

GRAIL: A Scalable Index for Reachability Queries in Very Large Graphs

Hilmi Yıldırım · Vineet Chaoji · Mohammed J. Zaki

Received: date / Accepted: date

Abstract Given a large directed graph, rapidly answering reachability queries between source and target nodes is an important problem. Existing methods for reachability trade-off indexing time and space versus query time performance. However, the biggest limitation of existing methods is that they do not scale to very large real-world graphs. We present a simple yet scalable reachability index, called GRAIL, that is based on the idea of randomized interval labeling, and that can effectively handle very large graphs. Based on an extensive set of experiments, we show that while more sophisticated methods work better on small graphs, GRAIL is the only index that can scale to millions of nodes and edges. GRAIL has linear indexing time and space, and the query time ranges from constant time to being linear in the graph order and size.

1 INTRODUCTION

Let $G = (V, E)$ be a directed graph, where V is the set of vertices and E is the set of directed edges, with $|V| = n$ and $|E| = m$. A *reachability query* asks if there exists a path p from a source node u to a target node v in the directed graph

G . If such a path exists, we say that u can reach v (or v is reachable from u), and denote it as $u \rightsquigarrow v$. If u cannot reach v , we denote it as $u \not\rightsquigarrow v$. The reachability query itself is denoted as $u \overset{?}{\rightsquigarrow} v$.

Answering graph reachability queries quickly has been the focus of research for over 20 years. Traditional applications include reasoning about inheritance in class hierarchies and testing concept subsumption in knowledge representation systems. However, interest in the reachability problem has revived in recent years with the advent of new applications which have very large graph-structured data that are queried for reachability repeatedly. The emerging area of Semantic Web is composed of RDF/OWL data which are indeed graphs with rich content, and there exist RDF data with millions of nodes and billions of edges. Reachability queries are often necessitated on these data to infer the relationships among the objects. In network biology, reachability play a role in querying protein-protein interaction networks, metabolic pathways and gene regulatory networks. As specific examples of graph reachability consider the following applications.

Semantic Query Engines: The vision of the semantic web is to allow machines to understand the meaning of the information on the world wide web [20]. The data stores and interchange formats used include RDF and XML. XML documents are typically trees, but with the use of ID/IDREF links it is possible to represent graphs. On the other hand, RDF documents represent relationships and are naturally graph structured. XQuery and SPARQL are the query languages for XML and RDF respectively. SPARQL lets the user define and search for graph patterns using SQL-like syntax, over RDF stores. With SPARQL 1.1 it is possible to search for patterns that include paths of arbitrary length via *property paths* [30]. For example, the following SPARQL query asks for all researchers named *Joe Example* who might

This work was supported in part by NSF Grants EMT-0829835, and CNS-0103708, and NIH Grant 1R01EB0080161-01A1.

H. Yıldırım
Rensselaer Polytechnic Institute
E-mail: yildih2@cs.rpi.edu

V. Chaoji
Yahoo! Labs, Bangalore India
E-mail: chaojv@yahoo-inc.com

M. J. Zaki
Rensselaer Polytechnic Institute
Tel.: +1-518-2766340
Fax: +1-518-2764033
E-mail: zaki@cs.rpi.edu

be influenced by the paper titled *Some Paper* in a citation database.

```

SELECT ?name
WHERE {
  ?x foaf:name ?name .
  ?y citation:title ?title .
  ?y citation:author ?x .
  ?x citation:cites+/citation:title ?pname .
  FILTER regex(?name, "Joe Example")
  FILTER regex(?pname, "Sample Paper")
}

```

Fig. 1 A SPARQL query that uses property paths

One strategy to answer this query is to use a reachability index to check if any paper of *Joe Example* can reach *Some Paper* by citations links (*cites+*). With SPARQL 1.1 the necessity for reachability indexes in RDF stores will become more apparent.

Concept Subsumption in Ontologies: In knowledge representation systems, ontologies define a set of concepts as well as their properties and inter-relationships. Ontologies are typically used to describe a domain or to reason about the entities within a domain. In semantic web, inference engines generate all the possible conclusions and store the resulting triples in the datastore as well. For instance, if a class *A* is a subclass of *B* which is a subclass of *C*, it can be inferred that *A* is a subclass of *C*, and this fact is stored in the database. For instance, the Gene Ontology [1], which consists of a directed acyclic graph of terms describing genes, also includes the inferred relationships to speed up querying [2]. This is equivalent to keeping the transitive closure of the graph in the database, which is not scalable for very large graphs. An alternate strategy is to avoid inference computation at least for transitive rules such as *rdfs:subPropertyOf* and *rdfs:subClassOf* [33]. In this case the inference engine should infer these relations on-demand using a reachability labeling mechanism. One of the test cases in our experimental evaluation in Section 6 is composed of the terms of the Gene Ontology and the gene products from the Uniprot [3] annotations database. The overall data has more than 7 million nodes and 56 million connections.

Biological Networks: With the advance of high-throughput data acquisition technologies, biologists have amassed a large amount of heterogeneous graph data such as metabolic pathway databases, gene regulatory networks, signal transduction networks, etc. In these datasets nodes represent entities such as proteins, genes, compounds and edges represent how they interact. A query whether a gene *A* regulates gene *B* directly or indirectly in a gene regulatory network maps to a reachability query. Similarly one can ask the question of whether a specific amino acid *A* is broken down to a specific chemical compound *B* in a metabolic pathway graph, which is also a reachability query.

In general, given the ubiquity of large graphs in the form of ontologies, social networks, biological networks, and so on, there is a crucial need for highly scalable indexing schemes for reachability queries.

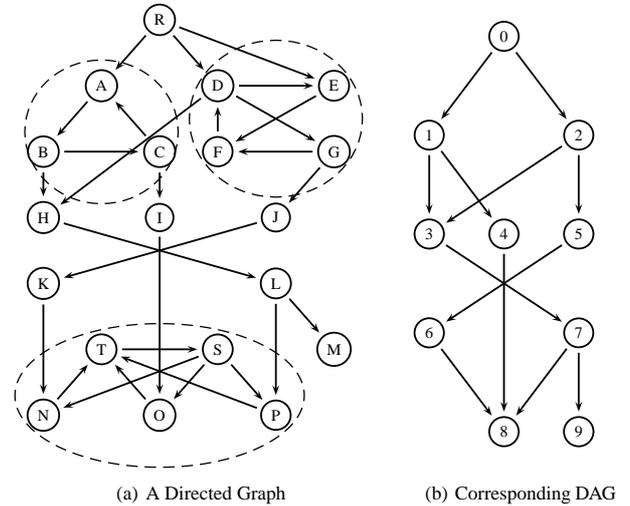


Fig. 2 Coalescing strongly connected components

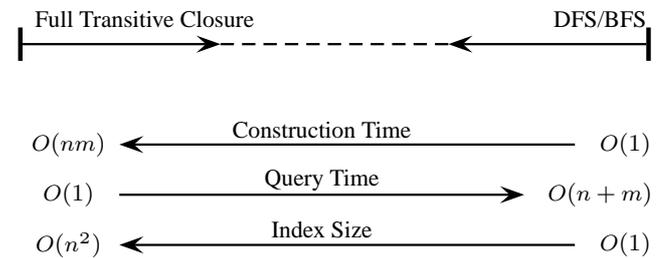


Fig. 3 Tradeoff between Query Time and Index Size

1.1 Reachability Index

It is worth noting at the outset that the problem of reachability on directed graphs can be reduced to reachability on directed acyclic graphs (DAGs). Given a directed graph *G*, we can obtain an equivalent DAG, *G'* (called the *condensation graph* of *G*), in which each node corresponds to a strongly connected component of the original graph, and an edge exists between two components if any node in one component can reach a node in the other component. To answer whether node *u* can reach *v* in *G*, we simply look up their corresponding strongly connected components, *S_u* and *S_v*, respectively, which are the nodes in *G'*. If *S_u* = *S_v*, then by definition *u* can reach *v* (and vice-versa). If *S_u* ≠ *S_v*, then we pose the question whether *S_u* can reach *S_v* in

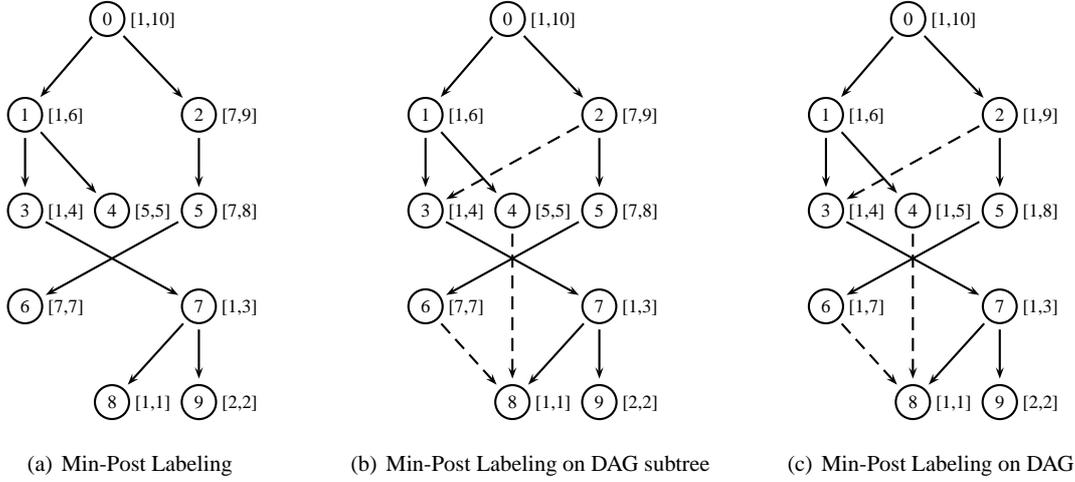


Fig. 4 Min-Post Labeling

G' . Thus all reachability queries on the original graph can be answered on the DAG. Figure 2 shows an example directed graph and the corresponding DAG obtained by coalescing the strongly connected components that are marked in dashed circles. In that graph $N \overset{?}{\rightsquigarrow} P$ is true since they are in the same component. The query of $G \overset{?}{\rightsquigarrow} P$ in the original graph is equivalent to the query of $2 \overset{?}{\rightsquigarrow} 8$ in the corresponding DAG. Henceforth, we assume that all input graphs have been transformed into their corresponding DAGs, unless otherwise stated.

There are two basic approaches to answer the reachability queries which lie at the two extremes of the index design space, as illustrated in Figure 3. Given a DAG G , with n vertices and m edges, one extreme (shown on left) is to pre-compute and store the full transitive closure; this allows one to answer reachability queries in constant time by a single lookup, but it unfortunately requires a quadratic space index, making it practically unfeasible for large graphs. On the other extreme (shown on right), one can use a depth-first (DFS) or breadth-first (BFS) traversal of the graph starting from node u , until either the target, v , is reached or it is determined that no such path exists. This approach requires no index, but requires $O(n + m)$ time for each query, which is unacceptable for large graphs. Existing approaches to graph reachability indexing lie in-between these two extremes.

While there is no single best indexing scheme for DAGs, the reachability problem on trees can be solved effectively by *interval labeling* [18], which takes linear time and space for constructing the index, and provides constant time querying. Given a tree/forest $T = (V, E)$ that has $|V| = n$ nodes and $|E| = m$ edges, interval labeling assigns each node u a label $L_u = [s_u, e_u]$ which represents an interval starting at s_u and ending at e_u . A desired labeling has to satisfy the condition that L_u subsumes L_v if and only if u reaches v .

A reachability query $u \overset{?}{\rightsquigarrow} v$ can then be answered by just comparing the corresponding intervals, i.e., $u \rightsquigarrow v$ if and only if $L_v \subseteq L_u$ ($s_u \leq s_v$ and $e_u \geq e_v$).

Consider the *min-post labeling* method for trees, which assigns $L_u = [s_u, e_u]$ to each node u where $e_u = \text{post}(u)$ is the post-order value of the node u , defined as the rank of the node u in a post order traversal of the tree. The starting value of the interval is defined as $s_u = \min\{s_x | x \in \text{children}(u)\}$ if u is not a leaf, and $s_u = e_u$ if u is a leaf. It is easy to see that s_u is equivalent to $\min\{e_x | x \in \text{descendants}(u)\}$, that is, s_u is the lowest post-order rank for any node x in the subtree rooted at u (i.e., including u). Figure 4(a) shows the min-post labeling for an example tree. It is easy to see that reachability queries in trees can be answered by interval containment. For example, $1 \rightsquigarrow 9$, since $L_9 = [2, 2] \subseteq [1, 6] = L_1$, but $2 \not\rightsquigarrow 7$, since $L_7 = [1, 3] \not\subseteq [7, 9] = L_2$.

Interval labeling can be generalized to DAGs. Existing methods [18, 31, 8, 23, 4] use either min-post labeling or pre-post labeling (see Section 2) on a subtree of the DAG, as shown in Figure 4(b) (the dashed edges are the non-tree edges). The labeling obtained is the same as in Figure 4(a), because the same tree is used in both approaches, and dashed edges are never followed. However, we can easily see that for DAGs, the labeling is not sufficient to answer reachability queries, since the labeling fails to cover some of the reachable pairs. For example, $2 \rightsquigarrow 7$ but $L_7 = [1, 3] \not\subseteq [7, 9] = L_2$. Existing methods thus have to supplement the index by auxiliary indexes or some other approaches to correctly answer the queries.

1.2 Our Contributions

In this paper we present a novel and scalable graph indexing approach for very large graphs called GRAIL¹, which stands for **Graph Reachability Indexing via RAnDomized Interval Labeling**. GRAIL applies min-post labeling directly on the DAG, as opposed to a subtree of the DAG, as in other interval labeling variants. Figure 4(c) shows the GRAIL labeling on the example DAG. This labeling captures all reachable pairs at the cost of falsely categorizing some pairs as reachable. For instance, compared to Figure 4(b), the label of node 6 is enlarged to $[1, 7]$ since node 8 is a descendant of node 6 and its label is $[1, 1]$. With this new labeling, $2 \rightsquigarrow 7$ can be answered using label containment (i.e., $L_7 = [1, 3] \subseteq [1, 9] = L_2$), whereas it cannot be answered using the labeling only on the DAG subtree as in Figure 4(b). However, the new labeling introduces false-positives, also called as *exceptions*. For example, $L_4 = [1, 4] \subseteq [1, 8] = L_5$ but $5 \not\rightsquigarrow 4$. GRAIL uses different strategies to handle such exceptions, as detailed in Section 3.

In this paper we make the following original contributions:

- To our knowledge GRAIL is the first direct DAG interval labeling approach. The key idea is to do very fast elimination for those pairs of query nodes for which non-reachability can be determined via the intervals. In other words, if $L_v \not\subseteq L_u$, we immediately return $u \not\rightsquigarrow v$.
- Instead of using a single interval per node, the basic approach GRAIL uses multiple intervals obtained via randomized DAG traversals, a technique we call *randomized multiple interval labeling*. This approach drastically reduces the number of exceptions. We also empirically compare alternative multiple interval labeling methods designed to minimize the number of exceptions.
- To guarantee correctness, GRAIL uses “smart” graph search methods with recursive label-based pruning. We study the impact of DFS/BFS search, and also explore the use of bi-directional BFS. Independently, these search methods also form optimized baseline methods for comparison against GRAIL. In addition, we propose an alternative method to directly maintain *exception lists* per node, which can eliminate false-positives completely. This can offer benefit for indexing smaller graphs, with pruning based search as the scalable alternative for massive graphs.
- We propose several enhancements to the basic GRAIL approach, which can dramatically improve the querying time. The *topological level filter* is a simple yet effective strategy to prune non-reachable pairs of nodes. The *positive cut filter* uses the subtree labelings induced by each one of GRAIL’s multiple DAG labelings to guarantee

reachability for a subset of all the reachable pairs. For these guaranteed reachable pairs u and v , if $L_v \subseteq L_u$, then GRAIL immediately returns $u \rightsquigarrow v$.

- Via an extensive set of experiments, we show that GRAIL outperforms existing methods on very large real and synthetic graphs, sometimes by over an order of magnitude. Whereas previous methods studied graphs mainly under 50K nodes and 70K edges when transformed into DAGs, we use large real and synthetic graphs with millions of nodes and edges (see Tables 5 and 6). In many cases, GRAIL is the only method that can even run on such large graphs.

