

Regular Path Queries on Large Graphs

André Koschmieder and Ulf Leser

Humboldt-Universität zu Berlin, Germany
Department of Computer Science
{koschmie,leser}@informatik.hu-berlin.de

Abstract. The significance of regular path queries (RPQs) on graph-like data structures has grown steadily over the past decade. RPQs are, often in restricted forms, part of graph-oriented query languages such as XQuery/XPath and SPARQL, and have applications in areas such as semantic, social, and biomedical networks. However, existing systems for evaluating RPQs are restricted either in the type of the graph (e.g., only trees), the type of regular expressions (e.g., only single steps), and/or the size of the graphs they can handle. No method has yet been developed that would be capable of efficiently evaluating general RPQs on large graphs, i.e., with millions of nodes/edges.

We present a novel approach for answering RPQs on large graphs. Our method exploits the fact that not all labels in a graph are equally frequent. We devise an algorithm which decomposes an RPQ into a series of smaller RPQs using rare labels, i.e., elements of the query with few matches, as way-points. A search thereby is decomposed into a set of smaller search problems which are tackled in a bi-directional fashion, supported by a set of graph indexes. Comparison of our algorithm with two approaches following the traditional methods for tackling such problems, i.e., the usage of automata, reveals that (a) the automata-based methods are not able to handle large graphs due to the amount of memory they require, and that (b) our algorithm outperforms the automata-based approach, often by orders of magnitude. Another advantage of our algorithm is that it can be parallelized easily.

1 Introduction

A general regular path query (RPQ) is a regular expression R over the (edge or node) labels of a graph G [28]. Its result is the set of all cycle-free paths in G whose concatenation of labels (edge or node) spells out R . Different flavors of RPQs are used in a wide range of applications. For instance, XPath supports a restricted form of RPQs on XML documents [25]. SPARQL supports a very simple form of RPQs for RDF graphs, and various proposals exist for enhancing SPARQL syntax with full RPQs (e.g., [4,5,9,20]). [11] describes a restricted form of RPQs important for graph pattern matching, and [32,33] describe languages supporting restricted forms of RPQs for studying social networks. A particularly important application domain for RPQs are the Life Sciences, where understanding the interactions of different biological entities is of great importance. Such

interactions are typically modeled as graphs, and RPQs are used to find specific biochemical pathways between distant nodes [23].

None of these works support general RPQs on large graphs, but all focus on restricted languages which typically allow for more efficient evaluation. Evaluating RPQs on arbitrary graphs is an NP-hard problem [28]. We illustrate the problem and our main idea by an example. Suppose a graph of researchers (nodes), either labeled as **P**rofessors or **S**Tudents, connected by directed edges such as **S**upervised or **J**oint work. In this graph, the query $P(JP)(JP)?$ finds all paths between a professor and direct or indirect co-workers. $(PS)(PS) + (P|T)$ finds all paths between a professor and his doctorate descendants. Now suppose we also model research prizes as nodes (such as **N**obel Prize or **S**igmoid **A**ward), and connect them to researchers with edges labeled **H**onored. Then, we can find the doctorate predecessors of all Nobel Prize winners using the query $(PS) + PHN$, and those of any prize winner using the query $(PS) + PH(N|A)$.

RPQs have been studied intensively for XML, where the predominant approach is to use automata [13]. Both the graph and the query are represented as automata, whose intersection automaton is the subgraph specified by the query. In this process, the graph needs to be translated into a DFA, which can be of exponential space and may need exponential construction time. Research in XML query languages has shown that automata-based RPQ evaluation works well for trees (XML) [30], but we will show that its space consumption is enormous on general graphs (see Section 6). Furthermore, automata-based approaches completely disregard the fact that certain labels are much more frequent than others, which can be exploited for speeding up query execution. For instance, to answer the query for Nobel Laureates, it is sensible to first search for nodes labeled with N , because (a) one such node must be in any matching path and (b) there are much less Nobel Laureates than professors. Having all N nodes, complete paths can be computed easily by traversing the graph.

Such reasoning is the basis of the approach we propose in this paper; however, finding a good evaluation strategy is not always as easy as in the example, as a query usually contains various labels with different frequencies in the graph. To this end, we first gather all labels in the query that only occur a few times in the graph and use these as fix points for a series of bi-directional searches. Thus, we split the query at these rare labels into smaller queries, and answer these individually. We search all matching paths between each two adjacent rare labels as well as at the start and end of the query, and combine the results to answer the original query. The main advantage of this approach is that we do not need to consider the whole graph, but only those fractions of it that lie between adjacent rare labels. In Section 6 we will show that this strategy, over a wide range of small to large synthetic and real-life graphs, is considerably faster and especially much less space-demanding than automata-based methods ([39,14]).

Figure 1 illustrates this idea. Suppose we want to answer the RPQ $a+ b c+$ on a graph (Fig. 1b). Since there is only one edge labeled with b , we use it as rare label and split the query there. Now, the two smaller queries $a+$ and $c+$ have to be answered, using the b edge as end point or start point, respectively.




Fig. 1. a) The RPQ $a+ b c+$ shown as a (nondeterministic) automaton; b) a path fulfilling the RPQ in a small exemplar graph

The result of the original query is the combination of the smaller queries and the rare label.

