between Turbo_{HOM++} and RDF engines

We report the elapsed times of the benchmark queries using a single thread. Since the server has a NUMA architecture, memory allocation is always done within one CPU’s local memory.

LUBM. Table 2 shows the number of solutions for all benchmark queries in all LUBM datasets. Table 3 shows experimental results for LUBM80, LUBM800, and LUBM8000. Note that Triplebit was not able to return correct answers for two queries over LUBM80/LUBM800 and for ten queries over LUBM8000. In Table 3, we use ‘X’ or the superscript ‘*’ over the elapsed times when TripleBit returns incorrect numbers of solutions.

In order to analyze results in depth, we classify the LUBM queries into two types. The first type of queries has a constant number of solutions regardless of the dataset size. Q1, Q3 ~ Q5, Q7, Q8, and Q10 ~ Q12 belong to this type. These queries are called *constant solution queries*. The other queries (Q2, Q6, Q9, Q13, and Q14) have increasing numbers of solutions proportional to the dataset size. These queries are called *increasing solution queries*.

Regarding the constant solution queries, only Turbo_{HOM++} achieves the ideal performance in LUBM, which means constant performance regardless of dataset size. This phenomenon is analyzed as follows. Each constant solution query contains a query vertex whose ID attribute is set to an entity in the RDF graph. Thus, Turbo_{HOM++} chooses that query vertex as a starting query vertex and generates a candidate region. Furthermore, in the LUBM

⁴<http://wifo5-03.informatik.uni-mannheim.de/bizer/berlinsparqlbenchmark/spec/ExploreUseCase/index.html>

Table 2: Number of solutions in LUBM queries.

Dataset	Q1	Q2	Q3	Q4	Q5	Q6	Q7	Q8	Q9	Q10	Q11	Q12	Q13	Q14
LUBM80	4	212	6	34	719	838892	67	7790	21872	4	224	15	380	636529
LUBM800	4	2003	6	34	719	8352839	67	7790	218261	4	224	15	3800	6336816
LUBM8000	4	2528	6	34	719	83557706	67	7790	2178420	4	224	15	37118	63400587

Table 3: Elapsed time in LUBM [unit: ms] (X: wrong number of solutions (# of solutions difference > 3), **: wrong number of solutions (# of solutions difference ≤ 3)).

	Q1	Q2	Q3	Q4	Q5	Q6	Q7	Q8	Q9	Q10	Q11	Q12	Q13	Q14
Turbo _{HOM++}	0.09	6.37	0.09	0.13	0.13	4.43	0.05	2.26	101.42	0.09	0.10	0.10	0.06	3.08
RDF-3X	3.09	188.90	4.09	12.37	14.74	375.04	91.06	58.32	770.32	3.19	2.35	3.52	15.08	262.41
TripleBit	2.56	86.09	12.82	5.26	18.92	165.93	24.76	48.22*	X	9.23	0.44	1.86	19.31	132.09
System-X	2.00	426.00	2.00	4.67	2.67	64.33	4.00	19.33	3512.00	2.00	2.33	4.67	5.67	47.00

(a) LUBM80.

	Q1	Q2	Q3	Q4	Q5	Q6	Q7	Q8	Q9	Q10	Q11	Q12	Q13	Q14
Turbo _{HOM++}	0.09	124.13	0.09	0.13	0.13	25.70	0.05	2.32	1239.46	0.09	0.10	0.09	0.12	19.72
RDF-3X	4.15	2473.01	5.17	16.50	25.02	5103.35	840.16	461.80	10033.57	3.83	7.48	7.09	100.16	3607.13
TripleBit	23.32	3548.58*	142.29	15.76	183.46	2309.57	187.39	181.20*	X	109.47	2.84	3.51	161.65	1818.52
System-X	2.67	4394.00	2.00	4.67	3.00	239.33	4.33	21.00	175040.33	2.00	2.33	4.00	29.00	186.33

(b) LUBM800.