GRAIL is a highly effective and light-weight index to answer reachability queries in massive graphs with millions of nodes and edges. It maintains d interval labels per node, and thus its index size is $O(dn)$. The index construction time for GRAIL is $O(d(n + m))$, since d random graph traversals are made to obtain those labels. For reachability queries, it takes only $O(d)$ time if the queries node pairs are not reachable. Likewise, if the node pairs are part of the guaranteed reachable pairs, then again it takes only $O(d)$ time. In other cases, GRAIL resorts to pruning-based recursive graph search, which can take $O(n + m)$ in the worst case. GRAIL thus retains the best properties on both ends of indexing extremes shown in Figure 3. Since d is typically a small constant (usually capped at $d = 5$ even for the largest graphs), GRAIL requires index construction time and space linear in the graph size and order. It is worth noting that since most large real-world graphs are very sparse, most (random) query node pairs are non-reachable, and these queries can be answered by GRAIL in constant time. Further, due to the positive cut filter, a significant portion of reachable pairs can also be answered in constant time. When either of these fail, GRAIL can take time linear in the graph to answer the query. However, here too the time increases gracefully, since interval-based pruning is applied recursively during the search.

2 RELATED WORK

Existing approaches for graph reachability combine aspects of indexing and pure search, trading off index space for querying time. Major approaches can be classified into two main groups, namely interval labeling [4, 31, 32, 23, 10], and 2HOP labeling [14, 29, 28, 35, 11, 19]. These methods are discussed below, and summarized in Table 1.

The interval labeling approaches use either the min-post labeling [4] (illustrated in Figure 4(b)), or pre-post labeling [18, 31, 8, 23] (illustrated in Figure 5), on a spanning subtree of the DAG. *Pre-post labeling* assigns $L_u = [s_u, e_u]$ to each node u where s_u and e_u are the pre-order and post-order ranks of node u in a DFS traversal of the DAG, start-

¹ A preliminary version of GRAIL appears in [34].

| | Construction Time | Query Time | Index Size |
|---------------------|--------------------|---------------|----------------|
| Opt. Tree Cover [4] | $O(nm)$ | $O(n)$ | $O(n^2)$ |
| GRIPP [31] | $O(m+n)$ | $O(m-n)$ | $O(m+n)$ |
| Dual Labeling [32] | $O(n+m+t^3)$ | $O(1)$ | $O(n+t^2)$ |
| PathTree [23] | $O(mk)$ or $O(mn)$ | $O(\log^2 k)$ | $O(nk)$ |
| 2HOP [14] | $O(n^4)$ | $O(\sqrt{m})$ | $O(n\sqrt{m})$ |
| HOPI [29] | $O(n^3)$ | $O(\sqrt{m})$ | $O(n\sqrt{m})$ |
| GRAIL (this paper) | $O(d(n+m))$ | $O(d)/O(n+m)$ | $O(dn)$ |

Table 1 Comparison of Approaches: n denotes number of vertices; m , number of edges; $t = O(m-n)$, number of non-tree edges; k number of paths/chains; and d number of intervals.

ing from the root(s), with the rank being incremented each time we enter a node or back-track from a node, as shown in Figure 5. In contrast, as noted above, in a min-post labeling, e_u is the post-order rank of u , and s_u is the minimum rank of any node under u , as illustrated in Figure 4. The interval-based indexing is not self-contained since the labeling can fail to cover some of the reachable pairs. For example, in Figure 5, $2 \rightsquigarrow 7$ but $L_7 = [4, 9] \not\subseteq [14, 19] = L_2$. Thus for completeness the existing methods supplement the label index with auxiliary indexes. Different interval-based graph indexing methods differ from each other in terms of how the subtree of the DAG is selected, and what kind of auxiliary index is used.

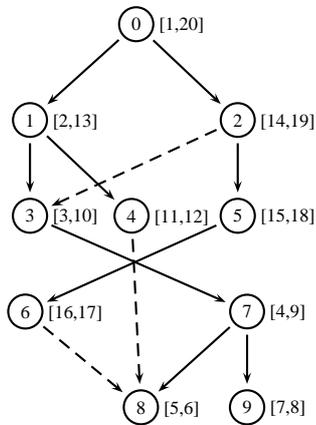


Fig. 5 Pre-Post Labeling on a DAG

Optimal Tree Cover [4] is the first known variant of interval labeling for DAGs. The approach first creates interval labels for a spanning tree of the DAG. This is not enough to correctly answer reachability queries, as mentioned above. To guarantee correctness, the method processes nodes in reverse topological order for each non-tree edge (i.e., an edge that is not part of the spanning tree) between u and v , with u inheriting all the intervals associated with node v . Furthermore node u inherits the intervals of its other children whose intervals are updated. Thus u is guaranteed to contain all of its children's intervals. Testing reachability is equivalent to deciding whether the interval list of the source node sub-

sumes the first interval of the target node. The construction complexity of this method is the same as a full transitive closure because selecting the tree that provides optimal labeling requires the precomputation of the transitive closure.

GRIPP [31] is another variant of interval labeling. Instead of inflating the index size for the non-tree edges as in [4], reachability testing is done via multiple containment queries. Given nodes u and v , if L_v is not contained in L_u , the non-tree edges (x, y) , such that x is a descendant of u , are fetched, and recursively a new query (y, v) is issued for every y , until either v is reachable from a y node or if all non-tree edges are exhausted. If one of the y nodes can reach v then u can reach v . Since there are $m-n$ non-tree edges, the query time complexity is $O(m-n)$.

Dual labeling [32] also uses interval labeling but it processes non-tree edges in a different way. Their main observation is that if there exists a single non-tree edge (x, y) in the path from u to v , it must be true that L_u contains x and L_y contains v . Based on this, a non-tree edge $e = (x, y)$ is connected to another non-tree edge $e' = (x', y')$ if and only if L_y contains $L_{x'}$. After labeling the selected tree, *dual labeling* computes the transitive closure of non-tree edges so that the entry for the edge pair $(e = (x, y), e' = (x', y'))$ being 1 implies that all nodes u , whose interval contains x , can reach all nodes v whose interval is contained by $L_{y'}$. Therefore for each query they scan relevant edge pairs to find out the reachability. With further optimizations, they reduce the query time to $O(1)$, however their index size is $O(n+t^2)$, and construction time is $O(n+m+t^3)$ where $t = O(m-n)$ denotes the number of non-tree edges. GRIPP and Dual Labeling thus lie on the opposite sides of the trade-off illustrated in Figure 3.

A chain decomposition approach was proposed in [21] to compress the transitive closure. The graph is split into node-disjoint chains. A node u can reach to node v if they exist in the same chain, and u precedes v . Each node also keeps the highest node that it can reach in every other chain. Thus the space requirement is $O(kn)$ where k is the number of chains. Such a chain decomposition is computed in $O(n^3)$ time. This bound was improved in [10], where they proposed a decomposition which can be computed in $O(n^2 + kn\sqrt{k})$ time. Recently, [9] further improved this scheme by using

general spanning trees in which each edge corresponds to a path in the original graph. [6] solves a variant of the reachability problem where the input is assumed to be a collection of non-disjoint paths instead of a graph.

PathTree [23] is the generalization of the tree cover approach. It extracts the disjoint paths of a DAG, then creates a tree of paths on which a variant of interval labeling is applied. That labeling captures most of the transitive information and the rest of the closure is computed in an efficient way. PathTree has very fast querying and construction times, but its index size might get very large on dense graphs (k in Table 1 denotes the number of paths in the decomposition). In a recent paper by the same authors, they proposed 3HOP [22] which addresses the issue of large index size. Although 3HOP has a reduced index size, the construction and query times degrade significantly.

The other major class of graph reachability indexing methods is based on *2HOP Indexing* [14, 29, 28, 35, 11, 19], where each node determines a set of intermediate nodes it can reach, and a set of intermediate nodes which can reach it. The query between u and v returns success if the intersection of the successor set of u and predecessor set of v is not empty. 2HOP was first proposed in [14], where they also showed that computing the minimum 2HOP cover is NP-Hard, and gave an $O(\log m)$ -approximation algorithm based on a greedy algorithm for the set-cover problem. Its $O(n^4)$ construction time was improved in [35] by using a geometric approach which produces slightly larger 2HOP cover than obtained in [14]. A divide-and-conquer strategy to 2HOP indexing was proposed in [29, 28]. HOPI [29] partitions the graph into k subgraphs, computes the 2HOP indexing within each subgraph and finally merges their 2HOP covers by processing the cross-edges between subgraphs. [28], by the same authors, improved the merge phase by changing the way in which cross-edges between subgraphs are processed. [11] partition the graph in a top-down hierarchical manner, instead of a flat partitioning into k subgraphs. The graph is partitioned into two subgraphs repeatedly, and then their 2HOP covers are merged more efficiently than in [28]. Their approach outperforms existing 2HOP approaches in large and dense datasets.

The HLSS [19] method is a hybrid of 2HOP and interval labeling. It first labels a spanning tree of the graph with interval labeling, and then extracts a remainder graph whose transitive closure is yet to be computed. The transitive closure of the remainder graph is computed and compressed by a variant of 2HOP labeling. The overall time complexity for constructing HLSS is $O(m^3)$ but it produces more compressed labelings.

Although there has been much interest in static transitive closure, not much attention has been paid to practical algorithms for the dynamic case, though several theoretical studies exist [27, 24, 16]. Practical works on dynamic transi-

tive closure [17, 25] and dynamic 2HOP indexing [7] have only recently been proposed.

A preliminary version of the GRAIL appeared in [34]. GRAIL is also utilizes interval labeling to construct the index. However, it is diametrically opposite to the other interval labeling approaches mentioned above. Previous methods provide an incomplete coverage of the reachable set, i.e., they use spanning trees, chain covers, path decompositions, and so on, to index the reachability for nodes only in a subtree (or a set of chains/paths) of the input DAG. If $u \rightsquigarrow v$, these methods do not guarantee that $L_v \subseteq L_u$; this is true only for the paths over tree edges. They thus have to use auxiliary indexes to ensure completeness. In contrast, GRAIL uses min-post traversals directly on the DAG, and it fully covers all reachable pairs, as well as some non-reachable pairs – the false positives. GRAIL guarantees that if $u \rightsquigarrow v$, then $L_v \subseteq L_u$. By contraposition, if the interval label of v is not contained in that for u , GRAIL can immediately return $u \not\rightsquigarrow v$. However, due to the false positives, GRAIL uses graph search, and other optimizations to answer reachability. The GRAIL approach is more favorable in large, sparse graphs, since it is more likely that a random pair of nodes is not reachable, and for these GRAIL can immediately prune the search. GRAIL can also handle higher density large graphs, since in these cases, the subtree/path based interval labeling of previous approaches covers an even smaller fraction of the reachable pairs. A detailed description of the basic GRAIL approach and other novel optimizations, and an extensive set of supporting experiments appears in the sections below.

3 GRAIL: Scalable Reachability Index for Large Graphs

In this section we present our novel and scalable GRAIL approach for indexing very large graphs. As opposed to other interval labeling variants, GRAIL uses min-post labeling directly on the directed acyclic graph. Furthermore, instead of using a single interval, GRAIL employs multiple min-post intervals that are obtained via random graph traversals (or other strategies). We use the symbol d to denote the number of intervals per node, which also corresponds to the number of graph traversals used to obtain the node label. For example, Figure 6(c) shows a DAG labeling using $d = 2$ intervals (the first interval assumes a left-to-right ordering of the children, whereas the second interval assumes a right-to-left ordering).

3.1 Basic GRAIL Approach

Our approach to reachability indexing is motivated by the observation that existing interval labeling variants identify

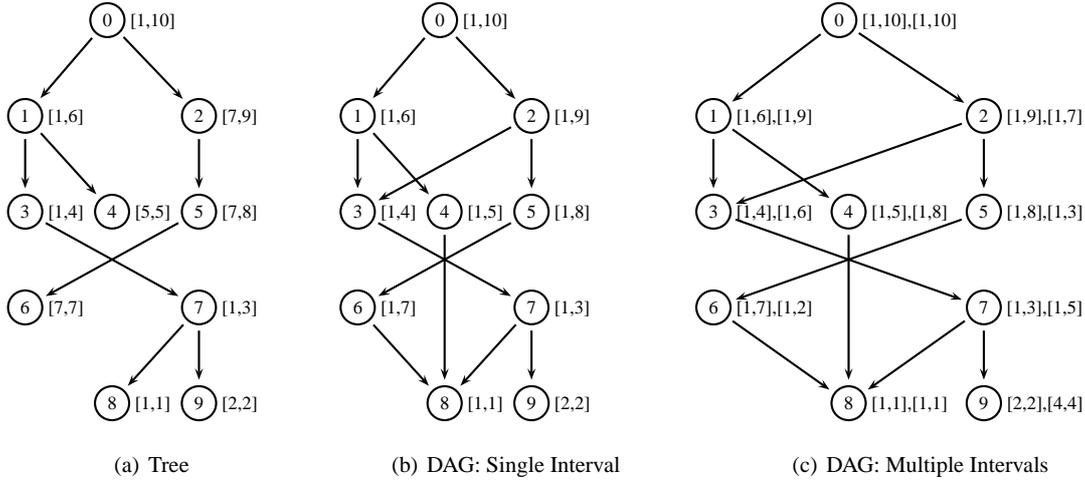


Fig. 6 Interval Labeling: Tree (a) and DAG: Single (b) & Multiple (c)

a subgraph of the DAG (i.e., trees in [4,32,31] and path-tree in [23]) in the first stage, and incorporate the remaining (uncovered) portion of the DAG, in the second phase of indexing or during the query time. However, most of the reachability information is captured in the first stage. The motivating idea in GRAIL is to use interval labeling multiple times to reduce the workload of the second phase of indexing or the querying. The multiple intervals yield a hyper-rectangle instead of single interval per node, and the containment check is over these hyper-rectangles.

| node | exceptions (E) | direct (E^d) | indirect (E^i) |
|------|--------------------|------------------|--------------------|
| 2 | {1, 4} | \emptyset | {1, 4} |
| 4 | {3, 7, 9} | {3, 7, 9} | \emptyset |
| 5 | {1, 3, 4, 7, 9} | \emptyset | {1, 3, 4, 7, 9} |
| 6 | {1, 3, 4, 7, 9} | {1, 3, 4, 7, 9} | \emptyset |

Table 2 Exceptions for DAG in Figure 6(b)

In GRAIL, for a given node u , the new label is given as $L_u = L_u^1, L_u^2, \dots, L_u^d$, where L_u^i is the interval label obtained from the i -th (random) traversal of the DAG, and $1 \leq i \leq d$, where d is the *dimension* or number of intervals. We say that L_v is contained in L_u , denoted as $L_v \subseteq L_u$, if and only if $L_v^i \subseteq L_u^i$ for all $i \in [1, d]$. If $L_v \not\subseteq L_u$, then we can conclude that $u \not\rightsquigarrow v$, as per the lemma below.

Lemma 1 *If $L_v \not\subseteq L_u$, then $u \not\rightsquigarrow v$.*

PROOF: *Given that $L_v \not\subseteq L_u$, there must exist a “dimension” i , such that $L_v^i \not\subseteq L_u^i$. Assume that $u \rightsquigarrow v$, and let x and y be the lowest ranked nodes under u and v , respectively, in the post-order traversal. In this case $L_v^i = [r_y, r_v]$ and $L_u^i = [r_x, r_u]$, where r_n denotes the rank of node n . But $u \rightsquigarrow v$ implies that $r_u > r_v$ in post-order, and further that $r_x \leq r_y$, which in turn implies that $L_v^i = [r_y, r_v] \not\subseteq [r_x, r_u] = L_u^i$.*

$[r_x, r_u] = L_u^i$. But this is a contradiction to our assumption that $u \rightsquigarrow v$. We conclude that $u \not\rightsquigarrow v$.

On the other hand, if $L_v \subseteq L_u$, it is possible that this is a false positive, i.e., it can still happen that $u \not\rightsquigarrow v$. We call such a false positive containment an *exception*. For example, in Figure 6(b), there are 15 exceptions in total, as listed in Table 2. For instance, for node 2, node 1 is an exception, since $L_1 = [1, 6] \subseteq [1, 9] = L_2$, but in fact $2 \not\rightsquigarrow 1$. The basic intuition in GRAIL is that using multiple random labels makes it more likely that such false containments, i.e., exceptions, are minimized. For example, when one considers the 2-dimensional intervals given in Figure 6(c), for the very same graph, 12 out of the 15 exceptions get eliminated! For instance, we see that 1 is no longer an exception for 2, since $L_1 = [1, 6], [1, 9] \not\subseteq [1, 9], [1, 7] = L_2$, since for the second interval we have $[1, 9] \not\subseteq [1, 7]$. We can thus conclude that $2 \not\rightsquigarrow 1$. However, note that 3 is still an exception for 4 since $L_3 = [1, 4], [1, 6] \subseteq [1, 5], [1, 8] = L_4$. For 4, nodes 7 and 9 also remain as exceptions. In general using multiple intervals drastically cuts down on the exception list, but is not guaranteed to completely eliminate exceptions.