Our idea can also be used for variations of the general RPQ problem, such as finding all shortest matching paths spelling out a regular expression, or finding all matching paths between two given nodes. Especially the latter is important if RPQs are used as predicates in a general query language (see, for instance, [20]), where variables for defining the start and the end of a path often already have been bound by other predicates before RPQs are evaluated. For space constraints, in this paper we only describe in detail the algorithm for solving the general RPQ problem (finding all acyclic paths matching a given regular expression) without bindings for the start and end nodes, as this is the most complex case. Furthermore, we only consider RPQs over edge labels; extensions to include node labels are straight-forward. Adaptations of our method targeting such variations of the RPQ problem are described in [21], which also contains a number of additional experiments we omit here for space constraints.

This paper is structured as follows. Section 2 gives an overview on related work. In Section 3, we define the basic concepts. In Section 4, we describe our novel RPQ evaluation algorithm and give implementation details in Section 5. We evaluate our method in Section 6 using various real and synthetic graphs and conclude in Section 7.

2 Related Work

The most common approach for answering RPQs is based on automata. One implementation of this idea are DataGuides by Goldman and Widom [13], based on Lorel [1]. Therein, the graph is considered as an NFA that is first converted into a DFA and then minimized. This minimized DFA (a DataGuide) is then used as an index. However, this index can become much larger than the original graph, which is a problem when dealing with arbitrary graphs (but not for regularly structured XML). Goldman and Widom therefore propose “Approximate DataGuides” [14] which reduce the index size using heuristics. As the implementation of Lorel (and DataGuides) is not supported any more since 2000, we re-implemented the algorithms and compare them to our approach in Section 6.

After the uptake of the semi-structured data model in XML, many proposals have been put forward to use automata in optimizing queries on XML (see e.g. [25,30]). However, these works mostly use tree automata [30] and are not applicable to arbitrary graphs. Additional index structures have been proposed, for instance, by Milo and Suciu [29] and Kaushnik et al. [19], but, again, they are designed to work with XML data and cannot be applied to non-tree graphs. Fernandez and Suciu present another interesting approach to speed up graph searching based on Graph Schemas [12]; however, these have to be created manually, a step that seems unfeasible for graphs with millions of nodes.

An area where RPQ queries on graphs are important is querying RDF data. However, SPARQL, the official W3C recommendation as an RDF query language, does not support regular path queries, which spurred research into extending SPARQL with RPQs. Alkhateeb et al. [4] developed the query language PSPARQL that includes RPQs, but the authors focus on formal semantics of RPQs on RDF and do only describe a proof-of-concept implementation based on backtracking. Detwiler et al. [9] present the GLEEN system, an extension to SPARQL including RPQs that is implemented as an extension to the ARQ library for SPARQL processing (see Section 6.2 for a comparison). Anyanwu et al. [5] present another extension to SPARQL that also includes RPQs, but no method for evaluating them is described. Zauner et al. [39] present a path language for RDF supporting RPQs and does provide an implementation; the system is based on automata, and we compare against it in Section 6.2. SPARQLeR is another RDF querying language encompassing RPQs, again based on automata, for which an implementation was described in [20], but is not publicly available. Note that the runtimes reported in the paper range in the order of seconds for queries with bound start and end nodes on moderately sized graphs, a setting in which our algorithm only needs milliseconds (see [21]).

There were also a number of proposals for general graph query languages that are not based on RDF. Leser [23] proposes a query language for querying biological pathways, which syntactically supports RPQs, but does not describe a scalable resolution technique. Graphs-at-a-time is a query language based on graph grammars that is capable of expressing RPQs, but the presented implementation does not cover such predicates [15]. The query language proposed in [10], which also includes a type of RPQs, is not accompanied by any implementation. Several systems have been designed to support extremely large graphs (hundreds of millions nodes and edges), such as Pregel [26], GRAIL [38], or DEX [27], however, none of these systems support RPQs. Finally, Sevón and Eronen [34] describe a method for querying paths in labeled graphs using context-free grammars. They traverse the graph breadth-first and use a context-free parser to find matches. While context-free grammars are more powerful than regular expressions, [34] only focus on finding paths between fix start and end nodes and do not provide optimization techniques as we do.

Another line of related work are concerned with graph pattern matching. Fan et al. [11] show that a language supporting reachability and a restricted form of RPQs allows for an evaluation in cubic time. Jin et al. [18] present

algorithms for finding paths which only consist of labels from a predefined set of labels. Ronen and Shmueli [32] describe a graph query language that supports conditions on labels being contained or not contained in paths and also supports ranking (provided that the edges are weighted) and aggregation over sets of paths. In contrast to these works, our approach supports full RPQs.

Zou et al. introduced the Distance-join [40] as a technique to find subgraphs with similar connections to a query pattern. The work most related to ours is probably [7], in which Cheng et al. present algorithms for finding subgraphs that are homeomorphic to a query graph, i.e., where each edge of the query may be mapped to a path of unbounded length in the graph. In some sense, this can be seen as a generalization of the RPQ problem we study in this paper; however, a transformation would be non-trivial (mapping of edges to nodes and vice versa; introduction of special edges which must be mapped to exactly one edge for labels without '?' or '*' modifier etc.). Besides, the source code of [7] seems to be not available for comparisons.