	Q1	Q2	Q3	Q4	Q5	Q6	Q7	Q8	Q9	Q10	Q11	Q12	Q13	Q14
Turbo _{HOM++}	0.10	309.74	0.09	0.12	0.13	191.52	0.05	1.61	5238.79	0.09	0.11	0.10	0.83	149.53
RDF-3X	4.31	30492.93	4.87	19.53	94.89	65453.67	8476.19	4201.81	131053.33	4.15	23.27	12.83	630.91	48285.17
TripleBit	X	X	X	X	2348.87	18974.80	X	X	X	1251.25	X	X	X	14197.47
System-X	2.67	41449.33	2.67	5.00	3.00	1519.67	4.33	42.67	3123629.67	2.67	2.33	5.00	88.00	1155.00

(c) LUBM8000.

datasets, although we increase the scaling factor in order to increase the database size, the size of the candidate region explored by every constant solution query remains almost the same.

In contrast, the elapsed times of RDF-3X increase as the dataset size increases. This is because the data size to scan for merge join increases as the dataset size increases. Thus, the performance gap between Turbo_{HOM++} and RDF-3X increases as the dataset size increases. In LUBM80, Turbo_{HOM++} is 23.50 (Q11) ~ 1821.20 (Q7) times faster than RDF-3X. In LUBM800, Turbo_{HOM++} outperforms RDF-3X by 42.56 (Q10) ~ 16803.20 (Q7) times. In LUBM8000, Turbo_{HOM++} outperforms RDF-3X by 43.10(Q1) ~ 169523.80 (Q7) times. TripleBit shows a similar trend as RDF-3X. Accordingly, Turbo_{HOM++} is 4.40 (Q11 in LUBM80) ~ 18068.23 (Q5 in LUBM8000) times faster than TripleBit. System-X shows constant elapsed times for these queries, although it is consistently slower than Turbo_{HOM++} by up to 86.60 times.

For the increasing solution queries (Q2, Q6, Q9, Q13, and Q14), Turbo_{HOM++} also shows the best performance in all LUBM datasets. Overall, the elapsed times of Turbo_{HOM++} are proportional to the number of solutions for these queries. Specifically, after type-aware transformation, Q13 has one query vertex whose ID attribute is set to an entity in the data graph. Thus, the number of candidate regions is one, which is similar to the constant solution query. However, as the dataset size increases, the candidate region size also increases. The other queries (Q2, Q6, Q9, Q14) do not have any query vertex whose ID attribute is set to an entity in the data graph. As the dataset increases, the number of candidate regions for these queries increases, while each candidate region size does not change. All systems show the increasing elapsed time as the dataset size increases. RDF-3X shows 7.60 (Q9 in LUBM80) ~ 760.13 (Q13 in LUBM8000) times longer elapsed times than Turbo_{HOM++}. TripleBit shows 13.51 (Q2 in LUBM80) ~ 1347.08 (Q13 in LUBM800) times longer elapsed time than Turbo_{HOM++}

when considering the queries which have the right number of solutions. System-X shows 7.72 (Q14 in LUBM8000) ~ 596.25 (Q9 in LUBM8000) times longer elapsed time than Turbo_{HOM++}. For the constant solution query, System-X seems to be the best competitor of Turbo_{HOM++}. However, regarding the most time-consuming queries (Q2, Q9), System-X shows poor performance.

BTC2012. Table 4 shows the exact number of solutions and elapsed times in BTC2012. Even though BTC2012 contains over 1-billion triples, all the engines process all BTC2012 queries quite efficiently. This is because the shapes of query graphs are simple (tree-shaped). Furthermore, like LUBM, Q2, Q4, and Q5 in the BTC2012 query set contain one query vertex whose ID attribute is set to an entity in the RDF graph. Still, Turbo_{HOM++} outperforms RDF-3X, TripleBit, and System-X by up to 422.60, 28.57, and 266.18 times, respectively.

Table 4: Number of solutions and elapsed time [unit: ms] in BTC2012.

	Q1	Q2	Q3	Q4	Q5	Q6	Q7	Q8
# of sol.	4	4	1	4	13	1	664	5996
Turbo _{HOM++}	0.12	0.16	0.96	0.89	0.18	2.49	36.81	1.99
RDF3X	6.67	7.52	10.42	13.07	69.97	22.75	392.73	841.96
TripleBit	1.56	1.81*	0.98	6.94	5.20	3.52	133.64*	X
System-X	8.00	4.67	5.00	12.33	4.67	663.67	110.67	351.67

BSBM. Table 5 shows the exact number of solutions and elapsed times in BSBM. The open source RDF engines, RDF-3X and TripleBit, are excluded as they do not support OPTIONAL and FILTER. Like BTC2012, even though BSBM contains about 1-billion triples, Turbo_{HOM++} processes most BSBM queries less than 5ms except Q5 and Q6. That is because they have a small number of solutions and contain one query vertex whose ID attribute is set to an entity in the RDF graph. For those ten queries,

Turbo_{HOM++} outperforms System-X by 2.37 ~ 7284.47 times. Q5 and Q6 take longer than the other queries because they use expensive filters such as join conditions (Q5) and a regular expression (Q6) and filter out a large number of solutions after basic graph pattern matching is finished. Before evaluating FILTER, Q5 (Q6) has 178030 (2848000) solutions from the query graph pattern and only qualifies 6803 (43508) final solutions.