There are two main issues in GRAIL: i) how to compute the d random interval labels while indexing, and ii) how to deal with exceptions, while querying. We will discuss these in detail below.

3.2 Index Construction

The index construction step in GRAIL is very straightforward; we generate the desired number of post-order interval labels by simply changing the visitation order of the children randomly during each depth-first traversal. Algorithm 1 shows an implementation of this strategy; an interval L_u^i is

denoted as

$$L_u^i = [s_u^i, e_u^i]$$

The number of possible labelings is exponential, but most graphs can be indexed very compactly with a small number of dimensions depending on the edge density of the graph. Furthermore, since it is not guaranteed that all exceptions will be eliminated, the best strategy is to cease labeling after a small number of dimensions (such as 5), with reduced exceptions, rather than trying to totally eliminate all exceptions, which might require a very large number of dimensions.

Algorithm 1: GRAIL Indexing: Randomized Intervals

```

RandomizedLabeling( $G, d$ ):
1 foreach  $i \leftarrow 1$  to  $d$  do
2    $r \leftarrow 1$  // global variable: rank of node
3    $Roots \leftarrow \{n : n \in roots(G)\}$ 
4   foreach  $x \in Roots$ , in random order do
5      $\lfloor$  Call RandomizedVisit( $x, i, G$ )
      $\rfloor$ 

RandomizedVisit( $x, i, G$ ) :
6 if  $x$  visited before then return
7 foreach  $y \in Children(x)$ , in random order do
8    $\lfloor$  Call RandomizedVisit( $y, i, G$ )
    $\rfloor$ 
9  $r_c^* \leftarrow \min\{s_c^i : c \in Children(x)\}$ 
10  $L_x^i \leftarrow [\min(r, r_c^*), r]$ 
11  $r \leftarrow r + 1$ 
  
```

In terms of the traversal strategies, we aim to generate labelings that are as different from each other as possible. We experimented with the following traversal strategies.

Randomized: This is the strategy shown in Algorithm 1, with a random traversal order for each dimension.

Randomized Pairs: In this approach, we first randomize the order of the roots and children, and fix it. We then generate pairs of labeling, using left-to-right (L-R) and right-to-left (R-L) traversals over each node's children. The intuition is to make the intervals as different as possible; a node that is visited first in L-R order is visited last in R-L order. An example of such a pair is shown in Figure 6(c).

Heuristic (Guided) Traversals: It is also possible to use deterministic approaches with the hope of eliminating exceptions as much as possible. Intuitively, during the i -th traversal, the idea would be to choose the node with the most number of exceptions in the first $i - 1$ dimensions. However computing the number of exceptions after each traversal is expensive; instead, we experiment with the following heuristic approaches. For each of the approaches, during the i -th traversal, from any node in the graph, we choose the child that maximizes the given objective:

- **Maximum Volume:** The volume of a node is defined as the volume of its hyper-rectangle in the first $i - 1$

traversals, given as

$$vol(u) = \prod_{j=1}^{i-1} (e_u^j - s_u^j)$$

Choosing u 's child in decreasing order of volume is based on the intuition that a larger volume may contain more exceptions than a smaller volume.

- **Maximum of Minimum Interval:** The volume of node u can be large even if in one of the first $i - 1$ traversals it has a tight interval. Therefore, another approach is to sort the children of u in decreasing order of their minimum interval ranges in the first $i - 1$ traversals. The objective is given as:

$$mi(u) = \min_{j=1}^{i-1} \{(e_u^j - s_u^j)\}$$

- **Maximum Adjusted Volume:** The nodes which have bigger reachable sets (the number of nodes they can reach) are expected to have larger intervals and thus larger volumes. Therefore a larger volume does not directly imply larger number of exceptions. In the ideal case, each node should have intervals whose lengths are equal to the size of its reachable set. Thus we adjust our computations to eliminate the effect of reachable set sizes by subtracting them from the actual values. The adjusted volume of a node is given as

$$adj_vol(u) = \prod_{j=1}^{i-1} (e_u^j - s_u^j - tcs(u))$$

where $tcs(u)$ is the size of the reachable set of u . Since computing the exact tcs values is not practical we use a linear-time estimation approach [12, 13] to approximate the tcs values.

- **Maximum of Adjusted Minimum Interval:** This is similar to the minimum interval, but using the maximum of the adjusted intervals. That is, we sort the children of each node u in decreasing order of

$$adj_mi(u) = \min_{j=1}^{i-1} \{(e_u^j - s_u^j - tcs(u))\}$$

With randomized or randomized pair traversal, the index construction time for GRAIL is $O(d(n + m))$, corresponding to the d traversals for the graph G . For the other traversal strategies, since we have to sort the children of each node the construction complexity is $O(d(n + m) + dn(o \log o))$ where o is the maximum out degree in the graph. Computing adjusted interval (in terms of tcs), does not increase the complexity as we used linear time reachability set estimation [12]. The space complexity of the GRAIL index is exactly $2dn = O(dn)$, since d intervals are kept per node.

3.3 Reachability Queries

To answer reachability queries between two nodes, u and v , GRAIL adopts a two-pronged approach. GRAIL first checks whether $L_v \not\subseteq L_u$. If so, we can immediately conclude that $u \not\rightsquigarrow v$, by Lemma 1. On the other hand, if $L_v \subseteq L_u$, nothing can be concluded immediately since we know that the index can have false positives, i.e., exceptions.

There are basically two ways of tackling exceptions. One approach is to explicitly maintain an *exception list* per node. Given node x , we denote by E_x , the list of exceptions involving node x , given as:

$$E_x = \{y : (x, y) \text{ is an exception, i.e., } L_y \subseteq L_x \text{ and } x \not\rightsquigarrow y\}$$

For example, for the DAG in Figure 6(b), we noted that there were 15 exceptions in total, as shown in Table 2. From the table, we can see that $E_2 = \{1, 4\}$, $E_4 = \{3, 7, 9\}$, and so on. If every node has an explicit exception list, then once we know that $L_v \subseteq L_u$, all we have to do is check if $v \in E_u$. If yes, then the pair (u, v) is an exception, and we return $u \not\rightsquigarrow v$. If no, then the containment is not an exception, and we answer $u \rightsquigarrow v$. We describe how to construct exceptions lists in section 4.

Unfortunately, keeping explicit exception lists per node adds significant overhead in terms of time and space, and further does not scale to very large graphs. Thus the default approach in GRAIL is to not maintain exceptions at all. Rather, GRAIL uses a “smart” DFS with recursive containment check based pruning to answer queries. This strategy does not require the computation of exception lists so its construction time and index size is linear in the graph.

Algorithm 2: GRAIL Query: Reachability Testing

```

Reachable( $u, v, G$ ):
1 if  $L_v \not\subseteq L_u$  then
2   return False //  $u \not\rightsquigarrow v$ 
3 else if use exception lists then
4   if  $v \in E_u$  then return False //  $u \not\rightsquigarrow v$ 
5   else return True //  $u \rightsquigarrow v$ 
6 else
7   // DFS with pruning
8   foreach  $c \in \text{Children}(u)$  such that  $L_v \subseteq L_c$  do
9     if Reachable( $c, v, G$ ) then
10    return True //  $u \rightsquigarrow v$ 
11 return False //  $u \not\rightsquigarrow v$ 

```

Algorithm 2 shows the pseudo-code for reachability testing in GRAIL. Line 1 tests whether $L_v \not\subseteq L_u$, and if so, returns false. Line 3 is applied only if exceptions lists are explicitly maintained, either complete or memoized (see Section 4): if $v \in E_u$, then GRAIL returns false, otherwise it returns true. Lines 7-10 implement the default strategy of

recursive DFS with pruning. If there exists a child c of u , that satisfies the condition that $L_v \subseteq L_c$, and we check and find that $c \rightsquigarrow v$, we can conclude that $u \rightsquigarrow v$, and GRAIL returns true (Line 9). Otherwise, if none of the children can reach v , then we conclude that $u \not\rightsquigarrow v$, and we return false in Line 10. As an example, let us consider the single interval index in Figure 6(b). Let $u = 2$, and let $v = 4$, and assume that we are not using exception lists. Since $L_4 = [1, 5] \subseteq [1, 9] = L_2$, we have to do a DFS to determine reachability. Both 3 and 5 are children of 2, but only 5 satisfies the condition that $L_4 = [1, 5] \subseteq [1, 8] = L_5$, we therefore check if 5 can reach 4. Applying the DFS recursion, we will check 6 and then, finally conclude that 5 cannot reach 4. Thus the condition in Line 8 fails, and we return false as the answer (Line 10), i.e., $2 \not\rightsquigarrow 4$.

Computational Complexity: It is easy to see that querying takes $O(d)$ time if $L_v \not\subseteq L_u$. If exception lists are to be used, and they are maintained in a hash table, then the check in Line 3 takes $O(1)$ time; otherwise, if the exceptions list is kept sorted, then the times is $O(\log(|E_u|))$. The default option is to perform DFS, but note that it is possible we may terminate early due to the containment based pruning. Thus the worst case complexity is $O(n + m)$ for the DFS, but in practice it can be much faster, depending on the topological level of u and depending on the effectiveness of pruning. Thus the query time ranges from $O(d)$ to $O(n + m)$.

4 Exception Lists

A naive approach for computing the exception lists is to precompute the transitive closure of each node, and then to check if each pair of nodes is an exception or not. Its overall complexity is $O(nm + n^2d)$. $O(nm)$ is for computing transitive closure and $O(n^2d)$ is the time spent for checking each pair. It also requires quadratic space. This is clearly not acceptable for large graphs. Instead we suggest another approach (see Algorithm 3) that has a better computational complexity and is faster in practice.

We categorize exceptions into two classes: i) If L_u contains v ², but none of the children of u contain v , then we call the exception between u and v a *direct exception*. ii) If at least one child of u contains v as an exception, then we call the exception between u and v an *indirect exception*. For example, in Figure 6(b) 3 is a direct exception for 4, but 1 is an indirect exception for 2, since 5 is a child of 2, and 1 is also an exception for 5. Table 2 shows the list of direct (denoted E^d) and indirect (denoted E^i) exceptions for the DAG in Figure 6(b).

² In this section, the phrases “ u contains v ”, “ L_u contains v ”, and “ u contains L_v ” are used interchangeably. All are equivalent to saying that L_u contains L_v .

The algorithm first computes the direct exceptions considering only the first dimension (or traversal). Then indirect exceptions from the first traversal are inferred via the method *ExtractIndirectExceptions*. After each node u has two lists, E_u^d and E_u^i , namely, the direct and indirect exception lists, these lists are updated after each additional traversal, via the method *ShrinkExceptionLists*.

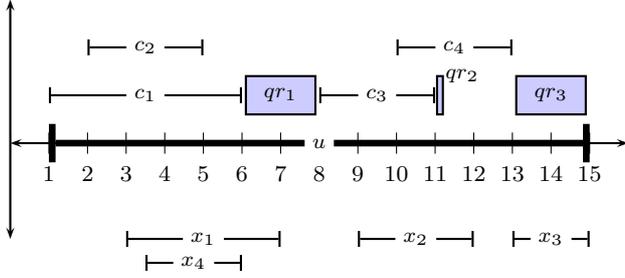


Fig. 7 Direct Exceptions: c_i denote children and x_i denote exceptions, for node u .

Direct Exceptions: Let us assume that $d = 1$, that is, each node has only one interval. GRAIL uses an *interval tree* [15] which keeps a list of intervals. Querying the interval tree for intervals intersecting a given range or interval of interest can be done in $O(\log n + K)$ time where n is the number of intervals and K is the number of results returned. Given the interval labeling, GRAIL constructs the exception lists for all nodes in the graph as described below. We illustrate how the algorithm works on the example in Figure 7, where we want to determine the exceptions for node u . c_i denote the children's intervals, whereas x_i denote the exceptions to be found.

The nodes are processed in increasing interval size order so that when a node u is being processed all the intervals that are contained in L_u are in the interval tree (line 7). To find the direct exceptions of u we look for the intervals that are not covered by any of the children of u using the interval tree. We first sort the children by increasing s_c^1 values and if there is a tie on these values they are ordered by decreasing e_c^1 values (line 9). It is clear that if an exception is contained completely within any one of the children intervals, it cannot be a direct exception. There are two cases where a direct exception can occur: i) The regions not covered by the union of the child intervals. We can see that these regions are $qr_1 = [6, 8]$ and $qr_3 = [13, 15]$ in our example. These regions are queried in the interval tree, and we find the nodes x_1 and x_3 as the direct exceptions of u . ii) Some exceptions might fall into the region of the union of two consecutive children as in the case of c_3 and c_4 in our example. x_2 is not covered by both of them but it is an exception for u so it should be a direct exception. To find such nodes, we query a very tiny interval at the end of the first child. In our

Algorithm 3: GRAIL : Exception List Extraction

```

ExceptionListExtraction( $G, L, d$ ):
1  $T \leftarrow$  Reverse Topological Order( $G$ )
2 Call FindDirectExceptions( $G, L, E^d$ )
3 Call ExtractIndirectExceptions( $G, L^1, E^d, E^i, C$ )
4 foreach  $i \leftarrow 2$  to  $d$  do
5    $\lfloor$  ShrinkExceptionLists( $G, L, E^d, E^i, C, dim$ )

FindDirectExceptions( $G, L, E^d$ ) :
6 Initiate an empty IntervalTree  $IT$ 
7  $T \leftarrow$  Sort by Increasing Interval Size
8 foreach node  $u \in T$  do
9    $c \leftarrow$  Sort children of  $u$  - with keys ( $s_{c_j}^1(\uparrow), e_{c_j}^1(\downarrow)$ )
10   $qr \leftarrow$  Find query regions by scanning  $c$ 
11  foreach region  $r_i \in qr$  do
12    Query  $IT$  for the range  $r_i$ 
13    Add resulting nodes to  $E_u^d$ 
14  Add the interval  $[s_{T_i}^1, e_{T_i}^1]$  to  $IT$ 

ExtractIndirectExceptions( $G, L, E^d, E^i$ ):
15 foreach node  $u$  in bottom up order do
16   foreach  $x \in (\bigcup_{j=1}^k E_{c_j})$  where  $c_j$  is a child of  $u$  do
17     if  $x$  is not a proper descendant of any child  $c$  then
18       Add  $x$  to  $E_u^i$ 
19        $C_u^x \leftarrow$  Number of children  $u$  that contain  $x$ 

ShrinkExceptionLists( $G, L, E^d, E^i, C, dim$ ):
20 foreach node  $u$  do
21   foreach exception  $x \in E_u^d$  do
22     if  $L_u^{dim} \not\supseteq L_x^{dim}$  then
23       Remove  $x$  from  $E_u^d$ 
24       foreach parent  $p$  of  $u$  where  $x \in E_p^i$  do
25          $C_p^x \leftarrow C_p^x - 1$ 
26         if  $C_p^x = 0$  then
27            $\lfloor$  Move  $x$  from  $E_p^i$  to  $E_p^d$ 

```

case, we query the region $[11, 11 + \delta]$ where δ is a constant less than 1. Note that we can skip the end regions of the children that are contained by other children such as in the case of c_1 and c_2 . For example if we had queried the region $[5, 5 + \delta]$ it would return x_4 which is indeed covered by c_1 so it is not a direct exception. Thus scanning the intervals in that particular order is sufficient to find out such query regions (line 10). Next, we query each region in the interval tree and add the resulting intervals to the direct exception list of the node being processed (line 11-13). Finally that node is added to the interval tree so that it can be found if it is an exception for other nodes. The complexity of this method is $O(no(\log n + p))$ where o is the maximum out degree and p is maximum number of exceptions returned by a query. This is because at each node we can query at most o times, each of which runs in $O(\log n + p)$.