Matching regular expressions (REs) on strings is a related problem that was recently picked up by the database community. Examples are [6,8], which both use index structures that are similar in spirit to our method, i.e., concentrating on rare characters in the query. Another line of related research is graph indexing. Here, the idea of using frequencies of labels has been used extensively, especially in mining and searching of subgraphs [36,22].

In summary, despite a large body of research around evaluating RPQs on graphs, we are aware of only two available implementations supporting full RPQs as we do [39,9]. In Section 6, we compare our approach to these methods and also to a re-implementation of the DataGuide system [13] and show that, for large graphs, they suffer from excessive memory consumption and are clearly outperformed by our method.

3 Terms and Definitions

We use labeled directed multigraphs, i.e., a graph G is a tuple $G = (V, E, f, l, \Sigma)$, where V is a finite set of nodes, E is a finite set of edges, $l : E \rightarrow \Sigma$ specifies the edge labels Σ , and $f : E \rightarrow V \times V$ is the connection function, specifying which nodes are connected by which edges.

The topological properties of a graph can be measured through node degree and label distribution. A graph is called scale-free if the number of nodes with degree k is $P(k) \sim k^{-\lambda}$ for large values of k . Scale-free graphs are the likely outcome of various random growth processes, and, indeed, many graphs discovered in biological research are scale-free [24]. The label distribution in a graph is called Zipfian if the frequency of the labels occurring in the graph follows the power law $F(k) \sim k^{-\delta}$, with $\delta \approx 1$.

A regular path query (RPQ) is a regular expression over Σ . We use the definition for regular expressions as in [16]. To evaluate regular path queries, a regular expression can be converted into an automaton that can be used to match paths. We assume definitions for deterministic (DFA) and non-deterministic automata (NFA) as in [3].

Several kinds of questions can be answered for a given regular expression R and a given graph G . (1) Does G contain any path fulfilling R ? (2) Which is the shortest path in G fulfilling R ? (3) Is there a path in G between two fixed nodes fulfilling R ? (4) Which paths in G fulfill R ? In this paper, we discuss our proposal using the latter problem as show case as it subsumes all other types of queries. In Section 6, we will show that, for instance, fixing the start and end nodes of an RPQ allows to drastically speed-up query processing compared to the more liberal problem of returning all paths in the graph.

4 Answering RPQs Using Rare Labels

In this section, we present our novel algorithm for efficiently answering RPQs. The basic idea is to search the graph while simultaneously advancing in the query automaton. Compared to an approach which converts both the graph and the query into automata, our method has several advantages: No preprocessing of the large graph is needed, it only uses space linear in the size of the graph, the search is easily parallelizable, and it can be enhanced with techniques that take label frequencies into account.

4.1 Rare Labels

Definition 1. Let G be a graph and R be a RPQ. We call a label occurrence in R *mandatory* iff it is not followed by a modifier other than $+$, i.e. it occurs in every possible result of R in G . We call a label *rare* iff it occurs at most m times in G and is mandatory in R (where m is a parameter of the method, see Section 4.3).

For example, in the regular expression $a b+ c* d?$, a and b are mandatory, while c and d are not. Finding all rare labels in a query is very fast if a list of all labels in the graph together with their frequencies is stored. If this list is indexed by labels, finding all rare labels for a given query is linear to the size of the query pattern.

If a query contains a rare label, then any match of the query in the graph must contain an occurrence of it. Therefore, we can use the occurrences of rare labels as way-points during the search process. If we can find two or more rare labels in a query, we can use a two-way search algorithm to find all matching paths between their matches in the graph. Every additional rare label further reduces the search space.

Note that our implementation (see Section 5.1) also treats disjunctions (regular expressions of the form $(a|b|...)$) as mandatory by searching for any occurrence of the labels. A further, not yet implemented optimization would be to rewrite expressions of the form $R S* T$, where R , S , and T are regular expressions, as two expressions RT and $R S+ T$. Thus, the original expression could be answered by evaluating two expressions, one of which is shorter and the other contains an additional mandatory label for our optimization.

4.2 Searching the Graph Using Rare Labels

For queries that include at least one rare label, we split the query at these rare labels and use them as fix points in the search. Searching the graph and advancing in the regular expression at the same time, we search all paths between each two adjacent rare labels, all paths from the first rare label backward to the start of the regular expression, and from the last rare label forward to the end. As shown above, the number of nodes that need to be visited during this search shrinks with every additional rare label but grows with increasing numbers of occurrences of rare labels.

Besides keeping the search space smaller, rare labels often also allow for early stops. If there is no path between any two adjacent rare labels, then there can be no path fulfilling the original query, and the search can be stopped immediately.

In the following, we use the term *first rare nodes* for all nodes that are starting point of an edge of which the label is the first rare label, according to the regular expression. Analogously, *last rare nodes* are the end nodes of all edges with the last rare label. Answering RPQs using rare labels is done in the following 6 steps.