Table 5: Number of solutions and elapsed time [unit: ms] in BSBM.

	Q1	Q2	Q3	Q4	Q5	Q6
# of sol.	79	17	202	142	6803	43508
Turbo _{HOM++}	0.58	0.15	8.15	1.27	344.66	3969.18
System-X	10	1092.67	19.33	21.67	589.67	9889.00
	Q7	Q8	Q9	Q10	Q11	Q12
# of sol.	2	1	21	3	10	1
Turbo _{HOM++}	0.25	0.16	0.11	0.23	0.14	0.12
System-X	23.33	12.33	4.00	11.00	3.00	8.00

6.3 Effect of Improvement Techniques

We measure the effect of the improvement techniques including the type-aware transformation (Section 4.1) and the four optimizations (Section 4.3). For this purpose, we use the largest LUBM dataset, LUBM8000. We first show the effect of the type-aware transformation because it is beneficial to all LUBM queries. We next show the effect of the four optimizations (Section 6.3.2).

6.3.1 Effect of Type-aware Transformation

Table 6 shows the elapsed times for the LUBM queries in LUBM8000 using the direct transformation (Turbo_{HOM}) and the type-aware transformation (Turbo_{HOM++} without optimizations). Compared with the direct transformation, the type-aware transformation improves the query performance by 1.01(Q1) to 27.22(Q6).

The obvious reason for performance improvement is the smaller query sizes after the type-aware transformation. The reduced sized query graph leads to smaller size candidate regions and shorter elapsed times. First of all, Q6 and Q14 benefit the most from the type-aware transformation. After the type-aware transformation, these queries become point-shaped. That is, solutions of these two queries are directly obtained by iterating the data vertices which have the vertex label of the query vertex, which corresponds to lines 2–4 in Algorithm 1. Q13 also benefits much from the type-aware transformation, since the type-aware transformation chooses a better starting query vertex than the direct transformation which chooses a query vertex having type information. Q1, Q3, Q4, Q5, Q7, Q8, Q10, Q11, and Q12 do not benefit from the type-aware transformation because they already have a small number of candidate vertices under the direct transformation.

Q2 benefits less than the other long running queries from the type-aware transformation. The following is the profiling result of Q2 with the direct/type-aware transformation. Q2 with direct transformation takes 26774.73 milliseconds in *ExploreCandidateRegion* and 31191.29 milliseconds in *SubgraphSearch*. Note that, with direct transformation, the starting vertex is arbitrarily chosen from u_0, u_1, u_2 in Figure 5b since they all have same vertex label frequency ($freq(g, L(u_i)) = 1, i = 0, 1, 2$) and the same degree of 1. In our implementation, the first query vertex u_0 is chosen and thus the label of the non-tree edge is *subOrganizationOf*. However, with type-aware transformation, the starting vertex is u_1 in Figure 8, and the label of the non-tree edge is *memberOf*. Although the number of

candidate regions with u_1 is the minimum among $u_0, u_1,$ and u_2 , the cost of *IsJoinable* calls for *memberOf* increases 1.30 times. Thus, Q2 with type-aware transformation takes 9523.60 milliseconds in *ExploreCandidateRegion* and 40469.47 milliseconds in *SubgraphSearch*. We achieve only 1.16 times performance improvement. However, the cost of the *IsJoinable* call is significantly reduced by using +INT. Thus, after applying type-aware transformation and the tailored optimizations, the final elapsed time for Q2 becomes 309.74ms, i.e., 187.15 times performance improvement compared with direct transformation only.