Indirect Exceptions: Given that we have the list of direct exceptions E_u^d for each node, the construction of the indirect exceptions (E_u^i) proceeds in a bottom up manner from

the leaves to the roots. Let $E_{c_j} = E_{c_j}^d \cup E_{c_j}^i$ denote the list of direct or indirect exceptions for a child node c_j of u . To compute E_u^i , for each exception $x \in E_{c_j}$ we check if there exists another child c that can reach x (line 17). This reachability check is easy since by that time we know the exception list of c (i.e. E_c). A child c_k reaches to x if $L_x^1 \subseteq L_c^1$ and $x \notin E_c$. On the other hand, if there is no such children, then x must be an indirect exception for u , and we add it to E_u (line 24). We also keep a counter for each exception x of the node u , C_u^x , which records the number of children of u that have x as exception. These counters are utilized while incorporating the other traversals. For example, consider node 2 in Figure 6(b). Assume we have already computed the exception lists for each of its children, $E_3 = E_3^d \cup E_3^i = \emptyset$, and $E_5 = E_5^d \cup E_5^i = \{1, 3, 4, 7, 8\}$. We find that for each $x \in E_5$, nodes 1, and 4 fail the test with respect to E_3 , since $L_1 \not\subseteq L_3$, and $L_4 \not\subseteq L_3$, therefore $E_2^i = \{1, 4\}$, as illustrated in Table 2. The complexity of this step is $O(neo^2)$ where o is the maximum out degree and e is maximum number of exceptions a child node has. This is because every exception of a child node c_i of a node u has to be checked in the exception list of every other child c_j .

Multiple Intervals: To find the exceptions when $d > 1$, GRAIL first computes the direct and indirect exceptions from the first traversal, as described above. We maintain the counters C_u^x to keep track of the number of children of u which has x in their exception list. It is worth noting that an exception can be removed only when C_u^x becomes zero. For computing the remaining exceptions after the i -th traversal, GRAIL processes the nodes in a bottom up order. For every direct exception $x \in E_u^d$, remove x from the direct exception list if x is not an exception for u for the i -th dimension (line 24), and further, decrement the counter for x in the indirect exceptions list E_p^i for each parent p of u (line 25). Also, if after decrementing, the counter for any indirect exception x becomes zero, then move x to the direct exception list E_p^d of the parent p , provided $L_x \subseteq L_p$ (line 26-27). In this way all exceptions can be found out for the i -th dimension or traversal. The complexity of this step is $O(n\phi e)$ as every direct exception is checked in constant time and the counters for the parent nodes are decremented if needed in $O(\phi)$ time, where ϕ is the maximum in-degree.

Therefore the overall complexity of exception list maintenance is $O(neo(\log n + e + o))$ with the reasonable assumption of maximum in-degree is in the order of maximum out-degree. It is expected to be better than $O(nm)$ as long as e is not close to n . In practice the bottleneck of this algorithm is the second step as there are many exceptions just after the first traversals. However the number of exceptions reduces drastically with the addition of second traversal, therefore the *shrinking exception lists* step runs very fast although it is called multiple times. It is possible to speed up

the algorithm if direct exceptions can be found by considering the first two traversals; we plan to explore this idea in the future.

5 GRAIL Optimizations

In this section we discuss several optimizations over the basic GRAIL approach. The proposed methods can be very effective in speeding up the query times.

5.1 Topological Level Filter

Once the graph is transformed into a directed acyclic graph, the nodes can be divided into levels so that the leaf nodes are placed into the first level and the interior nodes are placed in a higher level than all of their children. Therefore the level of u is

$$l_u = \begin{cases} 1 & \text{if } u \text{ is leaf} \\ 1 + \max_{v \in \text{children}(u)} \{l_v\} & \text{otherwise} \end{cases}$$

There is no possibility for a node u at level l_u to reach another node v at level l_v if $l_v \leq l_u$. Thus the topological level filter can be utilized during GRAIL's recursive DFS to prune unreachable node pairs. Even though this filter is very simple, our experimental results show that it is very effective in improving query times.

5.2 Different Search Strategies

We perform a depth-first search from the source node to the target node in algorithm 2 for reachability testing. However it is not always the best strategy. Breadth-first search (BFS) guarantees to find the shortest path with the cost of using more memory, though this is not a limitation in our case because the memory usage is limited by the graph size. In average, BFS is expected to be faster considering the cases where the target node is close to the source node that has a large reachability set. However this difference might not be apparent in sparse graphs where DFS is also very fast. In negative queries they both have to exhaust the whole reachable set to conclude non-reachability so they are expected to have similar performance.

Besides the standard graph search strategies of DFS and BFS, we also consider bidirectional breadth-first search (BBFS) which has a few advantages over BFS. In bidirectional search two BFS are started simultaneously utilizing separate queues for the query $u \rightsquigarrow v$. The first adds all the children $c_i \in \text{children}(u)$ to a BFS queue Q_C of children, whereas the second adds all the parents $p_i \in \text{inparents}(v)$ to another BFS queue Q_P of parents. Then BBFS alternately extracts one node from Q_C pushing its children back to Q_C and extracts

one node from Q_P pushing its parents back to Q_P . BBFS stops positively if any newly added node c_i is already in Q_P or any newly added node p_j is already in Q_C , and BBFS returns $u \not\rightsquigarrow v$ the moment either of the queues becomes empty. Furthermore BBFS prunes the branch of c_i by not adding into Q_C if $L_v \not\subseteq L_{c_i}$. Symmetrically p_j is not added to Q_P if $L_{p_j} \not\subseteq L_u$. The first advantage of bidirectional search is that it can answer positive queries faster when the average degree of the graph is high. In negative queries, the worst case scenario is that bidirectional search takes twice as much time as a BFS would take; however there are substantial amount of cases where bidirectional search is much faster than BFS. For instance, consider the case in which the source node has a large reachable set and the target has a small number of ancestors. In this scenario bidirectional search terminates after visiting the ancestors of the target node which is significantly smaller than the search space of a standard BFS which has to exhaust the large reachable set of the source node.

Our experiments show that BBFS can be a very effective strategy, especially in combination with the level filter. In fact, we use the three different search strategies DFS, BFS, and BBFS (with and without the level filter) as the baseline methods to compare against GRAIL as well as other indexing methods.

5.3 Positive Cut Filter

Recall that GRAIL use direct DAG interval labeling to guarantee that if $u \rightsquigarrow v$ then $L_v \subseteq L_u$. However, it can have false positives, i.e., for a pair of nodes with $L_v \subseteq L_u$ it may be that case that $u \not\rightsquigarrow v$. Recall also that other interval labeling methods try to label only a subtree (or a set of paths), and for paths restricted to the tree edges (which typically constitute only a small fraction of the DAG) reachability can be guaranteed. It is possible for GRAIL to also make use of the subtree edges, via a new pruning method we call *positive cut filter*.

Notice that from each of the d traversals in GRAIL we can easily keep track of additional interval labels restricted to only a (spanning) subtree of the input DAG. For example, in Figure 8(a) a min-post labeling which is produced by a traversal of GRAIL is shown. We call these set of labels the *exterior labels*, denoted as EL_u for node u . Figure 8(b) is the min-post labeling of the corresponding spanning tree, which is composed of the solid edges (the dashed edges are the non-tree edges). We call this set of labels the *interior labels*, denoted IL_u for node u . The figure shows only one label, but for d traversals, there are a set of d exterior and interior labels.

Given a query $u \overset{?}{\rightsquigarrow} v$, we can first check whether $IL_v \subseteq IL_u$. If yes, then we are guaranteed that $u \rightsquigarrow v$, and we can immediately return a positive answer. On the other hand, if

the interior check fails then we default to the basic GRAIL strategy of first checking the exterior set and then recursive search. That is, we next check if $EL_v \subseteq EL_u$. If yes, then we immediately return a negative answer $u \not\rightsquigarrow v$. If no, then we do a recursive DFS (or BBFS, etc.). This positive cut filter strategy allows GRAIL to substantially speed up the positive queries. In effect, GRAIL captures the core benefits of interval labels in both directions, allowing it to cut-off the search for guaranteed positive pairs (over the tree edges, the interior labels), and also allowing it to prune the non-reachable pairs (using the exterior labels). Figure 8(c) illustrates this positive cut filter via a Venn diagram over $d = 2$ traversals. The solid black line (rounded outer rectangle) demarcates the universal set of all possible ordered pairs of nodes (all possible queries) in the graph. The solid red line (inner rounded rectangle) denotes the transitive closure, i.e., the set of all reachable pairs of nodes. Each GRAIL traversal is associated with two label sets. The exterior set defined by the min-post labeling of the graph is a superset of the reachable set (denoted by the outer blue and green hatched regions). The interior set defined by the min-post labeling of the tree is a subset of the reachable set (inner solid blue and green regions). The intersection of the exterior sets (the cross hatched region) defines the ordered pairs for whom GRAIL is not sure whether they are reachable or not. The union of the interior sets defines the set of ordered pairs which GRAIL guarantees to be reachable. If the two checks over the interior or exterior labels fail, then GRAIL must perform a recursive search to determine reachability; this corresponds to the visible cross hatched region, which excludes the interior regions. The recursive search continues as long as it is in the crosshatched region. GRAIL prunes a branch whenever it exits the region and it halts with a positive result whenever it reaches a pair inside the union of interior sets.

Experimental results in section 6 show that in many cases using positive cut filter substantially improves the GRAIL query performance. It is especially very effective on positive queries. Note that it does not deteriorate the construction time as the exterior and interior labelings can be created simultaneously. Furthermore, instead of doubling the index size, the increase is only a factor of 1.5. This is because, the two intervals (interior and exterior) per node can be stored by three integers. The interval label for node u for the i^{th} dimension/traversal is given as $L_u^i = [s_u^i, m_u^i, e_u^i]$, where $EL_u^i = [s_u^i, e_u^i]$ and $IL_u^i = [m_u^i, e_u^i]$ are the exterior and interior intervals respectively. For example, in Figure 8, the exterior label $EL_5 = [1, 8]$ and interior label $IL_5 = [1, 7]$ for node 5 are stored as a single label $L_5 = [1, 7, 8]$. Algorithm 4 shows the GRAIL querying method using the positive cut filter. For clarity, the method is shown using separate EL_u and IL_u labels, but in reality only L_u is maintained.

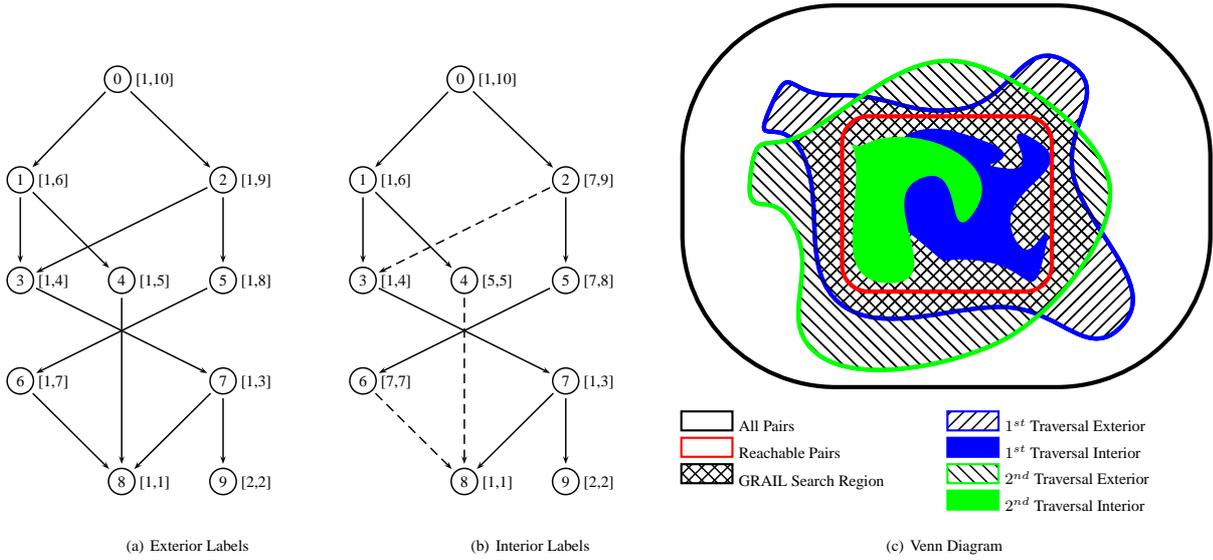


Fig. 8 Venn Diagram of Pairs wrt. Reachability

Algorithm 4: GRAIL+ Query: Reachability Testing

```

Reachable( $u, v, G$ ):
1 if  $IL_v \subseteq IL_u$  then
2   return True //  $u \rightsquigarrow v$ 
3 else if  $EL_v \not\subseteq EL_u$  then
4   return False //  $u \not\rightsquigarrow v$ 
5 else
6   // DFS with pruning
7   foreach  $c \in Children(u)$  such that  $EL_v \subseteq EL_c$  do
8     if Reachable( $c, v, G$ ) then
9       return True //  $u \rightsquigarrow v$ 
9   return False //  $u \not\rightsquigarrow v$ 

```

6 Experiments

We conducted extensive experiments to compare our algorithm with the best existing methods. All experiments are performed in a machine with x86_64 Dual Core AMD Opteron Processor 870, which has 8 processors and 32G ram. We compared our algorithm with several baseline search methods like DFS (depth-first), BFS (breadth-first) and BBFS (bidirectional BFS), as well as state-of-the-art reachability indexes, such as those based on interval labeling: Interval (INT) [4], GRIPP [31], and Dual Labeling (Dual) [32], based on path decomposition: PathTree (PT) [23], and variants of 2HOP indexing: HLSS [19] and 3HOP [22]. The code for these methods was obtained from the authors, though in some cases, the original code had been reimplemented by later researchers. We implemented GRIPP in-house.

We evaluate all methods in terms of the query time, index size, and index construction time. All query times re-

ported below are aggregate times for 100K queries. We generate 100K random query pairs, and issue the same set of queries to all methods. In the tables below, we use the notation $-(t)$, and $-(m)$, to note that the given method exceeds the allocated time (10M milliseconds (ms) for small sparse, and 20M ms – about 5.5hrs – for all other graphs; $M \equiv million$) or memory limits (32GB RAM; i.e., the method aborts with a bad-alloc error). All times reported are averages over three runs.

6.1 Datasets

We used a variety of real graph datasets, both small and large, as well as large synthetic ones, as described below. We present these graphs in 4 tables each of which lists the following properties: node size, edge size, average degree, average clustering coefficient (CC) and effective diameter (ED). As we are dealing with directed graphs, the average degree of a graph is equal to the ratio of the number of edges to the number of nodes. We used SNAP [26] software to compute the values of clustering coefficient and effective diameter.

Small-Sparse: These are small, real graphs, with average degree less than 1.2, taken from [23], and listed in Table 3. *xmark* and *nasa* are XML documents, and *amaze* and *kegg* are metabolic networks, first used in [31]. Others were collected from BioCyc (biocyc.org), a collection of pathway and genome databases. *amaze* and *kegg* have a different structure, in that they have a central node which has a very large in-degree and out-degree. Therefore these two

| Dataset | Nodes | Edges | Avg Deg | CC | ED |
|---------|-------|-------|---------|------|-------|
| agrocyc | 12684 | 13657 | 1.06 | 0.33 | 11.98 |
| amaze | 3710 | 3947 | 1.06 | 0.31 | 3.79 |
| anthra | 12499 | 13327 | 1.07 | 0.32 | 11.89 |
| ecoo | 12620 | 13575 | 1.08 | 0.32 | 12.23 |
| human | 38811 | 39816 | 1.03 | 0.39 | 11.74 |
| kegg | 3617 | 4395 | 1.22 | 0.26 | 4.08 |
| mtbrv | 9602 | 10438 | 1.09 | 0.30 | 11.90 |
| nasa | 5605 | 6538 | 1.17 | 0.07 | 14.22 |
| vchocyc | 9491 | 10345 | 1.09 | 0.07 | 11.91 |
| xmark | 6080 | 7051 | 1.16 | 0.00 | 11.34 |

Table 3 Small & Sparse Real Graphs

| Dataset | Nodes | Edges | Avg Deg | CC | ED |
|----------|-------|-------|---------|------|-------|
| arxiv | 6000 | 66707 | 11.12 | 0.35 | 5.48 |
| citeseer | 10720 | 44258 | 4.13 | 0.28 | 8.36 |
| go | 6793 | 13361 | 1.97 | 0.07 | 10.92 |
| pubmed | 9000 | 40028 | 4.45 | 0.10 | 6.32 |
| yago | 6642 | 42392 | 6.38 | 0.24 | 6.57 |

Table 4 Small & Dense Real Graphs

| Dataset | Nodes | Edges | AD | CC | ED |
|-------------|----------|----------|------|------|------|
| citeseer | 340945 | 312282 | 0.92 | 0 | 5.7 |
| uniprot22m | 1595444 | 1595442 | 1.00 | 0 | 3.3 |
| uniprot100m | 16087295 | 16087293 | 1.00 | 0 | 4.1 |
| uniprot150m | 25037600 | 25037598 | 1.00 | 0 | 4.4 |
| citeseerx | 6540399 | 15011259 | 2.30 | 0.06 | 8.4 |
| cit-patents | 3774768 | 16518947 | 4.38 | 0.09 | 10.5 |
| go-uniprot | 6967956 | 34770235 | 4.99 | 0 | 4.8 |

Table 5 Large Real Graphs: Sparse and Dense

| Dataset | Nodes | Edges | Avg Deg | CC | ED |
|------------|-------|-------|---------|----|-------|
| rand10m2x | 10M | 20M | 2 | 0 | 13.61 |
| rand10m5x | 10M | 50M | 5 | 0 | 8.75 |
| rand10m10x | 10M | 100M | 10 | 0 | 6.86 |
| rand100m2x | 100M | 200M | 2 | 0 | ? |
| rand100m5x | 100M | 500M | 5 | 0 | ? |

Table 6 Large Synthetic Graphs

graphs have smaller effective diameter although the number of reachable pairs is very high compared to other graphs.