1. Gather all rare labels for the query in the graph.
2. If more than one rare label exists, find the paths between the first and second rare label, the second and third etc. using a two-way search algorithm. If no path can be found in any of these search processes, stop the search and return an empty result for the query.
3. If more than one rare label exists: Using the results from step 2, find all paths from the *first rare nodes* to the *last rare nodes* and remove all rare nodes that are on no path, as these cannot be on a result path.
4. Beginning at all remaining *first rare nodes*, find all paths to the beginning of the regular expression, searching backward.
5. Beginning at all remaining *last rare nodes*, find all paths to the end of the regular expression (forward).
6. Using the results, enumerate all paths in the graph that fulfill the regular expression and return the result.

Figure 2 shows the principle of the algorithm. On a sample graph (edge labels and directions omitted), the query $a+ b c+ d e+$ is executed, assuming that b and




Fig. 2. Search process example for the query $a+ b c+ d e+$ in an arbitrary graph (edge labels and directions omitted)

d are rare labels. In step 1, rare label edges are gathered (b and d edges). In step 2, we search all paths between the end nodes of the b edges and the beginning of the d edges. These paths must fulfill the regular expression between the two rare labels, in this case $c+$. Here, two such paths can be found. For one rare edge, no path could be found, thus it is removed from further consideration in step 3.

In step 4, a one-way backward search is performed, starting at the start nodes of the b edges. The search ends once all paths have been found that fulfill the first part of the regular expression ($a+$). In step 5, we search all paths from the end of the last rare label to the end of the regular expression in forward direction. As a last step (not shown in the picture), we enumerate all paths by combining the results of the previous steps. In this case, there are 4 distinct paths. The result subgraph can be gathered by enumerating all nodes and edges of the result paths.

Our approach specifically aims at queries that include labels that do not occur often in the graph. While most queries used in Bioinformatics are interested in these rare labels, there are also queries in which no rare label is present. In such cases, our algorithm automatically switches to a brute force search, starting a search at every node in the graph (respectively at the given start/end nodes, if specified). In Section 6 we show that our implementation is faster than other approaches (in particular, the automata-based one) even for those cases.

4.3 Determining Rare Labels

The algorithm described above assumes a fixed value for m , the parameter determining which labels are considered rare. The choice of m needs to find a compromise between treating as many labels as rare as possible and keeping the number of occurrences of rare labels as small as possible. If a rare label has many occurrences in the graph, the search space increases because each occurrence needs to be included in the search (forward and backward). On the other hand, multiple different rare labels in a query speed up its execution, because partial paths that need to be searched are shorter.

We know of no simple way to determine the best value for m . It depends on the graph as well as on the query, so using a fixed value is not the best approach. We therefore use an adaptive heuristic for determining which labels to consider rare for a given graph and query. The idea is that if, for a given query, there are several possible rare labels, we set the threshold for rare labels higher than if only very few rare labels can be found. This, on the one hand, produces less queries without any rare labels. On the other hand, for queries with many potential rare labels, only labels with a small number of occurrences are included.

Our proposed heuristic works as follows. We first acquire a list of all potential rare labels for the query. We then reduce this list depending on the overall number of paths that would need to be searched in the current configuration. Labels that produce the most paths are removed first; we can compute the number of paths between any two adjacent rare labels r_1, r_2 as $|r_1| \cdot |r_2|$. The overall number of paths is the sum of all paths between all adjacent pairs. For the first and last rare labels, we also add their number of occurrences to account

for the search to the beginning and to the end of the path. We repeatedly remove the rare label that produces the most paths, until the sum of all paths is below a threshold.

5 Implementation

In this section, we give a short overview of the architecture of our system and the data structures used. We describe the implementation of our algorithms in some more detail and give a complexity analysis of our search algorithms.

We use a node based storage schema, which means that the nodes in the graph are represented explicitly, while the edges only exist as attributes of the nodes in the form of adjacency lists. Edges are stored in forward and backward direction, which enables two-way searching but almost doubles memory consumption. Labels are always stored as integers; if the labels are given as strings, a global mapping table is used to map them to numbers. To be able to efficiently gather rare labels from the graph, we also use an index on the edge labels also encoding their multiplicity. Figure 3 visualizes the architecture of our system.

Regular path queries are represented as NFAs. Converting a regular expression into an NFA is straight-forward. The automaton is stored as states and transitions, which are labeled with the number representing the label for the transition. Transitions are stored in both directions for two-way search. To speed up query execution, we create a list for every state with all labels that are accepted in that state, and a list showing into which states the automaton may transition for each label. We call these *labeled follow sets* as a reference to follow sets used in compiler construction [3].

Our current implementation requires about 2 GB of memory for a graph with 10 million nodes and 20 million edges, including all additional data structures described in the following. Thus, working with graphs even much larger than the ones we use for evaluation (see Section 6.1) would, in the first place, not be a problem of memory.

5.1 Search Algorithms

Answering RPQs involves several algorithms. In this section, we give an overview of the most important ones.




Fig. 3. Storage schema and query evaluation process of our system

RARE LABEL SEARCH. The first task in executing a query is to find the rare labels. To this end, we go through the automaton representing the query from start to end. For every state, we check if it is mandatory (i.e., does not have a modifier as * or ?). If it is mandatory, we look up the edge label the state represents in the edge label index (which requires constant time). This index gives us the number of occurrences of the label in the graph. If it is below the given threshold, the current state is added to the list of rare states. We consider alternatives (e.g., $a|b$) as rare if the sum of the number of occurrences of all alternatives is below the threshold. Items in brackets can only be mandatory if the bracket itself is mandatory.