6.3.2 Effect of Four Optimizations

In this experiment, we measure the effect of four optimizations of Turbo_{HOM++}. We use Q2 and Q9 in LUBM8000 since these two queries in LUBM8000 are the most time-consuming and exploit all optimizations. All the other queries are omitted since their elapsed times are too short, so that it is hard to recognize the effect of optimization. Note that the elapsed times of Q1, Q3 ~ Q5, Q7, Q8, Q10 ~ Q13 are too short ($< 2ms$), and Q6 and Q14 do not benefit from these optimizations since they are point-shaped.

Figure 12 shows the reduced times of Q2 and Q9 in LUBM8000 after applying these optimizations separately. The optimization techniques in X-axis are ordered by the reduced in a decreasing manner — +INT, -NLF, -DEG, and +REUSE. Interestingly, even though Q2 and Q9 have the same shape (i.e., triangular), the most effective optimizations were different. +INT was the most effective in Q2. -NLF was the most effective in Q9 since the size of each candidate region was very small. -DEG was more effective in Q9 than in Q2 since Q9 has more data vertices applied to the degree filter. +REUSE was effective in Q9 which has large number of candidate regions while Q2 did not benefit from +REUSE.




Figure 12: Reduced elapsed time of each optimization (Elapsed time of no-optimization: 50016.13ms (Q2) and 17829.50ms (Q9)).

7. CONCLUSION

The core function of processing RDF data is subgraph pattern matching. There have been two completely different directions for supporting efficient subgraph pattern matching. One direction is to develop specialized RDF query processing engines exploiting the properties of RDF data, while the other direction is to develop efficient subgraph isomorphism algorithms for general, labeled graphs. In this paper, we posed an important research question, “*Can subgraph isomorphism be tamed for efficient RDF processing?*” In order to address this question, we provided the first direct and comprehensive comparison of the state-of-the-art subgraph isomorphism method with representative RDF processing engines.

We first showed that a subgraph isomorphism algorithm requires minimal modification to handle a graph homomorphism with edge label mapping which is the RDF graph pattern matching semantics. We then provided a novel transformation method, called

Table 6: Effect of type-aware transformation in LUBM8000 (Performance gain = Direct transformation ÷ Type-aware transformation).

	Q1	Q2	Q3	Q4	Q5	Q6	Q7	Q8	Q9	Q10	Q11	Q12	Q13	Q14
Direct transformation (ms)	0.101	57966.93	0.11	0.16	0.43	5218.47	0.15	5.63	114116.33	0.10	0.21	0.30	21.48	3886.43
Type-aware transformation (ms)	0.100	50016.13	0.09	0.14	0.13	191.69	0.05	1.73	17829.50	0.09	0.11	0.10	1.33	149.60
Performance gain	1.01	1.16	1.23	1.09	3.34	27.22	2.80	3.25	6.40	1.14	1.95	3.01	16.17	25.98

type-aware transformation along with a series of optimization techniques. We next performed extensive experiments using RDF benchmarks in order to show the superiority of the optimized subgraph isomorphism over representative RDF processing engines. Experimental results showed that the optimized subgraph isomorphism method achieved consistent and significant speedup over those RDF processing engines.

This study drew a promising conclusion that a subgraph isomorphism algorithm tamed for RDF processing can serve as an in-memory accelerator on top of a commercial RDF engine for *real-time* RDF query processing as well. We believe that this approach opens a new direction for RDF processing, so that both traditional directions can merge or benefit from each other.

Acknowledgment

This work was supported in part by a gift from Oracle Labs' External Research Office. This work was also supported by the National Research Foundation of Korea(NRF) grant funded by the Korea government(MSIP) (No. NRF-2014R1A2A2A01004454) and the MSIP(Ministry of Science, ICT and Future Planning), Korea, under the "ICT Consilience Creative Program" (IITP-2015-R0346-15-1007) supervised by the IITP(Institute for Information & communications Technology Promotion).