Small-Dense: These are small, dense real-world graphs taken from [22] (see Table 4). *arxiv*³, *citeseer*⁴, and *pubmed*⁵ are all citation graph datasets. GO is a subset of the Gene Ontology⁶ graph, and *yago* is a subset of the semantic knowledge database YAGO⁷. Among these graphs, *go* is significantly different from others with its low average degree and clustering coefficient and high diameter. Since it is a graph representing a taxonomy, triangles in it implies the existence of redundant edges (e.g. if node *a* has *b* and *c* as neighbors and *b* has *c*, the edge between *a* and *c* is redundant in a taxonomy due transitivity). Therefore it is natural to have low clustering coefficient.

³ arxiv.org⁴ citeseer.ist.psu.edu⁵ www.pubmedcentral.nih.gov⁶ www.geneontology.org⁷ www.mpi-inf.mpg.de/suchanek/downloads/yago

Large-Real: To evaluate the scalability of GRAIL on real datasets, we collected 7 new datasets which have previously not been used by existing methods (see Table 5). *citeseer*, *citeseerx* and *cit-patents* are citations networks in which non-leaf nodes have 10 to 30 outgoing edges on average. However *citeseer* is very sparse because of data incompleteness. *citeseerx*⁸ is the complete citation graph as of March 2010. *cit-patents*⁹ includes all citations within patents granted in the US between 1975 and 1999. *go-uniprot* is the joint graph of Gene Ontology terms and the annotations file from the UniProt¹⁰ database, the universal protein resource. Gene ontology is a directed acyclic graph of size around 30K, where each node is a term. UniProt annotations consist of connections between the gene products in the UniProt database and the terms in the ontology. UniProt annotations file has around 7 million gene products annotated by 56 million annotations. The remaining uniprot datasets are obtained from the RDF graph of UniProt. *uniprot22m* is the subset of the complete RDF graph which has 22 million triples, and similarly *uniprot100m* and *uniprot150m* are obtained from 100 million and 150 million triples, respectively. These are some of the largest directed acyclic graphs ever considered for reachability testing. The reasons why these graphs have zero clustering coefficient vary. *uniprot* graphs have distinctive topology of having many roots that are connected to a single sink node via short paths. It is kind of a tree whose edges are reversed. Therefore there exists no triangles. For *go-uniprot* the reason is its being a taxonomy and for *citeseer* it is data incompleteness.

Large-Synthetic: To test the scalability with different density setting, we generated random DAGs, with 10M and 100M nodes, and with average degrees of 2, 5, and 10 (see Table 6). We first randomly select an ordering of the nodes which corresponds to the topological order of the final dag. Then for the specified number of edges, we randomly pick two nodes and connect them with an edge from the lower to higher ranked node. Since we randomly select the neighbors of the nodes, it is very unlikely to choose two nodes which are already connected among millions of nodes. Therefore clustering coefficient is zero in these graphs. SPAN was not able to scale to 50 million nodes for computing effective diameter so we show them with question mark.

6.2 Performance Comparison with Graph Indexes

Before studying the sensitivity of GRAIL to various parameters, and before evaluating the effectiveness of the various

⁸ citeseerx.ist.psu.edu⁹ snap.stanford.edu/data¹⁰ www.uniprot.org

| Dataset | GRAIL | HLSS | INT | Dual | PT | 3HOP | GRIPP |
|---------|-------|------------|----------|-----------|--------|------------|--------------|
| agrocyc | 18.33 | 12832.40 | 5188.56 | 12092.73 | 269.71 | 108761.47 | 4.72 |
| amaze | 5.67 | 684259.17 | 3074.14 | 4621.47 | 839.31 | 3387549.83 | 2.73 |
| anthra | 17.04 | 11935.63 | 4916.71 | 11375.35 | 259.08 | 95874.58 | 4.64 |
| ecoo | 20.26 | 12958.68 | 5058.26 | 11892.17 | 269.01 | 110114.53 | 4.76 |
| human | 52.33 | 129599.17 | 45352.85 | 128064.60 | 831.37 | -(m) | 16.02 |
| kegg | 5.83 | 1112604.10 | 3466.63 | 5381.69 | 968.50 | 4044018.00 | 2.56 |
| mtbrv | 14.46 | 3630.32 | 2517.65 | 3130.82 | 204.91 | 65721.12 | 3.53 |
| nasa | 9.29 | 1748.71 | 805.65 | 1037.23 | 124.67 | 50839.88 | 2.60 |
| vchocyc | 15.04 | 4840.56 | 2543.42 | 4350.59 | 201.27 | 66401.13 | 3.47 |
| xmark | 11.31 | 68687.95 | 1512.11 | 1788.98 | 260.53 | 197554.38 | 2.60 |

Table 7 Small Sparse Graphs: Construction Time (ms)

| Datasets | GRAIL | HLSS | INT | Dual | PT | 3HOP | GRIPP |
|----------|--------|--------|--------------|-------------|--------|-------|--------------|
| agrocyc | 88788 | 40097 | 27100 | 58552 | 39027 | 38631 | 27314 |
| amaze | 25970 | 17110 | 10356 | 433345 | 12701 | 10221 | 7894 |
| anthra | 87493 | 33532 | 26310 | 37378 | 38250 | 38112 | 26654 |
| ecoo | 88340 | 34285 | 26986 | 58290 | 38863 | 38449 | 27150 |
| human | 271677 | 109962 | 79272 | 54678 | 117396 | -(m) | 79632 |
| kegg | 25319 | 17427 | 10242 | 504K | 12554 | 10141 | 8790 |
| mtbrv | 67214 | 30491 | 20576 | 41689 | 29627 | 29324 | 20876 |
| nasa | 39235 | 20976 | 18324 | 5307 | 21894 | 18913 | 13076 |
| vchocyc | 66437 | 30182 | 20366 | 26330 | 29310 | 28988 | 20690 |
| xmark | 42560 | 23814 | 16474 | 16434 | 20596 | 21161 | 14102 |

Table 8 Small Sparse Graphs: Index Size (Num. Entries)

| Dataset | GRAIL | HLSS | INT | Dual | PT | 3HOP | GRIPP | DFS-L | BBFS-L | #Pos.Q |
|---------|-------|--------|--------|-------|--------------|--------|--------|---------|--------|--------|
| agrocyc | 20.83 | 71.10 | 162.82 | 70.93 | 8.42 | 122.63 | 70.95 | 34.26 | 142.31 | 133 |
| amaze | 16.46 | 99.20 | 103.05 | 62.91 | 7.08 | 50.51 | 73.21 | 693.76 | 196.66 | 17259 |
| anthra | 20.61 | 70.66 | 160.67 | 65.50 | 8.26 | 117.15 | 65.77 | 31.06 | 143.91 | 97 |
| ecoo | 20.74 | 74.35 | 161.41 | 68.22 | 8.07 | 120.40 | 74.94 | 35.03 | 143.47 | 129 |
| human | 22.69 | 80.36 | 237.73 | 76.57 | 16.93 | -(m) | 57.82 | 32.31 | 165.34 | 12 |
| kegg | 17.95 | 106.88 | 104.06 | 64.11 | 7.25 | 52.13 | 72.27 | 1086.81 | 265.52 | 20133 |
| mtbrv | 18.11 | 78.23 | 145.61 | 69.46 | 7.42 | 106.21 | 80.14 | 34.56 | 128.61 | 175 |
| nasa | 19.93 | 90.78 | 129.52 | 63.20 | 7.85 | 76.74 | 165.40 | 92.37 | 135.26 | 562 |
| vchocyc | 17.86 | 68.43 | 145.03 | 62.64 | 7.22 | 105.60 | 77.03 | 34.27 | 127.60 | 169 |
| xmark | 30.20 | 91.00 | 122.24 | 76.02 | 7.50 | 100.40 | 89.79 | 321.22 | 164.66 | 1482 |

Table 9 Small Sparse Graphs: Query Time (ms)

| Dataset | GRAIL | HLSS | INT | Dual | PT | 3HOP | GRIPP |
|----------|--------------|-----------|----------|-----------|---------|------------|--------------|
| arXiv | 29.71 | -(t) | 20097.87 | 386701.59 | 9770.60 | 8854231.00 | 22.57 |
| citeseer | 46.82 | 117865.33 | 6712.07 | 30213.80 | 710.29 | 110783.48 | 115.28 |
| go | 20.05 | 68030.83 | 1122.11 | 7491.06 | 219.43 | 29578.30 | 4.88 |
| pubmed | 35.05 | 142844.19 | 5043.40 | 25832.03 | 768.78 | 293161.20 | 181.80 |
| yago | 24.60 | 28259.51 | 2593.61 | 5329.56 | 506.53 | 30602.29 | 44.95 |

Table 10 Small Dense Graphs: Construction Time (ms)

| Datasets | GRAIL | HLSS | INT | Dual | PT | 3HOP | GRIPP |
|----------|--------|--------|--------|----------|--------|--------------|--------------|
| arxiv | 60000 | -(t) | 145668 | 14057239 | 86855 | 47472 | 133414 |
| citeseer | 107200 | 114088 | 142632 | 30615323 | 91820 | 51035 | 88516 |
| go | 67930 | 60287 | 40644 | 11000662 | 37729 | 27764 | 26722 |
| pubmed | 90000 | 102946 | 181260 | 15040251 | 107915 | 54531 | 80056 |
| yago | 66420 | 57003 | 57390 | 4371065 | 39181 | 27038 | 84784 |

Table 11 Small Dense Graphs: Index Size (Num. Entries)

| Dataset | GRAIL | HLSS | INT | Dual | PT | 3HOP | GRIPP | DFS-L | BBFS-L | #Pos.Q |
|----------|--------|--------|--------|-------|--------------|--------|------------|---------|--------|--------|
| arXiv | 380.94 | -(t) | 269.64 | 80.42 | 24.73 | 355.36 | 4041302.42 | 6599.51 | 665.56 | 15459 |
| citeseer | 64.04 | 327.45 | 222.12 | 80.58 | 24.61 | 184.21 | 9064.53 | 203.80 | 256.54 | 388 |
| go | 30.79 | 274.64 | 148.14 | 77.01 | 11.67 | 105.23 | 2015.98 | 106.07 | 164.00 | 241 |
| pubmed | 70.07 | 307.48 | 244.04 | 76.79 | 21.85 | 179.81 | 7551.03 | 202.00 | 190.14 | 690 |
| yago | 19.64 | 256.72 | 174.86 | 64.91 | 14.18 | 117.08 | 1122.39 | 29.86 | 161.17 | 171 |

Table 12 Small Dense Graphs: Query Time (ms)

| Dataset | Construction Time (ms) | | Index Size | |
|-------------|------------------------|-----------------|------------|-----------------|
| | GRAIL | GRIPP | GRAIL | GRIPP |
| cit-patents | 64676.35 | 41453.16 | 60396288 | 33037894 |
| citeseer | 5791.04 | 234166.72 | 11103152 | 624564 |
| citeseerx | 58746.74 | 1174080 | 104646416 | 30022520 |
| go-uniprot | 89821.39 | 41851.86 | 111487296 | 69540470 |
| uniprot22m | 2767.63 | 2642.50 | 11168108 | 3190884 |
| uniprot100m | 31919.36 | 33863.25 | 112611065 | 32174586 |
| uniprot150m | 49355.86 | 59986.65 | 175263200 | 50075196 |

Table 13 Large Real Graphs: Construction Time (ms) and Index Size

| Dataset | GRAIL | GRIPP | DFS-L | BFS-L | BBFS-L | #Pos.Q |
|-------------|----------------|----------|---------------|---------|---------|--------|
| cit-patents | 1369.02 | 31177982 | 25774.05 | 3407.33 | 5064.83 | 39 |
| citeseer | 27.04 | 362.91 | 25.749 | 86.10 | 157.49 | 0 |
| citeseerx | 113.43 | 27932.54 | 119260.48 | 4372.09 | 4162.48 | 239 |
| go-uniprot | 31.06 | 4864.87 | 37.78 | 157.56 | 271.21 | 0 |
| uniprot22m | 22.65 | 279.94 | 24.52 | 120.32 | 174.13 | 0 |
| uniprot100m | 59.35 | 496.59 | 63.47 | 155.44 | 220.87 | 0 |
| uniprot150m | 52.66 | 530.61 | 71.55 | 145.55 | 238.51 | 0 |

Table 14 Large Real Graphs: Query Times (ms)

| Size | Deg. | Construction Time (ms) | | | Index Size | | |
|----------|------|------------------------|-----------------|-----------------|------------|-------------|-------------|
| | | GRAIL | GRAIL* | GRIPP | GRAIL | GRAIL* | GRIPP |
| rand10m | 2 | 52782.9 | 53368.0 | 585639.3 | 70M | 40M | 40M |
| | 5 | 259631.1 | 244291.0 | 48338.5 | 160M | 100M | 100M |
| | 10 | 433898.4 | 412623.3 | 80471.1 | 160M | 100M | 200M |
| rand100m | 2 | 336722.2 | 313333.4 | 9920415.0 | 350M | 200M | 200M |
| | 5 | 1661943.6 | 1458824.1 | 292427.0 | 800M | 500M | 500M |

Table 15 Large Synthetic Graphs: Construction Time (ms) and Index Size

| Size | Deg. | GRAIL | GRAIL* | GRIPP | DFS-L | BFS-L | BBFS-L | #Pos.Q |
|----------|------|--------------|------------------|------------|------------|-----------|-----------|--------|
| rand10m | 2 | 174.8 | 259.8 | 1955.9 | 326.3 | 266.9 | 476.4 | 0 |
| | 5 | 6102.9 | 5764.5 | 54212516.0 | 75732.8 | 13011.2 | 17449.6 | 17 |
| | 10 | 1523248.0 | 1417776.5 | -(t) | 35372076.0 | 2815736.7 | 1734676.1 | 5522 |
| rand100m | 2 | 229.0 | 326.7 | 2646.6 | 532.0 | 343.0 | 603.4 | 0 |
| | 5 | 9281.9 | 7241.1 | -(t) | 130050.5 | 16229.6 | 23115.3 | 6 |

Table 16 Large Synthetic Graphs: Query Times (ms)

search strategies and optimizations, we compare GRAIL’s performance with existing state-of-the-art graph reachability indexes, as well as the baseline search methods.

Based on our experiments (outlined later) the default strategy for GRAIL is to use randomized pairs traversal to construct the labeling (with $d = 2$ for small-sparse graphs, $d = 3$ for small-dense graphs and $d = 5$ for large graphs). Furthermore, GRAIL uses the topological level and positive cut filters, and does not maintain exception lists. It uses DFS with recursive pruning as the default search strategy. We denote by GRAIL* the basic approach, i.e., randomized traversals, with no level or positive cut filters, and without exception lists (using DFS for recursive search).

6.2.1 Small Real Datasets: Sparse and Dense

Tables 7, 9, and 8 show the index construction time, query time, and index size for the small, sparse, real datasets. Tables 10, 12, and 11 give the corresponding values for the

small, dense, real datasets. The last column in Tables 9 and 12 shows the number of reachable query node-pairs, called *positive queries*, out of the 100K test queries; the query node-pairs are sampled randomly from the graphs and the small counts are reflective of the sparsity of the graphs. The best results are shown in bold.