TWO-WAY SEARCH. If two or more rare labels have been found, we use a two-way search algorithm to find paths between each two neighboring rare labels. Since rare labels can occur more than once, this is a many-to-many search, starting at the end nodes of all edges with one rare label, and ending at the start nodes of all edges with the next rare label. The search is performed breadth-first by iterating through the graph and the query automaton at the same time. A search state (one specific point during the search process) consists of the current position in the graph and the current state of the automaton. Different search states can be at the same position in the graph but in different states of the automaton, or vice versa. When traversing an edge in the graph, we check if its label is in the *labeled follow set* of the current state. In that case, new entries are added to the end of the list of search states to be processed. One search state is created for each entry in the follow set.

The aim of the search is to find all paths for each pair of fix start and end nodes. A path is found if a forward and a backward search meet at a node and are in the same state of the query automaton. We keep lists for every node where we store in which states a search passed the node. This is used to find completed paths as well as to prevent cycles in the result paths. The search ends once the list of unprocessed search states is empty.

START AND END SEARCH. Searching for the start and end of the path works much like the two-way search algorithm. The only differences are that we search one-way and that we do not end when finding specific nodes, but when hitting a finish state in the query automaton (or a start state, for the backward search).

5.2 Two-Way Search Complexity

Theoretically, all nodes might need to be searched during a two-way search, and every node could be visited in every state of the automaton, and from every start or end node, resulting in a complexity of $O(|V| \cdot |S| \cdot r)$, with S being the states of the automaton and r the number of start plus end nodes, i.e., the occurrences of the rare labels used for that search. In reality, however, the search space is limited because the labels on the path must match the automaton (a query should thus be as small as possible, e.g. not include $a**$ instead of $a*$). Also, the number of start and end nodes is much smaller than $|V|$ (depending on which labels are considered rare). Due to two-way search, the search space is

reduced further in most cases, since complete paths are only half as long from both directions.

Additional checks have to be made during the search, but they do not add to its complexity. For finish and cycle checks, we need to check for all nodes encountered on the path whether they have already been visited in the same state in the same direction (indicating a circle) or in the other direction (indicating a completed path). This can be done in constant time. However, cycle checks might be performed more than once per node, if a node has multiple incoming edges: In the worst case, the check is performed as often as the number of edges in the graph. Thus, the overall worst-case complexity sums up to $O((|V| + |E|) \cdot |S| \cdot r)$.

5.3 Parallelization

The search process can be parallelized in different ways. For a query that contains n rare labels, $n+1$ smaller queries have to be answered. This is performed as $n+1$ independent searches which are executed in parallel. Also, the search algorithm is a many-to-many search if the rare labels occur more than once in the graph. This can also run in parallel, with different threads processing different start points.

6 Experimental Results

In this section, we present an experimental evaluation of our method. We compare our rare-label based algorithm with other implementations available and show results for different graphs and different kinds of queries. Further experiments are devoted to scalability with regard to the size and density of graphs, to the effects of parallelization, different query types, and different label distributions. All tests were executed on a Quad-Core AMD Opteron machine with 16 GB of main memory. The execution times for queries given in the following were gathered by executing 10,000 queries and building the average.

As threshold for the number of path combinations in the rare label optimization (see Section 4.3), we used a value of 100; higher values did not yield any significant changes in runtimes, while lower values lead to slower queries.

6.1 Graphs and Queries

We use real graphs (from biological research) as well as artificially created graphs for the evaluation. We present results for two real-world graphs which we call *AliBaba* and *Extracts*. *AliBaba* is a network of protein-protein-interactions extracted by text mining on all of PubMed [31]. The graph has about 50,000 nodes and 340,000 edges. *Extracts* is a graph of enzymes and their relations, also extracted by text mining from biomedical abstracts, containing about 80,000 nodes and one million edges. Note that these graphs are not toys; such networks today are used regularly in Systems Biology, for instance to improve protein function prediction [17] or disease-gene identification [2]. Their size is roughly comparable

or larger to that of the largest databases of biological networks (the KEGG network currently contains approximately 45.000 nodes), but their density is considerably higher. Thus, they represent rather difficult cases for this domain. Results for other biological networks we tested were similar to those on these two graphs.

To systematically study the scalability of the algorithms with regard to various parameters, we created artificial graphs with sizes between 1000 and one million nodes, with different average degrees and different label distributions. All real and all synthetic graphs are scale-free and roughly have a Zipfian distribution of edge label frequencies. The influence of the type of graph and the label distribution is shown in [21].

For our evaluation, we used both real-life queries from the Bioinformatics domain as well as large, artificially generated sets of queries with similar properties. The biological queries on the AliBaba graph are used in a Systems Biology application analyzing networks of transcription regulation and protein-protein interactions [37]. They consist of 4-12 labels, which may or may not be rare. Artificial queries are used to evaluate the influence of the queries on the runtime (see Section 6.4). The queries have been created with special properties (e.g. number or occurrences of rare labels) and are structurally similar to the biological queries. See the full paper [21] for more details.

6.2 Comparing with other Implementations

We are aware of just two implementations of regular path queries on graphs and considered both as competitors of our algorithm.