References

- [1] D. J. Abadi et al. Sw-store: A vertically partitioned dbms for semantic web data management. *The VLDB Journal*, 385–406, 2009.
- [2] M. Atre et al. Matrix "bit" loaded: A scalable lightweight join query processor for rdf data. In *WWW '10*, 41–50.
- [3] C. Bizer and A. Schultz. The berlin sparql benchmark. *International Journal on Semantic Web and Information Systems (IJSWIS)*, 1–24, 2009.
- [4] J. Broekstra et al. Sesame: A generic architecture for storing and querying rdf and rdf schema. In *ISWC '02*, 54–68.
- [5] J. Cheng et al. Fg-index: Towards verification-free query processing on graph databases. In *SIGMOD '07*, 857–872.
- [6] W. Fan et al. Graph homomorphism revisited for graph matching. *VLDB '10*, 1161–1172.
- [7] A. Gubichev and T. Neumann. Exploiting the query structure for efficient join ordering in SPARQL queries. In *EDBT '14*, 439–450.
- [8] Y. Guo et al. Lubm: A benchmark for owl knowledge base systems. *Web Semant.*, 158–182, 2005.
- [9] W.-S. Han et al. Turbo_{ISO}: towards ultrafast and robust subgraph isomorphism search in large graph databases. In *SIGMOD '13*, 337–348.
- [10] A. Harth. Billion Triples Challenge data set. Downloaded from <http://km.aifb.kit.edu/projects/btc-2012/>, 2012.
- [11] H. He and A. K. Singh. Graphs-at-a-time: Query language and access methods for graph databases. In *SIGMOD '08*, 405–418.
- [12] J. Huang, D. J. Abadi, and K. Ren. Scalable sparql querying of large rdf graphs. *VLDB '11*, 1123–1134.
- [13] J. Kim et al. Taming subgraph isomorphism for rdf query processing. Tr, 2015. URL <http://wshan.net/download/TR/TurboHOM.pdf>.
- [14] J. Lee et al. An in-depth comparison of subgraph isomorphism algorithms in graph databases. *VLDB '12*, 133–144.
- [15] V. Leis et al. Morsel-driven parallelism: A numa-aware query evaluation framework for the many-core age. In *SIGMOD '14*, 743–754.
- [16] Y. Li, I. Pandis, R. Müller, V. Raman, and G. M. Lohman. Numa-aware algorithms: the case of data shuffling. In *CIDR*, 2013.
- [17] T. Neumann and G. Moerkotte. Characteristic sets: Accurate cardinality estimation for rdf queries with multiple joins. In *ICDE '11*, 984–994.
- [18] T. Neumann and G. Weikum. x-rdf-3x: fast querying, high update rates, and consistency for rdf databases. *VLDB '10*, 256–263.
- [19] T. Neumann and G. Weikum. The rdf-3x engine for scalable management of rdf data. *The VLDB Journal*, 91–113, 2010.
- [20] L. P. Cordella et al. A (sub)graph isomorphism algorithm for matching large graphs. *IEEE Trans. Pattern Anal. Mach. Intell.*, 1367–1372, 2004.
- [21] H. Shang et al. Taming verification hardness: An efficient algorithm for testing subgraph isomorphism. *VLDB '08*, 364–375.
- [22] B. Shao et al. Trinity: A distributed graph engine on a memory cloud. In *SIGMOD '13*, 505–516.
- [23] F. M. Suchanek et al. Yago: A large ontology from wikipedia and wordnet. *Web Semant.*, 203–217, 2008.
- [24] J. R. Ullmann. An algorithm for subgraph isomorphism. *J. ACM*, 31–42, 1976.
- [25] C. Weiss et al. Hexastore: sextuple indexing for semantic web data management. *VLDB '08*, 1008–1019.
- [26] K. Wilkinson and K. Wilkinson. Jena property table implementation. In *SSWS '06*, 35–46.
- [27] X. Yan et al. Graph indexing: A frequent structure-based approach. In *SIGMOD '04*, 335–346.
- [28] X. Yan et al. Graph indexing based on discriminative frequent structure analysis. *ACM Trans. Database Syst.*, 960–993, 2005.
- [29] P. Yuan et al. Triplebit: a fast and compact system for large scale rdf data. *VLDB '13*, 517–528.
- [30] K. Zeng et al. A distributed graph engine for web scale rdf data. *VLDB '13*, 265–276.
- [31] S. Zhang et al. Treepi: A novel graph indexing method. In *ICDE '07*, 966–975, .
- [32] S. Zhang et al. Gaddi: Distance index based subgraph matching in biological networks. In *EDBT '09*, 192–203, .
- [33] P. Zhao and J. Han. On graph query optimization in large networks. *VLDB '10*, 340–351.
- [34] L. Zou et al. A novel spectral coding in a large graph database. In *EDBT '08*, 181–192, .
- [35] L. Zou et al. gstore: answering sparql queries via subgraph matching. *VLDB '11*, 482–493, .