Small Sparse Graphs: On the sparse datasets (see Table 7), GRIPP has the smallest construction time among all indexing methods, though GRAIL (with $d = 2$ traversals) is a close second. PathTree is typically an order of magnitude slower than GRAIL, and the other methods are even more expensive (several orders of magnitude slower). 3HOP could not run on human, since it exhausted the memory (denoted $-(m)$). INT and GRIPP have the smallest index size (see Table 8). The other methods have comparable index sizes, though GRAIL has the largest number of entries (about 3 times larger than INT).

In terms of query times on the small sparse graphs (see Table 9), PathTree is the best, with GRAIL in the second place, typically being 2-4 times slower. On these small sparse datasets, it is worth noting that, with few exceptions, DFS-L (DFS with level filter) is faster than all indexing methods other than PathTree and GRAIL. The exceptions are *amaze*, *kegg*, and *xmark*, where DFS-L is not as effective; on these graphs BBFS-L is able to improve the query times by a factor of 2-4. The reason why DFS-L suffers in these graphs is the graph topology. These graphs have a central node which has a very high in degree and out degree, therefore in most of the cases a DFS has to scan the children of that central node to arrive the target. Whereas BBFS can terminate the search without scanning the children of the central node. Given the fact that the baseline graph search methods have no construction time or indexing size overhead, they are quite attractive for these small datasets. A more detailed comparison among DFS, BFS, and BBFS (with and without the level filter), appears in Tables 18, 19, and 20.

Small Dense Graphs: On the small dense datasets, GRAIL (with $d = 3$) has the smallest index construction times (see Table 10), being up to 6 times faster than the closest rival GRIPP. Other methods are orders of magnitude slower, and HLSS could not run on *arxiv*. The index size (see Table 11) is comparable for all methods except DUAL. 3HOP has the smallest index.

On these small dense graphs PathTree is still the fastest indexing method, with GRAIL in the second place, being typically 2-3 times slower than PathTree (see Table 12). Once again, DFS-L is comparable to the other indexing methods. On *arxiv* BBFS-L is able to remedy the shortcomings of DFS-L. It is also worth noting that GRIPP’s query performance acutely deteriorates as the average degree increases. One can conclude that for the small dense graphs, while pure search based methods are quite acceptable, indexing does deliver significant benefits in terms of query time performance.

6.2.2 Large Datasets: Real and Synthetic

Large Real Graphs: Table 13 shows the construction time, and index size on the large real graphs. On these graphs, with the exception of GRAIL and GRIPP, none of the other indexing methods were able to run. PathTree aborted with a memory limit error $(-m)$ on *cit-patents* and *citeseerx*, and it exceeded the 20M ms time limit $(-t)$ for the other datasets. It was able to run only on *citeseer* data (130406 ms for construction, and the index size was 2360732 entries). While GRIPP’s index size is smaller than GRAIL’s, its construction time can be up to 40 times slower than GRAIL.

The main index construction of GRIPP is linear on the number of edges, however by default it also extracts a list of special nodes (i.e. stop nodes) to speed up the querying. Therefore in some cases such as in *citeseer* its construction time does not reflect its linear computational complexity.

From the query times (see Table 14), we can observe that GRAIL can easily scale to large datasets (the only limitation being that it does not yet process disk-resident graphs). We can see that GRAIL outperforms GRIPP by orders of magnitude, and it is faster than BFS-L (the best among the search-based methods) by 3-40 times on the denser graphs: *go-uniprot*, *cit-patents*, and *citeseerx*. On the other datasets, that are very sparse, GRAIL is still the winner while pure DFS is the close second. On these graphs we use $d = 2$ traversals for GRAIL. It is worth noting that PathTree took 47.4 ms for querying on the sparser *citeseer* data, which is still 2 times slower than GRAIL. It is also interesting that GRIPP is still 10 times slower than GRAIL however it uses one third of index size. If memory is an issue, one should definitely choose with DFS with a simple topological filter on such sparse graphs.

Large Synthetic Graphs: We also tested the scalability of GRAIL on the large synthetic graphs, which have 10M and 50M nodes, with different average degrees: 2, 5, and 10. We used 2 traversals for GRAIL when the average degree is 2, and 5 traversals otherwise. Table 15 shows the construction time and index sizes. Once again, none of the indexing methods other than GRAIL and GRIPP could handle these large graphs. PathTree too aborted on all datasets, except for *rand10m2x* with avg. degree 2; it took 526004.7ms for construction, and its index size was 69378979 entries). The table also includes the performance of GRAIL*, the basic approach without any optimizations. We see that GRIPP construction time and index size is smaller for dense graphs. Because its construction time is linear on the number of edges and GRIPP’s index size is equal to twice the number of edges. However GRIPP’s construction is again an order of magnitude slower for sparser graphs. (the reason is discussed above) Looking at the query times (Table 16), GRAIL is orders of magnitude faster than GRIPP. In fact, when querying on *rand10m10x*, GRIPP exceeded the 20M ms time limit. Note that DFS also exceeded the time limit in that dataset. We can see that for these datasets GRAIL can be orders of magnitude faster than search-based methods. However for very dense graphs such as *rand10m10x* BBFS gets closer to GRAIL. PathTree ran only on *rand10m2x*, with query time 211.7 ms, where GRAIL took 174.8 ms. It is interesting to note that for these large, and sparse synthetic graphs, the positive cut and level filters can slightly slow down the query times for GRAIL, since the basic GRAIL* approach can be about 5-7% faster. Once again, we conclude that GRAIL is the most scalable reachability index for large

graphs, especially with increasing density. The fastest index PathTree does not scale to large graphs and the scalable index GRIPP cannot provide fast querying with increasing density. Indeed even pure search methods with level filters are preferable to them.

Synthetic Scale-Free Graphs: Additionally, we generated some random scale-free graphs using Albert-Barabasi preferential attachment network growth model [5] to examine the behavior of the leading methods on graphs that have power-law degree distribution. We generated 3 graphs of 10,000 nodes with average degrees of 2, 5 and 15, and 2 graphs of 100,000 nodes with average degrees of 2 and 5.

This particular set of experiments were performed on a system that has four 2.5 Ghz processors with 4 GB memory. Table 17 shows the comparison between GRAIL, GRIPP and PathTree. GRIPP is still fast at indexing, but even for average degree 2, it provides the worst query time. It gets drastically worse as the density increases. PathTree provides the fastest querying once it manages to construct the index, however, it is not scalable since it suffers memory problems as the density increases (e.g., for 100,000 node graph with average degree 5). GRAIL is still fast at indexing and querying in all cases. These results generally conform to the results in the previous experiments.

6.3 Graph Search Strategies: Baseline Methods

We saw above that pure graph search methods like DFS, BFS, and bidirectional BFS can be very effective in answering reachability queries, especially for the smaller graphs. In this section, we conduct a detailed evaluation of these strategies – the baseline methods – with and without the level filter.

We show results on three sets of queries. As one set we use the same 100K random query pairs used in all of the results above. However, since the graphs are very sparse, the vast majority of these pairs are not reachable. As an alternative, we generated 100K reachable pairs by simulating a random walk (start from a randomly selected source node, choose a random child with 99% probability and proceed, or stop and report the node as target with 1% probability). Finally, we generated an additional set of 100K deep positive queries, where by deep we mean longer path lengths between the source and destination nodes in the query. The frequency distribution for the number of hops between source and target nodes for the positive queries and deep positive queries is plotted in Figure 9. We can see the random positive queries have hop lengths from 1-18, but the deep positive pairs start and end at longer hop lengths. For example, for *cit-patents*, the hop length ranges from 1 to 12, but for the deep positive queries the hop length ranges from 10 to 15.

We compare the query time performance of depth-first, breadth-first and bidirectional breadth-first approaches both with and without the topological level filter in Tables 18, 19 and 20. We use DFS, BFS, and BBFS for the three search strategies without the level filter, and DFS-L, BFS-L, and BBFS-L, with the level filter.

Comparing the query times on the random query pairs (Table 18), we can observe that the simple topological level filter is extremely effective in pruning. It improves the performance of all the search methods. On the small graphs, the improvement is typically 2-4 times, whereas on the large graphs the improvement can be up to a factor of 10. One can also observe that DFS-L is invariably the preferred method for the sparse graphs. However, there are a few exceptions, such as *amaze*, *kegg*, *xmark* and *arxiv*, where BBFS-L can be an order of magnitude faster. The distinction of these graphs are their smaller effective diameter. On the large graphs, BFS-L seems to be the preferred method, since it is either the fastest method (by a factor of 2 or higher) or is not very far from the best method. BBFS-L is a close second.

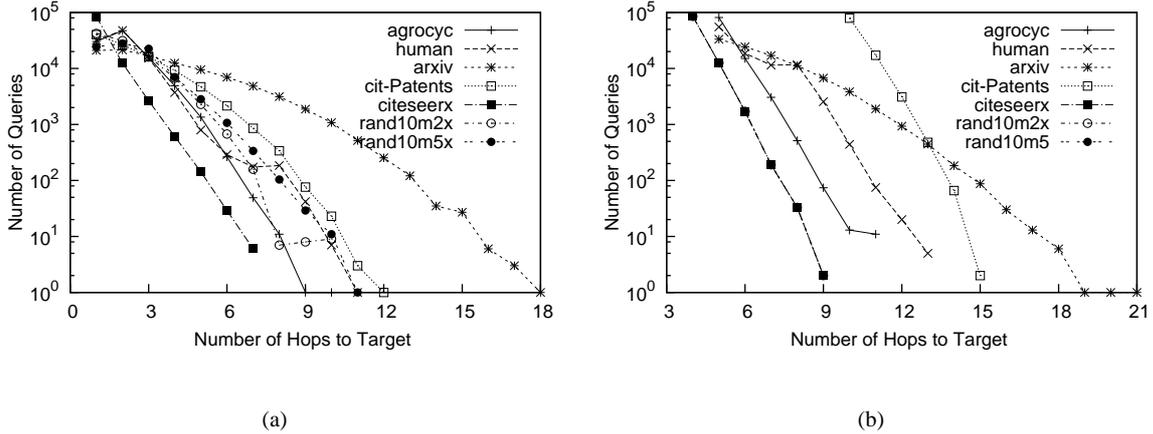
On the positive and deep positive queries (Tables 19 and 20) the effectiveness of the level filter is not very much. This is because of the fact that these are all reachable pairs of nodes, and cannot be filtered out. Thus the overhead of applying the level check consistently slows down the query times for all the three search methods, with some exceptions. On these (deep) positive queries, BBFS is the best overall method, being 2-100 times faster than alternative methods. The only exception is *go-uniprot*, where DFS/BFS does much better. This graph has around seven million roots which are connected to the remaining 30K nodes. Bidirectional search is slow in this graph because the reverse search that starts from the target node has to explore many irrelevant roots before finding a solution.

We can conclude that for small graphs DFS-L is the best overall method, whereas for the large graphs BFS-L is the best. However, BBFS is the clear winner if one expects more positive queries. It is for this reason that BBFS-L is the fastest on datasets like *kegg*, *amaze*, and *xmark* on the 100K random queries, since there are a relatively larger number of positive pairs in those query sets.

6.4 GRAIL: Effect of Parameters and Optimizations

In this section, we study the behavior of GRAIL under optimizations like the level and positive cut filter, as well as the choice of the graph search strategy (DFS/BFS/BBFS). We also evaluate the effectiveness of maintaining exceptions lists, and of the various traversal strategies to construct the interval labels.

| Dataset | | | Construction Time (ms) | | | Query Time (ms) | | | |
|---------|-----|-------|------------------------|--------|----------|-----------------|-----------|----------|----------|
| Nodes | Deg | PosQ | GRAIL | GRIPP | PathTree | GRAIL | GRIPP | PathTree | DFS |
| 10000 | 2 | 353 | 11.39 | 6.49 | 281.11 | 23.11 | 4026.49 | 9.14 | 209.35 |
| 10000 | 5 | 4287 | 16.22 | 12.13 | 3537.25 | 117.24 | 561482.62 | 34.16 | 2561.99 |
| 10000 | 15 | 21222 | 48.41 | 34.14 | 38716.01 | 658.91 | -(t) | 67.91 | 16251.73 |
| 100000 | 2 | 94 | 169.86 | 81.45 | 11581.34 | 46.86 | 16529.73 | 39.48 | 521.98 |
| 100000 | 5 | 1842 | 206.82 | 156.96 | -(m) | 511.88 | -(t) | -(m) | 12422.07 |

Table 17 Synthetic Scale-Free Graphs

Fig. 9 Reachability Queries: (a) Positive (b) Deep Positive Distance

6.4.1 Effect of Optimizations

We first consider the effect of the filters in combination with the recursive graph search strategy on the random, positive and deep positive query pairs. We compared eight variations of GRAIL, which are shown in tables 21, 22 and 23. Here GRAIL is the version which uses positive cut filter, whereas GRAIL* is the version without this filter. The suffix of LF is used when level filter is used, and BI is used when a bidirectional search is performed instead of DFS. Thus GRAIL is the approach that uses positive cut and the level filter, whereas GRAIL-LF is with level filter removed, GRAIL+BI uses bidirectional BFS instead of DFS, and GRAIL+BILF uses both BBFS and level filter. Similar extensions for GRAIL* show the variants of the basic approach, which uses no optimizations, and GRAIL*+LF uses the level filter. The column d represents how many traversals are used to index that graph.

It is clear that the positive cut filter is extremely effective, since GRAIL variants are invariably superior to the GRAIL* counterparts, by up to 100 times in some cases. The positive cut filter is very helpful for positive queries since if at any point of the search a node contains the guaranteed interval, it is sufficient to terminate the query with positive answer. However, it is also helpful in sparse graphs because most of the reachable pairs can be covered with the interior label set. In some exceptional cases GRAIL* is better than GRAIL, when the distance from the source to target

nodes is high as for `cit-patents` and `rand10m5x` in table 23. Due to the density of these datasets positive cut filter only works after some level until which the search compares interior intervals in vain. For instance, for `rand10m5x` the deep positive query lengths range from 10 to 16, whereas the search utilizes the positive cuts only after depth 5 on average, and thus up to depth 5 the search makes twice the number of comparisons. Even in these cases it is worth noting that the performance difference between GRAIL and GRAIL*+LF is not high.

We also observe that bidirectional search is not effective when used with GRAIL variants because the overhead of maintaining two queues and comparing twice the number of intervals at every step adds extra cost, which is not offset by additional pruning. We conclude that GRAIL (with positive cut and level filters, and DFS search) is the preferred method.