RPL [39] evaluates RPQs on RDF graphs using an automata-based implementation. Using the original code provided by the authors, we observed that the system works well on small RDF graphs. However, it cannot handle graphs of the size we target with our work. For answering queries on a graph with approx. 10,000 nodes and 20,000 edges, the system already required 16 GB of main memory. Queries on this graph required already several seconds per query, while our algorithm answers those queries in less than 100ms. All larger graphs, including our real-world biological graphs, could not be handled anymore. Therefore, we do not include this system in the following systematic evaluation.

GLEEN [9] is implemented as an extension to the ARQ library for SPARQL processing. Thus, a comparison of runtimes would be rather unfair, as the ARQ extension mechanism is based on an iterative load-and-verify of single edges of a graph, which cannot be compared to our approach of loading the entire graph into memory at once.

This leaves the re-implementation of the DataGuides system [14] (called AUT from now on) to compare with. Automata-based methods generally work as follows. First, the graph and the query are transformed into NFAs. Both NFAs are then converted into DFAs that are minimized in a third step. Then, both minimized DFAs are intersected. The result is an automaton that is equivalent to the subgraph (of the graph) embracing all paths that match the query. Extracting all matching paths from this subgraph is straight-forward. While this approach works very well for evaluating RPQs on schema-based (i.e. sequences of labels




Fig. 4. Average runtime (log scale) to answer one query on different graphs. The left diagram shows the results for queries without rare labels. On the right, each query contains at least one rare label. Queries on Extracts as well as on the synthetic graph with 100K nodes could not be executed with AUT.

follow a schema) or tree-like data (e.g. XML), it has problems with general graphs and schemaless label distributions, as in our case. For such structures, the DFA generated from the NFA often grows extremely in size (recall that the DFA for an NFA can be exponential in size of the NFA in the worst case). As we show shortly, even mid-sized graphs cannot be processed by this method.

For comparing our RL method with AUT, we use sets of queries with rare labels as well as completely random queries (which may contain an arbitrary number of rare labels or none). As we are not aware of any efficient parallelization scheme for automaton minimization and determinization, we only compare single-threaded versions of both implementations; the additional speed-up that is possible with our algorithm on current (multi-core) hardware will be evaluated in Section 6.5.

Figure 4 shows the average runtime (averaged over 10,000 queries) for answering an RPQ with RL and with AUT. The former is faster in all cases. For queries without rare labels, the runtimes do not differ much for small graphs, but differences get significant for larger graphs. For a graph of 10,000 nodes and 20,000 edges, RL already performs almost two orders of magnitude faster than AUT. For queries that contain at least one rare label, again, RL is always faster, and its superiority increases with graph size; differences for those queries are larger than for queries without rare labels.

Rare labels have a considerable influence on the runtime of RL. In contrast, we found that for AUT, the runtime for different queries on the same graph is about equal. However, time is only one problem of this method; the other is space. The NFA-DFA conversion in AUT incurs an exponential increase in the number of nodes. In our implementation, the whole process requires 350 MB for the graph with 1000 nodes and 2000 edges, but already 3.8 GB for the 10 times larger graph – which still is considerably less space than required by RPL (see above). Running AUT on the AliBaba or Extracts real-life graphs or on the 100K synthetic graph failed due to memory overflow.

6.3 Scalability: Graph Size and Density

To test different scalability aspects of RL, we used artificially created graphs with varying properties. All queries contain at least one rare label. For this




Fig. 5. Average runtime for answering one query on synthetic graphs with different numbers of nodes and edges. The left diagram shows results for a fixed node to edge ratio of 1:2, while the number of nodes is 100,000 on the right.

evaluation, we used the multi-threaded implementation with a fixed number of four threads. We evaluated the effect of our rare-label optimization by comparing it to a baseline method, which performs a brute-force search starting at every node (also in parallel) without considering label frequencies.

Figure 5 (left) shows how RL scales with the size of the graph at a fixed node/edge ratio. The smallest graph has 10K nodes and 20K edges, and the largest graph has 1 million nodes and 2 million edges. Clearly, the scaling of the implementation with the rare-label optimization is much better than that of the baseline. Even for the largest graphs we tested, RL can answer a regular path query in few seconds on average.

Figure 5 (right) shows scalability of RL with the average graph degree. Using multiple artificial graphs with 100,000 nodes, we increased the number of edges (and thus the average degree), leaving all other properties equal. Again, the increase in execution time is favorable compared to the baseline.

Both experiments also show that execution times grow super-linearly with increasing graph size and graph density. Also, the absolute times cannot be compared to those achieved for answering, for instance, reachability queries on graphs of similar size [38,35]. But one should not forget that evaluating RPQs on graphs is a NP-hard problem, whereas reachability can be answered in $O(n^3)$.

6.4 Influence of Query Types

To evaluate the influence of our optimizations on different types of queries, we created ten sets of 1000 queries each differing in the number of occurrences of rare labels. One set does not contain any rare labels (set 0). All other sets contain exactly one rare label, but with an increasing number of occurrences in the graph. The rare labels found in query set 1 appear only once in the graph, rare labels from set 2 exactly twice, and so on. The runtimes for the different query sets are shown in Figure 6.




Fig. 6. Left: average runtime for answering one query from the respective query set. Right: Number of queries taking up to 10 ms, 100ms, 1s, and 10s to execute. The synthetic graph that was used has 10,000 nodes and 20,000 edges.