6.4.2 Exception Lists

Table 24 shows the effect of using exception lists. Here GRAIL is the default GRAIL strategy with reversed pairs traversals, level and positive cut filters, and without exception lists, GRAIL* is the basic approach without any optimizations, and GRAIL*+E is the basic approach with exception lists. Results are shown only for small, real, sparse and dense graphs, since computing the exception lists on the large graphs is too expensive. We used $d = 2$ for sparse and $d = 3$ for dense graphs.

| Dataset | DFS | DFS-L | BFS | BFS-L | BBFS | BBFS-L |
|-------------|--------------|---------------|-----------|-----------------|----------|----------------|
| agrocyc | 43.68 | 33.06 | 105.23 | 76.35 | 141.32 | 138.65 |
| amaze | 1737.61 | 692.88 | 1844.58 | 406.93 | 283.17 | 191.74 |
| anthra | 38.79 | 29.83 | 100.68 | 71.65 | 139.55 | 139.62 |
| ecoo | 50.29 | 33.83 | 110.59 | 78.01 | 142.37 | 141.85 |
| human | 34.09 | 32.43 | 93.27 | 80.39 | 156.31 | 164.86 |
| kegg | 2150.63 | 1085.38 | 2049.80 | 606.48 | 388.31 | 267.29 |
| mtbrv | 54.70 | 33.32 | 117.74 | 77.56 | 126.74 | 125.47 |
| nasa | 131.12 | 93.12 | 190.31 | 85.47 | 142.67 | 132.49 |
| vchocyc | 55.43 | 33.80 | 104.10 | 64.50 | 126.68 | 124.34 |
| xmark | 388.60 | 324.93 | 370.84 | 214.58 | 204.89 | 158.41 |
| arxiv | 13237.37 | 6816.73 | 10635.92 | 1513.44 | 3491.59 | 664.77 |
| go | 121.75 | 106.21 | 161.35 | 121.39 | 179.10 | 164.64 |
| pubmed | 323.84 | 207.70 | 443.40 | 94.84 | 202.40 | 188.04 |
| yago | 116.43 | 29.52 | 133.26 | 90.70 | 152.03 | 162.44 |
| cit-Patents | 38616.12 | 30317.47 | 40983.66 | 3403.68 | 7997.03 | 5045.11 |
| citeseer | 46.20 | 25.97 | 101.31 | 87.58 | 139.00 | 153.54 |
| citeseerx | 224954.60 | 153682.93 | 259932.93 | 4428.25 | 6098.53 | 4218.07 |
| go-uniprot | 37.87 | 156.11 | 268.58 | 372.05 | 320.92 | 235.17 |
| rand10m2x | 553.33 | 391.02 | 625.60 | 266.27 | 537.53 | 477.06 |
| rand10m5x | 91262.17 | 74963.43 | 96374.10 | 12598.14 | 29365.89 | 17110.69 |

Table 18 Query Time Comparison of the Baseline Methods: Random Queries

| Dataset | DFS | DFS-L | BFS | BFS-L | BBFS | BBFS-L |
|-------------|---------------|----------|---------|---------|----------------|------------|
| agrocyc | 392.27 | 456.49 | 226.55 | 296.55 | 150.49 | 176.07 |
| human | 924.84 | 933.10 | 323.98 | 528.19 | 182.20 | 255.21 |
| arxiv | 1390.69 | 1399.79 | 858.13 | 1006.30 | 240.03 | 296.98 |
| citeseerx | 79843.72 | 80639.62 | 2133.04 | 3160.10 | 147.55 | 183.47 |
| cit-Patents | 6851.40 | 6859.36 | 2852.23 | 3623.24 | 605.37 | 944.77 |
| go-uniprot | 151.69 | 168.31 | 221.77 | 243.11 | 86679.75 | 2058323.75 |
| rand10m2x | 430.95 | 461.72 | 299.84 | 369.90 | 292.26 | 384.79 |
| rand10m5x | 49874.03 | 50063.26 | 3570.48 | 4738.75 | 1106.92 | 1627.96 |

Table 19 Query Time Comparison of the Baseline Methods - Positive Queries

| Dataset | DFS | DFS-L | BFS | BFS-L | BBFS | BBFS-L |
|-------------|----------------|------------|-----------|-----------|-----------------|---------------|
| agrocyc | 699.56 | 710.95 | 1306.06 | 1785.99 | 299.60 | 387.30 |
| human | 1359.66 | 1263.17 | 2105.32 | 3164.45 | 357.66 | 309.99 |
| arxiv | 1084.50 | 1010.98 | 2188.39 | 2348.89 | 380.77 | 503.91 |
| citeseerx | 1510195.37 | 1481738.00 | 375087.90 | 471321.15 | 9488.83 | 17209.19 |
| cit-Patents | 50584.21 | 49820.53 | 96294.84 | 128953.70 | 11692.85 | 18848.20 |
| go-uniprot | 173.61 | 188.34 | 256.48 | 290.80 | 184129.15 | 3214053.00 |
| rand10m2x | 1069.12 | 1120.13 | 1413.06 | 1801.97 | 1124.22 | 1512.55 |
| rand10m5x | 174483.20 | 174269.43 | 213657.48 | 268942.18 | 62386.87 | 81867.81 |

Table 20 Query Time Comparison of the Baseline Methods - Deep Positive Queries

On both the sparse and dense graphs the construction time for exception lists blows up; GRAIL*+E is 2-3 orders of magnitude slower than GRAIL/GRAIL*. It is interesting to note that for the sparse graphs, the number of exceptions is not very large. In fact, the total index size for GRAIL*+E (which is equal to the index size of GRAIL* plus the number of exceptions) can be smaller than the index size for GRAIL (which has to keep additional information for the positive cut and level filters). However, computing the exception lists is very expensive. On the dense graphs, the number of exceptions blows up by over a factor of 100, and thus the construction time for GRAIL*+E increases even more rapidly. In terms of query time, for the sparse graphs, it is interesting to note that GRAIL is faster than GRAIL*+E by a factor of 2. However, GRAIL*+E is faster than the basic GRAIL*

method (sometimes by a factor of 40). This means that using the positive cut and level filters is much more effective than using exception lists, though exception lists can deliver better performance than DFS search. On the other hand, on the dense graphs, the exception lists lead to faster query times, even when compared to positive cut and level filters. For example, on arXiv GRAIL*+E is faster than GRAIL by a factor of 6. However, this comes at the cost of 3 orders of magnitude slowdown in construction times. We can conclude that whereas using exceptions does help in some cases, the substantially higher overhead of construction time, and the large size overhead of storing exception lists, do not justify the relative small gains in query times. Furthermore, exceptions could not be constructed on the large real graphs.

| Dataset | d | GRAIL* | GRAIL*+LF | GRAIL*+BI | GRAIL*+BILF | GRAIL-LF | GRAIL | GRAIL+BI | GRAIL+BILF |
|-------------|-----|---------|----------------|-----------|-------------|----------|----------------|----------|------------|
| agrocyc | 2 | 66.76 | 66.08 | 151.75 | 151.62 | 51.38 | 21.99 | 151.39 | 150.05 |
| amaze | 2 | 818.53 | 834.61 | 340.41 | 318.06 | 25.17 | 15.81 | 135.78 | 136.08 |
| anthra | 2 | 62.50 | 62.23 | 148.21 | 146.85 | 49.17 | 22.37 | 153.76 | 154.35 |
| ecoo | 2 | 68.52 | 64.44 | 175.18 | 158.57 | 54.63 | 22.95 | 156.08 | 154.04 |
| human | 2 | 86.15 | 84.44 | 160.19 | 161.84 | 64.03 | 23.45 | 161.95 | 177.25 |
| kegg | 2 | 1105.01 | 1079.61 | 503.31 | 496.06 | 25.60 | 17.57 | 150.93 | 159.23 |
| mtbrv | 2 | 61.40 | 58.80 | 151.47 | 151.52 | 47.42 | 20.40 | 148.28 | 165.92 |
| nasa | 2 | 37.75 | 35.68 | 176.42 | 146.64 | 35.70 | 21.98 | 147.30 | 146.45 |
| vchocyc | 2 | 60.00 | 57.74 | 156.44 | 155.75 | 46.43 | 19.85 | 138.99 | 137.54 |
| xmark | 2 | 97.61 | 88.75 | 180.60 | 176.62 | 42.46 | 28.04 | 148.17 | 165.71 |
| arxiv | 4 | 457.89 | 436.09 | 1149.78 | 1126.67 | 369.36 | 331.72 | 732.75 | 718.59 |
| go | 3 | 54.55 | 50.27 | 159.23 | 142.51 | 52.65 | 34.11 | 141.77 | 144.07 |
| pubmed | 3 | 89.56 | 85.77 | 216.76 | 215.18 | 94.07 | 66.23 | 219.38 | 219.36 |
| yago | 3 | 50.29 | 45.62 | 216.07 | 213.03 | 53.95 | 19.10 | 221.40 | 233.04 |
| cit-Patents | 5 | 1425.87 | 1363.89 | 1995.60 | 1982.56 | 1537.45 | 1412.75 | 2021.37 | 1999.32 |
| citeseer | 2 | 99.99 | 93.29 | 224.15 | 225.36 | 105.15 | 26.30 | 248.23 | 245.85 |
| citeseerx | 5 | 8173.68 | 8259.32 | 2583.46 | 2381.76 | 193.41 | 115.42 | 1044.88 | 1044.86 |
| go-uniprot | 2 | 163.10 | 131.59 | 465.90 | 507.32 | 171.82 | 29.63 | 489.76 | 469.97 |
| rand10m2x | 3 | 222.25 | 210.52 | 338.70 | 347.04 | 234.54 | 155.93 | 350.90 | 344.73 |
| rand10m5x | 5 | 6479.56 | 6322.14 | 7914.99 | 6969.14 | 6196.16 | 5990.32 | 7446.50 | 7376.99 |

Table 21 Query Time Comparison of the GRAIL methods

| Dataset | d | GRAIL* | GRAIL*+LF | GRAIL*+BI | GRAIL*+BILF | GRAIL-LF | GRAIL | GRAIL+BI | GRAIL+BILF |
|-------------|-----|---------------|-----------|-----------|-------------|---------------|---------------|----------------|------------|
| agrocyc | 2 | 479.69 | 561.59 | 325.44 | 330.82 | 36.47 | 38.55 | 150.99 | 159.66 |
| human | 2 | 1377.98 | 1163.86 | 745.24 | 748.22 | 35.84 | 37.39 | 146.89 | 148.35 |
| arxiv | 4 | 321.82 | 330.89 | 1070.86 | 1222.06 | 253.14 | 253.09 | 816.84 | 826.17 |
| cit-Patents | 5 | 3109.31 | 3303.92 | 3037.28 | 2694.92 | 3195.76 | 3089.42 | 2452.24 | 2475.95 |
| citeseerx | 5 | 27400.50 | 26402.77 | 327.15 | 317.41 | 123.55 | 117.37 | 232.21 | 234.51 |
| go-uniprot | 2 | 262.75 | 275.12 | -(t) | -(t) | 319.57 | 327.29 | -(t) | -(t) |
| rand10m2x | 3 | 550.03 | 549.80 | 775.11 | 842.06 | 546.64 | 522.34 | 685.14 | 698.08 |
| rand10m5x | 5 | 17997.88 | 17555.28 | 3417.46 | 3840.27 | 19593.36 | 19371.47 | 3152.93 | 3175.09 |

Table 22 Query Time Comparison of the GRAIL methods with Positive queries

| Dataset | d | GRAIL* | GRAIL*+LF | GRAIL*+BI | GRAIL*+BILF | GRAIL-LF | GRAIL | GRAIL+BI | GRAIL+BILF |
|-------------|-----|---------------|-----------------|-----------|-------------|---------------|----------------|----------|------------|
| agrocyc | 2 | 492.98 | 512.95 | 1001.90 | 840.36 | 66.32 | 69.89 | 459.63 | 467.71 |
| human | 2 | 670.90 | 709.45 | 547.53 | 508.68 | 49.64 | 50.91 | 242.98 | 247.03 |
| arxiv | 4 | 371.42 | 376.47 | 2484.63 | 2513.01 | 282.28 | 292.93 | 1512.91 | 1488.77 |
| cit-Patents | 5 | 30969.72 | 30174.54 | 46411.89 | 45814.12 | 33551.28 | 33093.57 | 43740.33 | 44198.10 |
| citeseerx | 5 | 1131971.50 | 1090715.25 | 63204.19 | 62324.88 | 3711.59 | 3529.66 | 15287.94 | 15432.34 |
| go-uniprot | 2 | 290.22 | 301.95 | -(t) | -(t) | 349.56 | 361.48 | -(t) | -(t) |
| rand10m2x | 3 | 1294.02 | 1323.16 | 2363.52 | 2343.72 | 1195.18 | 1177.69 | 1914.48 | 1951.57 |
| rand10m5x | 5 | 64076.53 | 63784.27 | 90399.75 | 80405.13 | 68951.63 | 70447.32 | 74483.97 | 75279.82 |

Table 23 Query Time Comparison of the GRAIL methods with Deep Positive queries

6.4.3 Label Traversal Strategies

We implemented different traversal strategies for interval labeling, as explained in Section 3.2. Random is the default randomized strategy, whereas ReversePairs denotes the randomized pairs strategy. *MaxVol*, *MaxInt*, *MaxAdjVol* and *MaxAdjInt* denote the deterministic methods that prioritize the nodes that have larger volume, minimum interval, adjusted volume, and minimum adjusted interval, respectively. For the adjusted methods, we implemented the transitive closure size estimation algorithm in [12] to get the approximate size of the reachable set for each node. In these experiments we used GRAIL*.

As seen in Table 25, random labeling has the worst performance. Though the differences among the methods are

not substantial, they are more pronounced for the large graphs. On the small sparse datasets deterministic methods provide up to %20 improvement over the randomized labeling. MaxVol gives best results most of the time, closely followed by ReversePairs. We use ReversePairs as the default strategy for GRAIL, since it is usually better for large graphs, and does not need the estimation of the transitive closure.

In Table 26, we take a closer look at the impact of traversal strategy by comparing the number of exceptions remaining after performing the traversals. It is interesting to note that Random labeling produces significantly more exceptions compared to other labellings whereas the difference is not that apparent in query time comparison. It also suggests that the query time does not necessarily have to be directly correlated with the number of exceptions remained. For in-

| Dataset | Construction Time (ms) | | | Index Size | | | | Query Time (ms) | | |
|----------|------------------------|--------|-----------|------------|--------|----------|-------------|-----------------|---------|--------------|
| | GRAIL | GRAIL* | GRAIL*+E | GRAIL | GRAIL* | GRAIL*+E | #Exceptions | GRAIL | GRAIL* | GRAIL*+E |
| agrocyc | 18.37 | 18.33 | 7615.42 | 88788 | 50736 | 79378 | 28642 | 20.81 | 72.10 | 53.10 |
| amaze | 5.72 | 5.72 | 2221.14 | 25970 | 14840 | 45720 | 30880 | 16.91 | 867.85 | 25.36 |
| anthra | 17.15 | 17.17 | 14521.61 | 87493 | 49996 | 65320 | 15324 | 20.62 | 61.22 | 52.87 |
| ecoo | 20.27 | 20.36 | 7865.13 | 88340 | 50480 | 79576 | 29096 | 20.89 | 64.38 | 53.86 |
| human | 52.75 | 52.71 | 15756.67 | 271677 | 155244 | 176663 | 21419 | 22.90 | 81.84 | 69.88 |
| kegg | 5.89 | 5.87 | 2409.59 | 25319 | 14468 | 56717 | 42249 | 18.84 | 1137.29 | 25.81 |
| mtbrv | 14.53 | 14.60 | 4764.51 | 67214 | 38408 | 57085 | 18677 | 18.39 | 64.99 | 46.90 |
| nasa | 9.31 | 9.38 | 2952.96 | 39235 | 22420 | 191771 | 169351 | 20.16 | 35.73 | 30.95 |
| vchocyc | 15.12 | 15.29 | 4790.82 | 66437 | 37964 | 54234 | 16270 | 17.91 | 56.55 | 46.80 |
| xmark | 9.87 | 9.83 | 7134.75 | 42560 | 24320 | 1368326 | 1344006 | 30.30 | 99.72 | 37.71 |
| arXiv | 29.71 | 29.68 | 57595.03 | 60000 | 36000 | 4203463 | 4167463 | 380.94 | 485.79 | 57.56 |
| citeseer | 46.82 | 47.19 | 169422.86 | 107200 | 64320 | 8399563 | 8335243 | 64.04 | 88.66 | 58.89 |
| go | 20.05 | 20.09 | 14653.48 | 67930 | 40758 | 656031 | 615273 | 30.79 | 47.69 | 39.39 |
| pubmed | 35.05 | 34.95 | 94894.25 | 90000 | 54000 | 5980867 | 5926867 | 70.07 | 92.96 | 54.92 |
| yago | 24.60 | 24.53 | 56401.45 | 66420 | 39852 | 2179247 | 2139395 | 19.64 | 49.51 | 46.08 |

Table 24 GRAIL: Effect of Exceptions

| Dataset | d | Random | ReversePairs | MaxVol | MaxInt | MaxAdjVol | MaxAdjInt |
|-------------|-----|---------|----------------|----------------|---------------|---------------|---------------|
| agrocyc | 2 | 78.39 | 72.33 | 65.45 | 65.53 | 64.92 | 64.93 |
| anthra | 2 | 66.44 | 60.98 | 65.55 | 65.53 | 65.07 | 65.67 |
| amaze | 2 | 838.09 | 872.46 | 726.80 | 713.04 | 726.26 | 717.27 |
| mtbrv | 2 | 73.07 | 64.57 | 62.24 | 62.19 | 62.16 | 62.13 |
| arxiv | 3 | 535.05 | 490.00 | 595.52 | 568.92 | 673.33 | 589.49 |
| arxiv | 4 | 477.71 | 417.76 | 495.20 | 452.59 | 598.03 | 497.92 |
| pubmed | 3 | 98.15 | 92.20 | 86.52 | 88.90 | 86.07 | 86.91 |
| pubmed | 4 | 89.33 | 86.32 | 79.68 | 82.36 | 80.64 | 82.48 |
| yago | 3 | 56.25 | 49.33 | 46.66 | 47.15 | 46.98 | 46.82 |
| yago | 4 | 53.01 | 48.32 | 45.80 | 46.35 | 45.86 | 46.38 |
| go | 3 | 60.93 | 46.99 | 45.25 | 45.62 | 48.86 | 49.20 |
| go | 4 | 48.19 | 45.25 | 43.50 | 43.26 | 46.61 | 46.45 |
| citeseerx | 4 | 8905.04 | 7797.91 | 7848.54 | 7850.85 | 8142.71 | 8093.69 |
| citeseerx | 5 | 8404.13 | 8148.96 | 7374.71 | 7455.17 | 7942.13 | 7680.86 |
| cit-Patents | 4 | 1694.08 | 1626.33 | 1748.10 | 1926.49 | 1890.82 | 1915.62 |
| cit-Patents | 5 | 1458.75 | 1350.29 | 1387.13 | 1620.15 | 1593.68 | 1617.17 |
| go-uniprot | 4 | 158.72 | 137.59 | 136.71 | 136.52 | 136.65 | 135.37 |
| go-uniprot | 5 | 149.08 | 141.82 | 135.72 | 135.62 | 135.89 | 135.50 |

Table 25 Comparison of different traversal strategies in GRAIL: Query Times (ms)

stance, in Table 26 ReversePairs have about 20% more exceptions than MaxAdjInt for arxiv graph, however its query time is still 20% faster than MaxAdjInts as seen in Table 25.