As expected, answering queries without rare labels is considerably slower than for queries containing rare labels. The difference between queries without rare labels and those with one rare label appearing exactly once in the entire graph is almost three orders of magnitude. Also, queries with a rare label that occurs less frequently in the graph generally are executed faster. However, this trend is partly out-weighted by noise generated through the random complexity of the queries in the workload.

To further study how runtimes change for different queries, we categorized the runtimes of all queries from our query set on a given graph. As Figure 6 shows, more than 90% of all queries are answered in less than 10ms, and more than 95% are answered in less than 100ms. But a few queries take exceptionally longer than all others. These pathological queries are queries that contain only very frequent labels, leading to exceptionally large result sets. For instance, the most difficult query from our set took about 4 seconds and generated a result set of 230 million matching paths. However, we believe that such queries rarely occur in any real application.

We also investigated the influence of query lengths, i.e., the number of labels in the regular expression of the query. Comparing the execution times of queries depending on their lengths, we found that the query length only has a minor influence on execution speed (data not shown). The reason is that, although longer queries are more complicated to answer in general, they often do not have as many matches in the graph, which reduces the search space. For execution speed, the number of occurrences of the labels and the number of rare labels are much more important than the sheer number of labels in the query.

6.5 Parallelization

Figure 7 shows the effect of using multiple threads for RL. The scale-up is very good for up to four threads, but the additional advantage of adding more threads levels out for more than 4 threads. This behavior can be explained by the fact that our current implementation uses additional threads only for additional rare




Fig. 7. Average runtime for answering a query containing rare labels on a graph with 10,000 nodes and 20,000 edges against the number of threads used

labels. For example, a query containing two rare labels uses three threads: One for searching the paths between the rare labels, one for searching to the start and one for searching to the end of the query. Since few queries contain more than three rare labels, the scaling diminishes. However, it would be possible to enhance the implementation by also running the searches between two instances of a rare label in parallel. Since most rare labels appear more than once in the graph, finding paths between them is a many-to-many search and could use different threads starting at different nodes.

7 Conclusion

We presented a novel approach for answering regular path queries on large graphs. Our main idea is to structure a graph traversal along those labels from a query that are infrequent in the graph, but guaranteed to occur in any matching path. We use these rare labels as start-, end-, and way-points during traversal, thus essentially breaking up a very large search space into many much smaller ones. We compare our novel method with a traditional approach using automata and find that the former outperforms the latter over a wide range of different graphs and queries; furthermore, it requires only linear preprocessing and is able to handle much larger graphs. We also showed that using the rare-label optimization considerably improves scalability with regard to the size of the graph and to graph density.

References

1. Abiteboul, S., Quass, D., McHugh, J., Widom, J., Wiener, J.L.: The lorel query language for semistructured data. *Int. Journal on Digital Libraries* 1, 68–88 (1997)
2. Aerts, S., Lambrechts, D., Maity, S., Van Loo, P., et al.: Gene prioritization through genomic data fusion. *Nat. Biotechnol.* 24(5), 537–544 (2006)

3. Aho, A.V., Sethi, R., Ullman, J.D.: Compilers: principles, techniques, and tools. Addison-Wesley Longman Publishing Co., Boston (1986)
4. Alkhateeb, F., Baget, J.-F., Euzenat, J.: Extending SPARQL with regular expression patterns (for querying RDF). *Web Semant.* 7(2), 57–73 (2009)
5. Anyanwu, K., Maduko, A., Sheth, A.: Sparq2l: towards support for subgraph extraction queries in rdf databases. In: WWW 2007, Banff, Alberta, Canada, pp. 797–806 (2007)
6. Chan, C.-Y., Garofalakis, M., Rastogi, R.: Re-tree: an efficient index structure for regular expressions. *The VLDB Journal* 12(2), 102–119 (2003)
7. Cheng, J., Yu, J.X., Ding, B., Yu, P.S., Wang, H.: Fast graph pattern matching. In: ICDE 2008, pp. 913–922. IEEE (2008)
8. Cho, J., Rajagopalan, S.: A fast regular expression indexing engine. In: ICDE 2002, p. 0419 (2002)
9. Detwiler, L.T., Suciu, D., Brinkley, J.F.: Regular paths in sparql: Querying the nci thesaurus. *American Medical Informatics Association*, 161–165 (2008)
10. Dries, A., Nijssen, S., De Raedt, L.: A query language for analyzing networks. In: CIKM 2009, New York, NY, USA, pp. 485–494 (2009)
11. Fan, W., Li, J., Ma, S., Tang, N., Wu, Y.: Adding regular expressions to graph reachability and pattern queries. In: ICDE, pp. 39–50 (2011)
12. Fernandez, M.F., Suciu, D.: Optimizing regular path expressions using graph schemas. In: ICDE 1998, pp. 14–23. IEEE, Washington, DC (1998)
13. Goldman, R., Widom, J.: Dataguides: Enabling query formulation and optimization in semistructured databases. In: VLDB 1997, pp. 436–445 (1997)
14. Goldman, R., Widom, J.: Approximate dataguides. In: Workshop on Query Processing (1999)
15. He, H., Singh, A.K.: Graphs-at-a-time: query language and access methods for graph databases. In: SIGMOD 2008, New York, USA, pp. 405–418 (2008)
16. Hopcroft, J.E., Ullman, J.D.: Introduction to Automata Theory, Languages, and Computation. Addison-Wesley, Reading (1979)
17. Jaeger, S., Gaudan, S., Leser, U., Rebholz-Schuhmann, D.: Integrating protein-protein interactions and text mining for protein function prediction. *BMC Bioinformatics* 9(suppl. 8), S2 (2008)
18. Jin, R., Hong, H., Wang, H., Ruan, N., Xiang, Y.: Computing label-constraint reachability in graph databases. In: Proceedings of the 2010 International Conference on Management of Data, SIGMOD 2010, New York, NY, USA, pp. 123–134 (2010)
19. Kaushik, R., Bohannon, P., Naughton, J.F., Korth, H.F.: Covering indexes for branching path queries. In: SIGMOD Conference, pp. 133–144 (2002)
20. Kochut, K.J., Janik, M.: SPARQLeR: Extended Sparql for Semantic Association Discovery. In: Franconi, E., Kifer, M., May, W. (eds.) ESWC 2007. LNCS, vol. 4519, pp. 145–159. Springer, Heidelberg (2007)
21. Koschmieder, A., Leser, U.: Regular Path Queries on Large Graphs. In: Ailamaki, A., Bowers, S. (eds.) SSDBM 2012. LNCS, vol. 7338, pp. 177–194. Springer, Heidelberg (2012)
22. Kuramochi, M., Karypis, G.: An efficient algorithm for discovering frequent subgraphs. *IEEE Trans. on Knowl. and Data Eng.* 16(9), 1038–1051 (2004)
23. Leser, U.: A query language for biological networks. *Bioinformatics* 21(2), 33–39 (2005)
24. Li, L., Alderson, D., Tanaka, R., Doyle, J.C., Willinger, W.: Towards a theory of scale-free graphs: Definition, properties, and implications (ext. version). *Internet Mathematics* 2(4), 431–523 (2006)