6.4.4 Number of Traversals/Intervals (d)

In Figure 10 we plot the effect of increasing the dimensionality of the index, i.e., increasing the number of traversals d , on three small sparse (agrocyc, amaze, ecoo), two small dense (arxiv, pubmed), two large real (cit-patents, citeseerx), and two large synthetic (rand10m2x, rand10m5x) graphs. In each plot we show the query times of GRAIL and GRAIL*. Since GRAIL is usually much faster than GRAIL*, we had to use two different y-axis for some of the results. The y-axis on the left hand side shows the query time of GRAIL while the one on the right side displays the query time for GRAIL*. Note that the plots for GRAIL and GRAIL* appear similar in most of the graphs even though their query times are in different scales. For instance, in amaze GRAIL

rises to 26 ms from 19 ms whereas GRAIL* rises to 940 ms from 840 ms in a very similar manner. It is clear that increasing the number of intervals increases the construction time and index size, while decreasing the query time. However increasing d does not progressively decrease query times, since at some point the overhead of checking a larger number of intervals negates the potential reduction in exceptions. In small sparse graphs the minimum query time is obtained at some point between 2 and 4. However we decided to use $d = 2$ traversals for these graphs since adding each dimension increases the index size significantly while the gain in query time is comparatively very small. As the graphs get denser, the optimum number of traversals increases. For example, for arxiv we get the minimum query times between 12 to 16, and adding more traversals degrades query performance. Considering this graph has an average degree of 11.12, using 10 traversals seems acceptable. That makes the index size close to the order of the graph size and the query

| Dataset | d | Random | ReversePairs | MaxVol | MaxInt | MaxAdjVol | MaxAdjInt |
|----------|-----|----------|---------------|----------------|---------------|----------------|----------------|
| agrocyc | 2 | 1514319 | 28642 | 29943 | 29943 | 29672 | 29672 |
| amaze | 2 | 685366 | 30880 | 108472 | 108472 | 41280 | 41280 |
| anthra | 2 | 78560 | 15324 | 14874 | 14874 | 14436 | 14436 |
| ecoo | 2 | 2701160 | 29096 | 28475 | 28475 | 27576 | 27576 |
| human | 2 | 2003874 | 21419 | 23756 | 23746 | 21720 | 21720 |
| kegg | 2 | 651196 | 42249 | 100449 | 100449 | 36604 | 36604 |
| mtbrv | 2 | 1612124 | 18677 | 20240 | 20240 | 19734 | 19734 |
| nasa | 2 | 674175 | 169351 | 225161 | 225161 | 208253 | 208253 |
| vchocyc | 2 | 961784 | 16270 | 16131 | 16131 | 15759 | 15759 |
| xmark | 2 | 2298605 | 1344006 | 1376535 | 1376535 | 1233041 | 1233041 |
| arxiv | 3 | 4944155 | 4167463 | 3387488 | 3452679 | 3251799 | 3174875 |
| citeseer | 3 | 12645509 | 8335243 | 6560907 | 6777400 | 6547945 | 6801282 |
| go | 3 | 1279303 | 615273 | 640928 | 608102 | 663908 | 662851 |
| pubmed | 3 | 8255165 | 5926867 | 4643569 | 4768787 | 4621694 | 4714258 |
| yago | 3 | 4729341 | 2139395 | 1054214 | 1198431 | 1055417 | 1198552 |

Table 26 Comparison of different traversal strategies in GRAIL: Number of Exceptions Remained

time close to the minimum. However we used much smaller d in our experiments to be able to compete with other algorithms in terms of index size as well. Therefore one can choose the value d according to the needs by balancing the index size and query time. On very large sparse graphs such as `citeseerx` and `rand10m2x` the query times approach to the minimum level after 4 traversals. However for the denser large graphs we see that the query times continue to decrease up to 25 traversals. It is also interesting to see that for these graphs positive cut filter does not improve query performance and GRAIL has almost the same timing with GRAIL*. For `cit-Patents` we obtained 1369 ms with 5 traversals and this plot suggests that it can be improved to 650 ms by increasing its index size by 4 times. Therefore even in these cases it may not worth to keep adding traversals unless the memory is not a constraint or the query performance is critical.

Consequently, estimating the number of traversals that minimize the query time, or that optimize the index size/query time trade-off, is not straightforward. However, for any practical benefit it is imperative to keep the index size smaller than the graph size. This loose constraint restricts d to be less than the average degree. In our experiments, we found out that the best query time is obtained when $d = 5$ or smaller (when the average degree is smaller).

7 Conclusion

We proposed GRAIL, a relatively simple indexing scheme, for fast and scalable reachability testing in very large graphs, based on randomized multiple interval labeling. GRAIL has linear construction time and index size, and its query time ranges from constant to linear time per query. Based on an extensive set of experiments, we conclude that for the class of smaller graphs (both dense and sparse), while more sophisticated methods give a better query time performance, a

DFS/BFS search is often good enough, with the added advantage of having no construction time or index size overhead. On the other hand, GRAIL outperforms all existing methods, as well as pure search-based methods on large real graphs; in fact, for these large graphs existing indexing methods are simply not able to scale. Although GRIPP is scalable in indexing, it is not able to compete even with pure search methods in denser graphs.

As a future work we plan to address reachability labeling in dynamic graphs. GRAIL seems promising for the problem of dynamic reachability since its labeling has relaxed invariants. The updates on the graph such as edge addition and deletion could be reflected on the labeling via local modifications to the labels. The reachability problem in labeled graphs is also an interesting direction.

References

1. <http://www.geneontology.org/>. Accessed 4 December, 2010.
2. Go database guide. http://www.geneontology.org/GO.database.shtml#schema_notes. Accessed 4 December, 2010.
3. Uniprot. <http://www.uniprot.org/>. Accessed 4 December, 2010.
4. R. Agrawal, A. Borgida, and H. V. Jagadish. Efficient management of transitive relationships in large data and knowledge bases. *SIGMOD Rec.*, 18(2):253–262, 1989.
5. A. L. Barabási and R. Albert. Emergence of scaling in random networks. *Science*, 286(5439):509–512, 1999.
6. P. Bouros, S. Skiadopoulos, T. Dalamagas, D. Sacharidis, and T. Sellis. Evaluating reachability queries over path collections. In *SSDBM*, page 416, 2009.
7. R. Bramandia, B. Choi, and W. K. Ng. On incremental maintenance of 2-hop labeling of graphs. In *WWW*, 2008.
8. L. Chen, A. Gupta, and M. E. Kurul. Stack-based algorithms for pattern matching on dags. In *VLDB*, 2005.
9. Y. Chen. General spanning trees and reachability query evaluation. In *Canadian Conference on Computer Science and Software Engineering*, Montreal, Quebec, Canada, 2009.
10. Y. Chen and Y. Chen. An efficient algorithm for answering graph reachability queries. In *ICDE*, 2008.

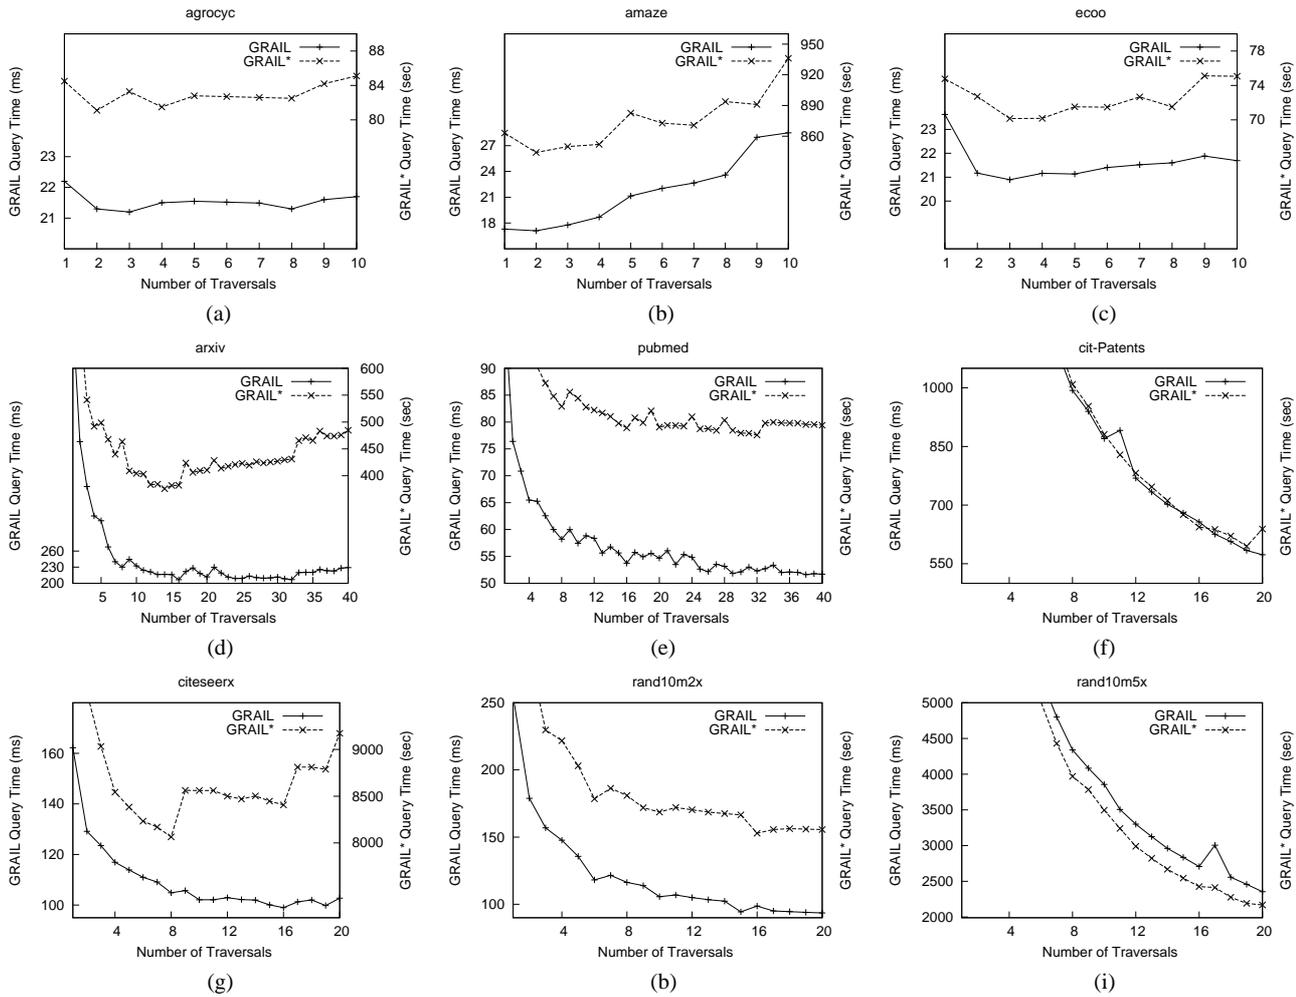


Fig. 10 Effect of Increasing Number of Intervals: (a) agrocyt, (b) amaze, (c) ecoo, (d) arxiv, (e) pubmed, (f) cit-patents, (g) citeseerx, (h) rand10m2x, (i) rand10m5x

- J. Cheng, J. X. Yu, X. Lin, H. Wang, and P. S. Yu. Fast computing reachability labelings for large graphs with high compression rate. In *EBDT*, 2008.
- E. Cohen. Estimating the size of the transitive closure in linear time. In *Proceedings of the 35th Annual Symposium on Foundations of Computer Science*, pages 190–200, 1994.
- E. Cohen. Size-estimation framework with applications to transitive closure and reachability. *Journal of Computer and System Sciences*, 55(3):441–453, 1997.
- E. Cohen, E. Halperin, H. Kaplan, and U. Zwick. Reachability and distance queries via 2-hop labels. *SIAM Journal of Computing*, 32(5):1335–1355, 2003.
- T. H. Cormen, C. E. Leiserson, R. L. Rivest, and C. Stein. *Introduction to Algorithms*. MIT Press, 2001.
- C. Demetrescu and G. Italiano. Fully Dynamic Transitive Closure: Breaking through the $O(n^2)$ Barrier. In *FOCS*, 2000.
- C. Demetrescu and G. Italiano. Dynamic shortest paths and transitive closure: Algorithmic techniques and data structures. *Journal of Discrete Algorithms*, 4(3):353–383, 2006.
- P. F. Dietz. Maintaining order in a linked list. In *STOC*, 1982.
- H. He, H. Wang, J. Yang, and P. S. Yu. Compact reachability labeling for graph-structured data. In *CIKM*, 2005.
- I. Herman. W3c semantic web faq. <http://www.w3.org/2001/sw/SW-FAQ#>. Accessed 4 December, 2010.
- H. V. Jagadish. A compression technique to materialize transitive closure. *ACM Trans. Database Syst.*, 15(4):558–598, 1990.
- R. Jin, Y. Xiang, N. Ruan, and D. Fuhr. 3-hop: a high-compression indexing scheme for reachability query. In *SIGMOD*, 2009.
- R. Jin, Y. Xiang, N. Ruan, and H. Wang. Efficient answering reachability queries on very large directed graphs. In *SIGMOD*, 2008.
- V. King and G. Sagert. A fully dynamic algorithm for maintaining the transitive closure. *J. Comput. Syst. Sci.*, 65(1):150–167, 2002.
- I. Krommidas and C. Zaroliagis. An experimental study of algorithms for fully dynamic transitive closure. *Journal of Experimental Algorithmics*, 12:16, 2008.
- J. Leskovec. Snap network analysis library. <http://snap.stanford.edu/snap/index.html>. Accessed 4 December, 2010.
- L. Roditty and U. Zwick. A fully dynamic reachability algorithm for directed graphs with an almost linear update time. In *STOC*, 2004.
- R. Schenkel, A. Theobald, and G. Weikum. HOPI: an efficient connection index for complex XML document collections. In *EBDT*, 2004.

29. R. Schenkel, A. Theobald, and G. Weikum. Efficient creation and incremental maintenance of the hopi index for complex xml document collections. In *ICDE*, 2005.
30. G. Steve Harris. Sparql 1.1 query language. <http://www.w3.org/TR/sparql11-query/#propertypaths>. Accessed 4 December, 2010.
31. S. Trissl and U. Leser. Fast and practical indexing and querying of very large graphs. In *SIGMOD*, 2007.
32. H. Wang, H. He, J. Yang, P. Yu, and J. X. Yu. Dual labeling: Answering graph reachability queries in constant time. In *ICDE*, 2006.
33. Z. Wu, G. Eadon, S. Das, E. I. Chong, V. Kolovski, M. Annamalai, and J. Srinivasan. Implementing an inference engine for rdfs/owl constructs and user-defined rules in oracle. In *ICDE*, pages 1239–1248, 2008.
34. H. Yildirim, V. Chaoji, and M. J. Zaki. Grail: Scalable reachability index for large graphs. *PVLDB*, 3(1):276–284, 2010.
35. J. X. Yu, X. Lin, H. Wang, P. S. Yu, and J. Cheng. Fast computation of reachability labeling for large graphs. In *EBDT*, 2006.