25. Li, Q., Moon, B.: Indexing and querying XML data for regular path expressions. In: VLDB 2001, Roma, Italy, pp. 361–370 (2001)
26. Malewicz, G., et al.: Pregel: a system for large-scale graph processing. In: PODC 2009, New York, NY, USA, p. 6 (2009)
27. Martínez-Bazan, et al.: Dex: high-performance exploration on large graphs for information retrieval. In: CIKM 2007, New York, NY, USA, pp. 573–582 (2007)
28. Mendelzon, A.O., Wood, P.T.: Finding regular simple paths in graph databases. *SIAM Journal on Computing* 24(6), 1235–1258 (1995)
29. Milo, T., Suciu, D.: Index Structures for Path Expressions. In: Beeri, C., Brune-man, P. (eds.) ICDT 1999. LNCS, vol. 1540, pp. 277–295. Springer, Heidelberg (1998)
30. Neven, F.: Automata theory for xml researchers. *SIGMOD Rec.* 31(3), 39–46 (2002)
31. Palaga, P., Nguyen, L., Leser, U., Hakenberg, J.: High-performance information extraction with alibaba. In: EDBT 2009, New York, USA, pp. 1140–1143 (2009)
32. Ronen, R., Shmueli, O.: SoQL: A language for querying and creating data in social networks. In: ICDE 2009, Shanghai, China, pp. 1595–1602 (2009)
33. San Martín, M., Gutierrez, C.: Representing, Querying and Transforming Social Networks with RDF/SPARQL. In: Aroyo, L., Traverso, P., Ciravegna, F., Cimi-ano, P., Heath, T., Hyvönen, E., Mizoguchi, R., Oren, E., Sabou, M., Simperl, E. (eds.) ESWC 2009. LNCS, vol. 5554, pp. 293–307. Springer, Heidelberg (2009)
34. Sevón, P., Eronen, L.: Subgraph queries by context-free grammars. *Journal of Integrative Bioinformatics* 5(2), 100 (2008)
35. Trißl, S., Leser, U.: Fast and practical indexing and querying of very large graphs. In: SIGMOD 2007, New York, NY, USA, pp. 845–856 (2007)
36. Yan, X., Yu, P.S., Han, J.: Graph indexing: a frequent structure-based approach. In: SIGMOD 2004, New York, NY, USA, pp. 335–346 (2004)
37. Yeger-Lotem, E., Sattath, S., Kashtan, N., et al.: Network motifs in integrated cellular networks of transcription-regulation and protein-protein interaction. *Proc. Natl. Acad. Sci. USA* 101(16), 5934–5939 (2004)
38. Yildirim, H., Chaoji, V., Zaki, M.J.: Grail: Scalable reachability index for large graphs. In: VLDB 2010. VLDB Endowment (2010)
39. Zauner, H., Linse, B., Furche, T., Bry, F.: A RPL through RDF: Expressive Navigation in RDF Graphs. In: Hitzler, P., Lukasiewicz, T. (eds.) RR 2010. LNCS, vol. 6333, pp. 251–257. Springer, Heidelberg (2010)
40. Zou, L., Chen, L., Özsu, M.T.: Distance-join: Pattern match query in a large graph database. *PVLDB* 2(1), 886–897 (2009)